

---

# TWIN AUGMENTED ARCHITECTURES FOR ROBUST CLASSIFICATION OF COVID-19 CHEST X-RAY IMAGES

---

**Kartikeya Badola, Sameer Ambekar, Himanshu Pant, Sumit Soman**

Department of Electrical Engineering  
Indian Institute of Technology Delhi, India  
{kartikeya.badola,ambekarsameer,panthimanshu17,sumit.soman}@gmail.com

**Anuradha Sural**

Department of Radiology  
Max Hospital, Vaishali, India  
anuradha.sural@gmail.com

**Rajiv Narang**

Department of Cardiology  
All India Institute of Medical Sciences, New Delhi, India  
r\_narang@yahoo.com

**Suresh Chandra**

Department of Mathematics  
Indian Institute of Technology Delhi, India  
chandras@maths.iitd.ac.in

**Jayadeva**

Department of Electrical Engineering  
Indian Institute of Technology Delhi, India  
jayadeva@ee.iitd.ac.in

February 17, 2021

## ABSTRACT

The gold standard for COVID-19 is RT-PCR, testing facilities for which are limited and not always optimally distributed. Test results are delayed, which impacts treatment. Expert radiologists, one of whom is a co-author, are able to diagnose COVID-19 positivity from Chest X-Rays (CXR) and CT scans, that can facilitate timely treatment. Such diagnosis is particularly valuable in locations lacking radiologists with sufficient expertise and familiarity with COVID-19 patients. This paper has two contributions. One, we analyse literature on CXR based COVID-19 diagnosis. We show that popular choices of dataset selection suffer from data homogeneity, leading to misleading results. We compile and analyse a viable benchmark dataset from multiple existing heterogeneous sources. Such a benchmark is important for realistically testing models. Our second contribution relates to learning from imbalanced data. Datasets for COVID X-Ray classification face severe class imbalance, since most subjects are COVID -ve. Twin Support Vector Machines (Twin SVM) and Twin Neural Networks (Twin NN) have, in recent years, emerged as effective ways of handling skewed data. We introduce a state-of-the-art technique, termed as Twin Augmentation, for modifying popular pre-trained deep learning models. Twin Augmentation boosts the performance of a pre-trained deep neural network without requiring re-training. Experiments show, that across a multitude of classifiers, Twin Augmentation is very effective in boosting the performance of given pre-trained model for classification in imbalanced settings.

## 1 Introduction

COVID-19 is caused by respiratory virus SARS-CoV-2, and may be termed as a special kind of viral pneumonia. Subsequent to its first identification in Wuhan, China in December 2019, it has caused massive disruption worldwide. It has caused over 1,000,000 unfortunate deaths as of October, 2020. Current testing methods are slow, expensive and not widely accessible in many developing countries. The gold standard is a RT-PCR test that starts with collecting a nasal or oral swab. Test results are not immediate, and the turn-around time impacts treatment. Many pneumonic patients exhibiting COVID-19 symptoms test negative with RT-PCR. Chest X-Ray (CXR) and CT scans of such hospitalized

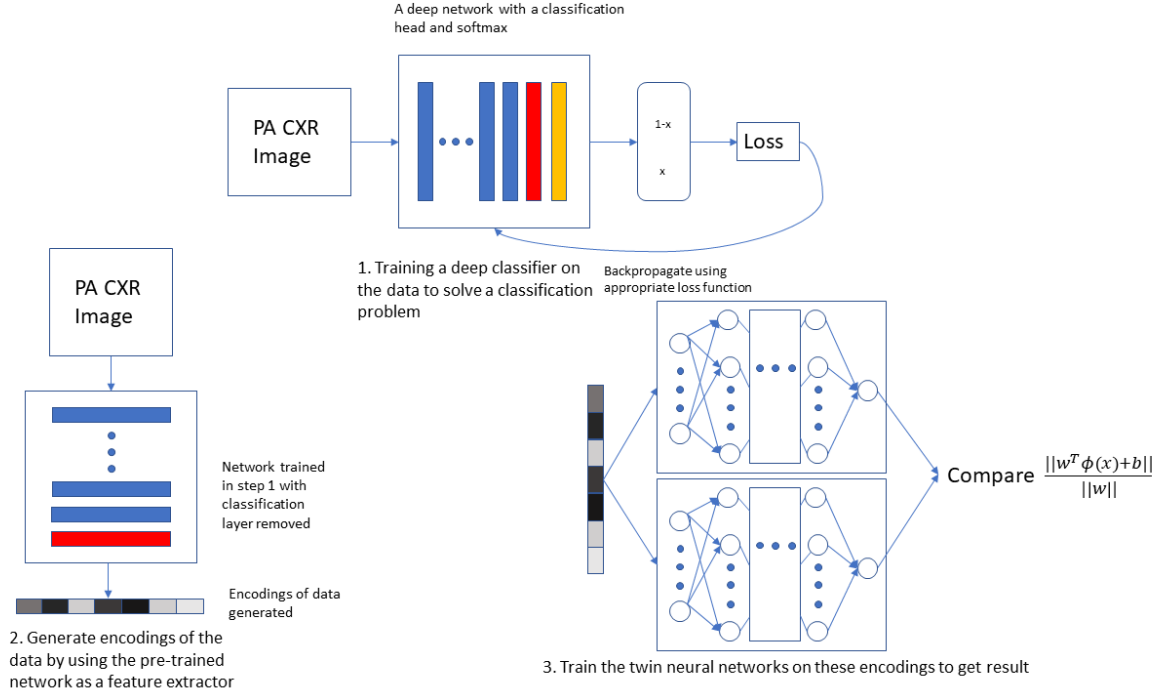


Figure 1: We follow a three step process for classifying chest X-Rays: 1. A deep classifier is trained on the dataset (pre-training). 2. We freeze the architecture and remove the final linear classification layer to obtain the encodings of the data after the penultimate layer 3. We initialize the Twin Neural Network which is then trained on these encodings. At test time, we process data using the truncated base network and use the Twin setup as the final layer (as opposed to a usual linear layer+softmax)

subjects often display features common among COVID-19 patients. Deviations in collection, storage, and transportation of samples have been cited as reasons for negative test results in such subjects. Inadequacy of the viral load in the test sample may also be a factor. Expert radiologists, including one co-author of this paper, have therefore suggested COVID-19 treatments for such patients on the basis of Chest X-Ray (CXR) or CT-scan images. Such a clinical input is particularly valuable in locations lacking expert radiologists or adequate RT-PCR testing facilities. This paper focusses on CXR based COVID-19 diagnosis using deep learning methods.

There have been several attempts at CXR based respiratory condition identification, such as by Irvin *et al.* [1], Cohen *et al.*[2]. Cohen *et al.*'s CXR dataset [3], that has been widely used for identifying COVID positivity, comprises only around 200 posteroanterior (PA) CXR images of COVID +ve subjects. It is built from varied sources, and has significant structural and visual diversity. The dataset also includes a smaller number of CXR images of COVID -ve subjects with bacterial, viral and fungal pneumonia. Owing to the small number of COVID -ve subjects in [3], negative samples are usually compiled from diverse sources. A careful analysis of these datasets reveals that some of them are very ill-suited for the task of identifying COVID positivity. This is mainly because the choice of negative samples leads to a dataset where the classes are easily separated. In fact, misleadingly high accuracies can be obtained.

In binary classification, class imbalance arises when one class (minority group) contains significantly fewer samples than the other class (majority group) [4]. In the current context, the number of COVID +ve subjects is much smaller than the number of COVID -ve ones. Classifiers trained on such data tend to over-fit on the majority class while erroneously classifying samples of the minority class [4]. Measures such as accuracy are misleading performance metrics, since they would be dominated by performance on the majority class.

Consider a binary classification dataset where the majority group is 99% of the total sample size, while the minority group forms 1% of the dataset. A naive learner trained on this imbalanced data, would achieve an accuracy of 99% by classifying all the samples as belonging to the majority group. This is invariably the case with severely imbalanced classification problems, with any state-of-the-art (SoTA) classifier employing naive losses, or no data augmentation. In such cases, performance metrics such as precision, recall and F1 are commonly used for imbalanced classification problems as they focus on the correct classification of minority class samples.

We present a new technique, called Twin Augmentation, to boost any neural network based classifier’s performance on imbalanced data, and report that this technique significantly boosts the accuracy, precision, recall and F1 of a variety of classifiers. In this approach, we replace the classification head of existing SoTA classifiers which are pre-trained on the given dataset using any training algorithm, with a Twin NN [5] block. The weights of the pre-trained architecture are frozen, and only the Twin NN block is trained (Fig. 1). We show, that despite being a simple extra step, Twin Augmentation consistently improves the performance of pre-trained networks by a significant margin. We also show, that this method is consistently robust with a variety of pre-trained classifiers, trained using a multitude of training algorithms. It involves minimal additional computation.

We compare our method with techniques like ADASYN [6], weighted cross-entropy loss and focal loss [7] which are widely used to solve the imbalance issue. We report that our method consistently outperforms these methods on the COVID dataset we compiled in this paper. The key contributions of this paper are as follows:-

- Analyze common sources used for COVID -ve samples, and compile a dataset that is diverse and tests generalizability.
- Introduce Twin Augmentation, that is a post-processing step for improving the performance of existing deep architectures. It is quick and easy to implement, requires minimal computation, and imparts robustness to the base model. It can work with any deep network trained with any training algorithm.

Subsequent sections are organized as follows. Section 2, discusses the literature on rectifying imbalance in classification tasks. In Section 3, we discuss prior work on CXR based COVID positivity prediction, and discuss choosing an appropriate dataset. In Section 4, we explain the motivation behind twin learning techniques such as the Twin SVM [8] and the Twin NN [5]), and introduce Twin Augmentation. Section 5 deals with experimental results. Section 6 contains concluding remarks.

## 2 Related Work

Skewed data distributions are widely found in applications where one of the classes is less common, e.g. data pertaining to disease diagnosis [9], fraud detection [10], [11], and image recognition [12]. Imbalance can be further be divided into intrinsic imbalance and extrinsic imbalance. Intrinsic imbalance occurs naturally, while extrinsic imbalance is caused through factors such as collection or storage procedures [4]. Approaches for handling class imbalance may be categorized as data level, algorithm based, and hybrid [4]. For our comparison, we focus on both data level methods and algorithmic methods.

### 2.1 Data level methods

Data level methods comprise data sampling methods, which can subsample the majority class or oversample the minority one. Oversampling may lead to the generation of spurious samples and increase classifier training time. SMOTE [13] is one such technique that interpolates using samples and their neighbours from the minority class. ADASYN [6] is SMOTE’s extension, and it creates more samples near the decision boundary. It often outperforms SMOTE on classification tasks, and is a useful comparison point.

### 2.2 Algorithmic methods

These techniques alter the learning process to increase the priority accorded to minority class samples. This may be done by assigning weights in the loss function for different classes. This can increase recall of the minority class but often at the cost of decreased precision. Other algorithmic methods employ a novel loss function. Focal loss [7] is a recent, widely cited technique, that reshapes cross entropy loss to reduce the importance of well-classified samples. This leads to a significant improvement on tasks such as classification and object detection [14]. In our experiments, we compare Twin Augmentation with weighted cross entropy and focal loss.

All the above methods rely on modifying the existing training process by sample addition, loss function modification, or both. Twin Augmentation, on the other hand, doesn’t need explicit re-training with new data, or any modification of the training process. We show that our method can boost the performance of any deep neural network, trained using any training algorithm, for the application at hand. This process is extremely fast to train and requires minimal GPU resources because of which time taken for hyperparameter tuning is also significantly reduced. Since Twin Augmentation is a post-processing step, it can benefit from any improvement in the base model or the training algorithm used to train the base model. This fact further expands the utility and longevity of the Twin Augmentation process in

the domain of deep learning since any improvement in classification task using deep neural networks can be further enhanced by using Twin Augmentation. Experimental results show that Twin Augmentation consistently outperforms other approaches on a variety of models for classification. For the task of classification, True Positive (TP) : positive data correctly classified as positive, False negative (FN) : positive data classified as negative, False Positive (FP) : negative data classified as positive and True Negative (TN) : negative data correctly classified as positive [15] F-measure is a performance measure that is computed based on precision and recall. It is used for imbalanced dataset, because it avoids using true positive, which tends to be extremely large in an imbalanced dataset [15]. Recall is the ratio where tp is the number of true positives and fn the number of false negatives. Precision is defined as the number of true positives over the number of true positives plus the number of false positives. F1 score is defined as the harmonic mean of precision and recall.

### 3 Analysis of Sources for COVID Negative Pneumonia Dataset




Figure 2: t-SNE plot of randomly selected datapoints from different sources: Circle-COVID +ve from Cohen’s repository, Cross-COVID -ve Pneumonia from Kermany’s dataset.

COVID-Net [16] has been widely cited for COVID-19 X-Ray classification. The arXiv version appeared in late March. Much of the literature followed the dataset used in COVID-Net. COVID-Net used Cohen’s repository [3, 17] for COVID +ve samples. Negative samples are from the pneumonia dataset by Kermany *et al.* (popularly known as Paul Timothy Mooney’s Kaggle Pneumonia dataset) [18]. It is important to note that in the current version of that paper (version 4 at the time of writing), the authors have completely changed the dataset they used originally. They have now compiled the dataset from a variety of sources. A recent review paper [19] surveys highly cited papers on CXR based COVID-19 diagnosis. Three out of four techniques mentioned in [19], viz. [20], [21, 16], use Kermany’s dataset for COVID -ve samples. Since these three papers are highly cited, many researchers have used the same datasets to carry out their experiments. For our study, we ran baselines on the same dataset (Cohen’s dataset for COVID +ve and Kermany’s dataset for COVID -ve) to diagnose COVID positivity from CXR images. We found that Mobilenet v2 [22] pretrained on ImageNet [23] yielded nearly 100% accuracy on this dataset. Closer examination using t-SNE reveals that samples in the Kermany’s dataset are significantly different from the samples in the Cohen’s dataset, and hence it is almost trivial to determine which source an image belongs to, by visual inspection alone. Fig. 2 contains a t-SNE plot [24] of samples from Kermany’s and Cohen’s datasets. It is evident that the two datasets form distinct clusters and are easily separable. Therefore, training for COVID positivity using their combination would be unwise. We highlight this issue by training a classifier on this dataset, and testing it on samples from other sources. This analysis raises the question of selecting an appropriate dataset for evaluating approaches to CXR based COVID-19 diagnosis. We therefore compiled pneumonic CXR images of COVID negative subjects from a diversity of sources, viz. CheXpert [1], Kermany [18], Cohen [3, 17], NIH [25] and Open-i [26]. Throughout this paper, we focus on the task of differentiating COVID +ve CXRs from pneumonic COVID -ve CXRs. Table 1 provides an overview.

Dataset Name	No. of images selected	Remarks
Kermany et al [18]	3875	All pneumonia images from the train set in Mooney’s Kaggle page
CheXpert [1]	991	Pneumonia images from downsampled version of CheXpert
NIH [25]	1431	Pneumonia images from NIH dataset. Also known as RSNA’s Pneumonia detection challenge dataset
Open-i [26]	68	Pneumonia images compiled from 5 different sources
Cohen’s repository COVID -ve [3, 17]	41	Contains fungal pneumonia, bacterial pneumonia and non-COVID viral pneumonia (12,17,12 during our experiments)

Table 1: An overview of COVID -ve pneumonia datasets used for our analysis.




Figure 3: A t-SNE plot of randomly selected datapoints from different sources: Circle-COVID +ve from Cohen’s repository, Cross-COVID -ve Pneumonia from CheXpert, Square-COVID -ve Pneumonia from Kermany’s dataset, Plus-COVID -ve Pneumonia from Cohen’s repository, Diamond-COVID -ve Pneumonia from NIH, Triangle-COVID -ve Pneumonia from Open-i.

Further, we examine t-SNE plots of these datasets. In Fig. 3 we can see that samples from CheXpert dataset [1] (denoted by crosses) and Kermany’s dataset [18] (denoted by squares) form clusters which are distant from positive samples (denoted by circles). This suggests that the images from these two datasets are very different from COVID +ve samples in Cohen’s dataset. A t-SNE plot of CheXpert [1] vs positive images from Cohen’s [3, 17] repository confirms our assertion (Fig. 4).

Therefore, we compiled COVID -ve CXR images from NIH, Open-i and Cohen’s datasets to construct the negative class. We took 70% of the samples for training, 10% for validation and 20% for testing, from each of these sources, after randomly shuffling the datasets. The same split was maintained for the positive class as well. In Fig. 5, we see that the t-SNE plot of our dataset does not exhibit the clear separability seen in Fig. 3. This indicates that a realistic choice of images from the general population yields a more challenging classification task. The final overview of our compiled dataset is given in Tables 2 and 3

We now validate our decision to eliminate the CheXpert and Kermany datasets. We took two classifiers pre-trained on ImageNet, and trained them on Cohen’s positive vs CheXpert and Cohen’s positive vs Kermany’s dataset. The loss function for each test is categorical cross-entropy, the optimizer is Adam [28] with learning rate=0.001, betas=0.9, 0.999, epsilon = 1e-08 and the batch size is 16. The models were trained for 100 epochs with early stopping (patience=20). We used PyTorch [29] for our experiments. For these experiments, the training and validation set of negative class from our compiled dataset were replaced with images from either Kermany’s dataset or CheXpert. For Kermany’s




Figure 4: A t-SNE plot of randomly selected datapoints from different sources: Circle-COVID +ve from Cohen’s repository, Cross-COVID -ve Pneumonia from CheXpert dataset.




Figure 5: t-SNE plot of randomly selected datapoints in the final dataset: Circle-COVID +ve from Cohen’s repository, Cross-COVID -ve Pneumonia from Cohen’s repository, Square-COVID -ve Pneumonia from NIH, Plus-COVID -ve Pneumonia from Open-i.

experiment, 3875 images were used for training, while 390 were used for validation of the negative class. In the CheXpert experiment, 699 images were used for training, and 99 for validation of the negative class. Once the model is trained, the weights corresponding to the lowest validation loss were taken, and the model was tested on the test set of our compiled dataset. The results of the experiments are summarized in Tables 4 and 5 .

It is evident that the trained models fare poorly on our test data. Recall is high, only because these models are basically classifying almost every image in the test set as COVID +ve. This is because Open-i and NIH images are close to the positive source. We also note, that CheXpert performed significantly worse than Kermany’s dataset, since

Dataset Name	Train/Val/Test
Cohen’s COVID +ve [3, 17]	105/15/31
Cohen’s COVID -ve Pneumonia [3, 17]	28/4/9
NIH COVID -ve Pneumonia [25]	1001/143/287
Open-i COVID -ve Pneumonia [26]	47/7/14

Table 2: An overview of our compiled dataset.

Set	Positive/Negative
Train	105/1076
Val	15/154
Test	31/310

Table 3: An overview of positive and negative class samples.

Model Name	Val Accuracy Kermamy’s vs Cohen’s Positive (15/390)	Test accuracy on our new dataset (31/310)	Precision on our dataset	Recall on our dataset	F1 on our dataset
Mobilenetv2 [22]	1.000	0.4076	0.1298	0.9677	0.2290
Resnet18 [27]	0.9951	0.2375	0.0982	0.9032	0.1772

Table 4: Train on Kermamy’s dataset and Cohen’s positive, test on our new compiled dataset.

Model Name	Val Accuracy CheXpert vs Cohen’s Positive (15/99)	Test accuracy on our new dataset (31/310)	Precision on our dataset	Recall on our dataset	F1 on our dataset
Mobilenetv2 [22]	1.000	0.0909	0.0909	1	0.1666
Resnet18 [27]	0.9912	0.1935	0.0986	0.9677	0.1791

Table 5: Train on CheXpert dataset and Cohen’s positive, test on our new compiled dataset.

Kermamy’s dataset has more images. It may be concluded that models trained on Kermamy or CheXpert are not useful for classifying COVID-19 X-Ray images. On compiling the optimal dataset, we notice that our dataset also faces severe class imbalance. The subsequent sections of our paper will focus on a new approach that is motivated by the Twin family of classifiers [8, 5]. We show how a simple step can demonstrably boost any classifier’s performance on an imbalanced dataset.

#### 4 The Twin Family of Classifiers and Twin Augmentation

For a binary classification problem, a Support Vector Machine (SVM) finds a single maximum margin hyperplane to separate the two classes [30],[31],[32],[33]. In a binary imbalanced setting, this formulation fails to generalize since the parameters of the single hyperplane are influenced by the samples in the majority class. The Twin Support Vector Machine (Twin SVM) [8] was formulated to solve this issue. Instead of solving for a single hyperplane, the Twin SVM solves two smaller Quadratic Programming Problems (QPPs) in order to find two non-parallel hyperplanes, one for each class. Each hyperplane is generated such that it is closer to samples of one of the two classes, and distant from samples of the other. This mitigates the problem caused due to class imbalance since each hyperplane is concerned with samples from one class only.

Let the training set  $X$  for a binary class problem be divided into samples of class 1 (denoted by  $X_1$ ), and those of class  $-1$  (denoted by  $X_{-1}$  resp.). The Twin SVM solves the following optimization problems ([5],[8]):

$$\min_{w_1, b_1, \xi} \frac{1}{2} \|X_1 w_1 + e_1 b_1\|_2^2 + C_1 e_{-1} \xi \quad (1)$$

subject to,

$$-(X_{-1} w_1 + e_{-1} b_1) + \xi \geq e_{-1}, \xi \geq 0 \quad (2)$$

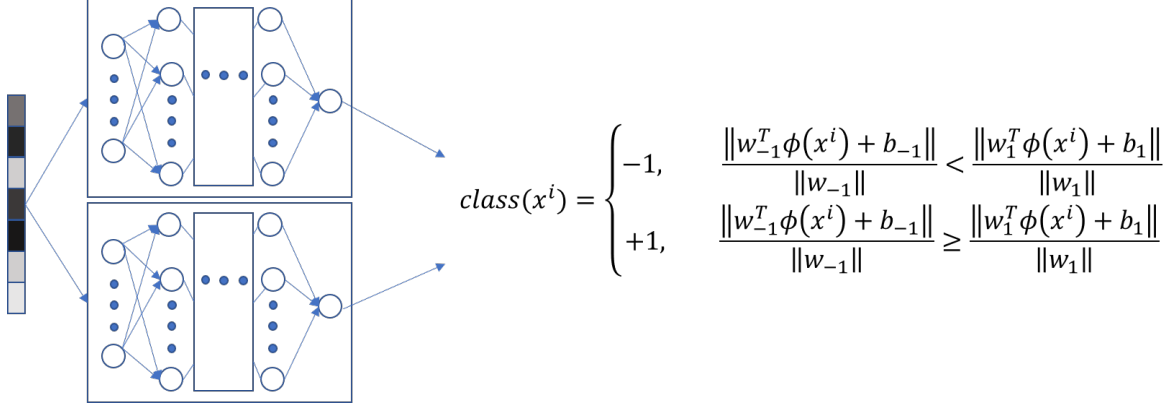


Figure 6: Diagrammatic representation of class inference using Twin Neural networks (Twin NN). Each network is assigned a class and models a hyperplane classifier corresponding to the assigned class. During inference, distances from each hyperplane is measured and the sample is assigned the class whose hyperplane is the closest to it.

and

$$\min_{w_{-1}, b_{-1}, \eta} \frac{1}{2} \|X_{-1} w_{-1} + e_{-1} b_{-1}\|_2^2 + C_{-1} e_1 \eta \quad (3)$$

subject to,

$$-(X_1 w_{-1} + e_1 b_{-1}) + \eta \geq e_1, \eta \geq 0 \quad (4)$$

Here, the separating hyperplanes are given by  $(w_1^T x + b_1 = 0)$  and  $(w_{-1}^T x + b_{-1} = 0)$ .  $C_1$  and  $C_{-1}$  are hyperparameters for the individual optimization problems, while  $\xi$  and  $\eta$  represent slack variables,  $e_1$  and  $e_{-1}$  represent vectors of ones, respectively. Since the total number of constraints in these two optimization problems is equal to the original number of constraints, Twin SVM is solving two smaller Quadratic Programming Problems (QPPs). Because of this Twin SVM runs faster compared to the standard SVM as shown in [8]. Another important point to note is that the hyperplanes generated by solving these optimization problems are non-parallel. Because of this flexibility, these hyperplanes can be very close to the cluster of samples of the corresponding class hence increasing generalizability. However, the Twin SVM is not scalable to large datasets, as it requires computation of kernel matrices.

Twin Neural Networks (Twin NN) [5] were formulated as a neural network implementation motivated by the Twin SVM formulation. Any neural network classifier can be considered as an encoder followed by a hyperplane classifier in the final (fully connected) layer. The Twin neural network exploits this fact and uses an ensemble of neural networks with the  $\tanh(\cdot)$  activation to model different hyperplanes in a parameterized kernel space. The aim for the classifying hyperplanes in the output layer is again the same as in the Twin SVM formulation. Not only are Twin neural networks scalable, they also perform better since backpropagating through the hidden layers in these networks allows for implicit kernel optimization. A Twin NN based architecture can therefore benefit from any improvements in deep learning and the advantages of Twin SVMs.

In a binary classification setting, the two neural networks in the twin setup minimize error functions  $E_{(1)}$  and  $E_{(-1)}$ , as given by (5)-(6) ([5]).

$$E_{(1)} = \frac{1}{2N_1} \sum_{i=1}^{N_1} (y^i - o^i)^2 + \frac{C_1}{2N_1} \sum_{i=1}^{N_1} (w_1^T \phi(x_1^i) + b_1)^2 \quad (5)$$

$$E_{(-1)} = \frac{1}{2N_1} \sum_{i=1}^{N_1} (y^i - o^i)^2 + \frac{C_{-1}}{2N_1} \sum_{i=1}^{N_1} (w_{-1}^T \phi(x_{-1}^i) + b_{-1})^2 \quad (6)$$

Here,  $o^i = f(net_i)$  represents the output of the corresponding neuron, where  $net_i = w^T \phi(x^i) + b$ , and  $f(\cdot)$  is an activation function whose value is bounded by  $\pm 1$ , such as  $\tanh(\cdot)$ . Further,  $N_1$  and  $N_{-1}$  represent the number of samples of classes 1 and  $-1$  in the training set respectively. Intuition for the loss function is provided in the Appendix A.



For class inference, we simply measure the distance of the datapoint (encoded representation) from the two hyperplanes and assign the class of the closer hyperplane (see Fig. 6).

The Twin NN setup can also be easily modified for a multi-class classification problem. We simply assign one neural network for each class and all of them solve the same optimization problem in a one vs rest fashion.

In this paper we also modify the twin formulation by allowing for multiple hyperplanes per neural network which can account for disjoint clusters of the same class. This is controlled using the number of planes argument and offers better performance than the original Twin Neural Network [5]. The details and the new loss function have been provided in Appendix B. Apart from this, having  $N$  deep networks (one for each of the  $N$  classes) in a classification setting may have prohibitive computational requirements.

Most deep architectures require significant computation to complete training on large datasets. We hence present Twin Augmentation, a technique that makes use of pre-trained models and confers the benefits of Twin NN without the need to retrain a new architecture. This potentially enables any existing pre-trained architecture to be extended to imbalanced datasets with minimal computation. On the other hand, our technique also allows us to use any architecture and training algorithm that promises SoTA performance for the training of base model; Twin Augmentation can simply work to push that performance even further.

In Twin Augmentation, we use a deep neural network pre-trained on the given dataset as an encoder to reduce the dimensionality of the data. This is simply done by removing its output layer (classification layer). This is illustrated in Fig. 1. The encoding obtained from the penultimate layer of the deep network hence becomes the input to the twin setup. The Twin Neural Network is then solved to obtain the set of optimal hyperplane(s), using (5)-(6). For a test sample, we compute the encoding using the pre-trained (deep) network and use it to obtain the predicted label using the Twin Neural Network model. In the following section, we evaluate our proposed model on datasets and show that it results in improved generalization.

### 5 Experiments and Results

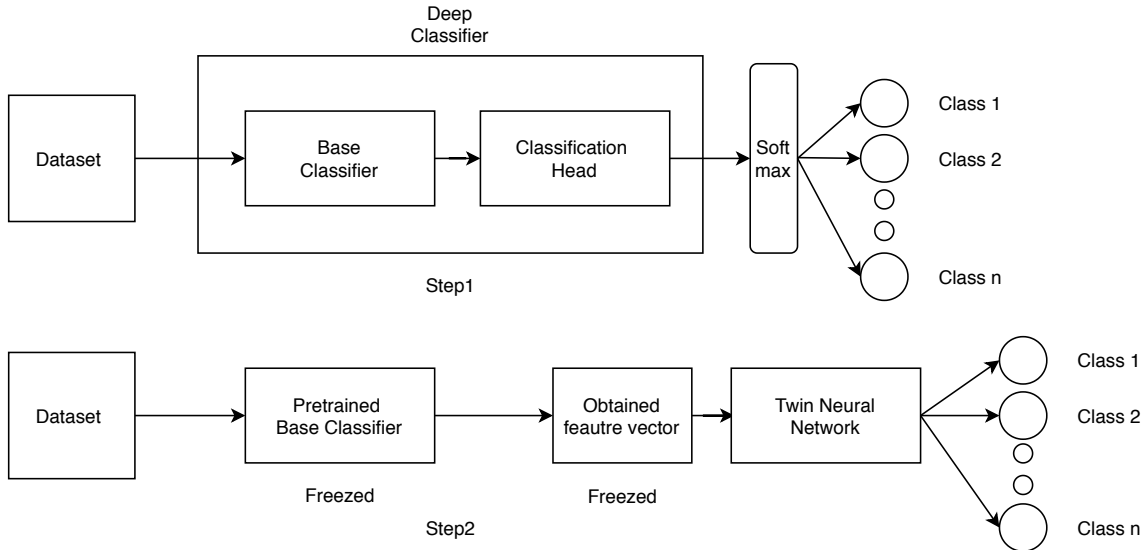


Figure 7: Flowchart which illustrates the pipeline. 1. A deep neural network classifier is trained initially using a selected training algorithm. 2. By removing the final layer, this classifier is turned into an encoder and encodings of the data are generated. These encodings are freezed and is used to train the twin. Final inference is done via the combination of pre-trained encoder and trained twin architecture.

For our experiments on the proposed Twin Augmented deep-learning architectures, we have compared the performance with baselines on the COVID dataset we compiled in Table 6. Results of Twin Augmented architectures are presented

Experiment	Accuracy	Precision	Recall	F1
MobilenetV2 UWC	0.9696 $\pm$ 0.0036	0.8614 $\pm$ 0.0196	0.7956 $\pm$ 0.0548	0.8259 $\pm$ 0.0261
MobilenetV2 Twin UWC	<b>0.9716 <math>\pm</math> 0.0036</b>	<b>0.8108 <math>\pm</math> 0.0425</b>	<b>0.9032 <math>\pm</math> 0.0263</b>	<b>0.8532 <math>\pm</math> 0.0143</b>
MobilenetV2 WC	0.9648 $\pm$ 0.0063	0.7677 $\pm$ 0.0495	0.8924 $\pm$ 0.0846	0.8211 $\pm$ 0.0329
MobilenetV2 Twin WC	<b>0.9736 <math>\pm</math> 0.0063</b>	<b>0.8582 <math>\pm</math> 0.0163</b>	<b>0.8494 <math>\pm</math> 0.0662</b>	<b>0.8528 <math>\pm</math> 0.0391</b>
MobilenetV2 Focal	0.9745 $\pm$ 0.0036	0.8618 $\pm$ 0.0194	0.8602 $\pm$ 0.0662	0.8590 $\pm$ 0.0264
MobilenetV2 Twin Focal	<b>0.9736 <math>\pm</math> 0.0023</b>	<b>0.8129 <math>\pm</math> 0.0201</b>	<b>0.9247 <math>\pm</math> 0.0608</b>	<b>0.8635 <math>\pm</math> 0.0183</b>
MobilenetV2 ADASYN	0.9618 $\pm$ 0.0082	0.8563 $\pm$ 0.0388	0.7096 $\pm$ 0.1602	0.7628 $\pm$ 0.0728
Resnet18 UWC	0.9579 $\pm$ 0.0036	0.8114 $\pm$ 0.0435	0.7096 $\pm$ 0.0912	0.7516 $\pm$ 0.0358
Resnet18 Twin UWC	<b>0.9589 <math>\pm</math> 0.0023</b>	<b>0.7388 <math>\pm</math> 0.0185</b>	<b>0.8494 <math>\pm</math> 0.0152</b>	<b>0.7900 <math>\pm</math> 0.0101</b>
Resnet18 WC	0.9540 $\pm$ 0.0055	0.6888 $\pm$ 0.0257	0.9032 $\pm$ 0.0263	0.7815 $\pm$ 0.0251
Resnet18 Twin WC	<b>0.9667 <math>\pm</math> 0.0013</b>	<b>0.7904 <math>\pm</math> 0.0340</b>	<b>0.8709 <math>\pm</math> 0.0696</b>	<b>0.8257 <math>\pm</math> 0.0139</b>
Resnet18 Focal	0.9530 $\pm$ 0.0082	0.8011 $\pm$ 0.0583	0.6451 $\pm$ 0.0526	0.7140 $\pm$ 0.0515
Resnet18 Twin Focal	<b>0.9550 <math>\pm</math> 0.0036</b>	<b>0.7491 <math>\pm</math> 0.0276</b>	<b>0.7634 <math>\pm</math> 0.0548</b>	<b>0.7546 <math>\pm</math> 0.0243</b>
Resnet18 ADASYN	0.9521 $\pm$ 0.0036	0.7377 $\pm$ 0.0409	0.7419 $\pm$ 0.0263	0.7383 $\pm$ 0.0079

Table 6: Results of Twin Augmented Architecture on our COVID-19 dataset along with comparisons.

in boldface. We used a platform with a 6 core Intel®Xeon @2.3 GHz processor, 32GB RAM and NVIDIA ®Tesla ®K80 for all our experiments.

We used two classifiers pre-trained on ImageNet [23] (MobileNet v2 [22] and ResNet 18 [27]). The input image to the classifiers used are of the size  $224 \times 224 \times 3$ . Number of classes are 2 (COVID-19 Pneumonia vs Non-COVID Pneumonia). We again used PyTorch [29] for all our tests. The batch size was kept as 16 and Adam [28] was used optimizer with learning rate=0.001, betas = 0.9,0.999, epsilon = 1e-08. The models were trained for 100 epochs with early stopping (patience=20). We chose the weights corresponding to the least validation loss, and then tested the model on our test set(unknown data). This generated the results for the base classifier. For the Twin Augmented Network, we removed the final classification layer of this trained model, and passed all train, validation and test samples to generate encodings. These encodings were then used to train, validate and test the Twin Augmented setup. For training the Twin NN block in this Twin Augmented Network, we used mini-batch stochastic gradient descent with lr=0.002, batch size = 30 for 200 epochs. We selected the weights corresponding to the best validation performance, and then tested the Twin Augmented Network on the test data. The numbers generated by following this pipeline (Fig. 7) highlight the functioning of Twin Augmentation as a post-processing boosting step, since each run generates two numbers (one base and one Twin Augmented) using the exact same base model. We perform three such runs and present our results in terms of mean  $\pm$  standard deviation. The results (Table 6) incontrovertibly demonstrate, that there is a significant increase in Precision, Recall, and F1 scores in comparison with the baselines.

For results labelled as UWC, in Table 6 we used the categorical cross-entropy loss function for training the model without any weights. For results labelled WC, weighted categorical cross-entropy was used for training. The weights are 1 for the positive class and 1/10.247 for the negative class, since 10.247 is the class imbalance factor. For results labelled as Focal, we use focal loss with  $\gamma = 2$ . For results labelled as ADASYN [6], we augmented positive class to be exactly of the same size as negative class using the ADASYN algorithm and trained using categorical cross-entropy loss.

For results additionally labelled as Twin, we use the pre-trained base model from the corresponding base test (UWC, WC or Focal) to generate encodings of the data and trained the Twin NN with them hence performing the Twin Augmentation.

From the results, it can clearly be seen that Twin Augmentation outperforms the corresponding base models in F1 score and accuracy. We also report that Twin Augmented networks outperformed the base model in each run consistently. Interestingly for UWC and Focal tests, the major contribution of Twin Augmentation is towards increasing recall by a large factor, whereas in WC tests, its contribution is towards increasing precision by a very high margin. This

suggests that these loss functions stress on very specific aspects of performance. Twin Augmentation, on the other hand, stabilizes the performance by equally stressing both precision and recall. It is important to note that the variance of F1 scores across the runs is lower with Twin Augmentation. This highlights the robustness of Twin Augmented Networks on imbalanced data.

The poor performance of ADASYN suggests that it is not well suited for high-dimensional data. On the above-mentioned system configuration, the time taken for MobileNet v2 [22] tests and ResNet 18 [27] tests (both w/o ADASYN) was around 12300 ms and 10800 ms respectively per epoch. Compared to this, the training time for the Twin NN block to perform Twin Augmentation on MobileNet and ResNet was around 823 ms and 693 ms per epoch. We conclude that time taken for Twin Augmentation is a small fraction of the time taken by the base model, hence showcasing its utility as an efficient post-processing step for boosting performance on imbalanced datasets. We also provide an ablation study in Appendix C.

## 6 Conclusion

We have analysed the usage of heterogeneous sources for CXR based COVID-19 diagnosis using deep learning. Datasets used in many existing approaches fail to capture the diversity in CXR images for the task at hand. We curate a meaningful dataset using t-SNE plots which provides a more accurate representation of the diversity in CXR images for identifying COVID positivity.

We address the problem of class imbalance which is naturally inherent in many datasets, or may be induced when compiling the dataset into train, validation, and test splits. The proposed Twin Augmentation works as an efficient, robust and general post processing step of boosting a classifier’s performance on an imbalanced dataset. We show that Twin Augmentation outperforms popular techniques used to tackle class imbalance. Twin Augmentation is generic enough to work on a variety of classifiers, without any change in the training pipeline. It also takes a small fraction of computational time compared to training the base model. The post-processing nature makes it extremely flexible to be used for a variety of tasks, as it would benefit from the improvement in performance of the existing base model or training algorithm.

## References

- [1] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 590–597, 2019.
- [2] Joseph Paul Cohen, Paul Bertin, and Vincent Frappier. Chester: A web delivered locally computed chest x-ray disease prediction system. *arXiv preprint arXiv:1901.11210*, 2019.
- [3] Joseph Paul Cohen, Paul Morrison, and Lan Dao. Covid-19 image data collection. *arXiv 2003.11597*, 2020.
- [4] Justin M Johnson and Taghi M Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1):27, 2019.
- [5] Himanshu Pant, Mayank Sharma, Sumit Soman, et al. Twin neural networks for the classification of large unbalanced datasets. *Neurocomputing*, 343:34–49, 2019.
- [6] Haibo He, Yang Bai, Eduardo A Garcia, and Shutao Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*, pages 1322–1328. IEEE, 2008.
- [7] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [8] Jayadeva, Reshma Khemchandani, and Suresh Chandra. Twin support vector machines for pattern classification. *IEEE Transactions on pattern analysis and machine intelligence*, 29(5):905–910, 2007.
- [9] R Bharat Rao, Sriram Krishnan, and Radu Stefan Niculescu. Data mining for improved cardiac care. *ACM SIGKDD Explorations Newsletter*, 8(1):3–10, 2006.
- [10] Wei Wei, Jinjiu Li, Longbing Cao, Yuming Ou, and Jiahang Chen. Effective detection of sophisticated online banking fraud on extremely imbalanced data. *World Wide Web*, 16(4):449–475, 2013.
- [11] Matthew Herland, Taghi M Khoshgoftaar, and Richard A Bauder. Big data fraud detection using multiple medicare data sources. *Journal of Big Data*, 5(1):29, 2018.

- [12] Miroslav Kubat, Robert C Holte, and Stan Matwin. Machine learning for the detection of oil spills in satellite radar images. *Machine learning*, 30(2-3):195–215, 1998.
- [13] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- [14] Shoujin Wang, Wei Liu, Jia Wu, Longbing Cao, Qinxue Meng, and Paul J Kennedy. Training deep neural networks on imbalanced data sets. In *2016 international joint conference on neural networks (IJCNN)*, pages 4368–4374. IEEE, 2016.
- [15] Haoyue Liu, MengChu Zhou, and Qing Liu. An embedded feature selection method for imbalanced data classification. *IEEE/CAA Journal of Automatica Sinica*, 6(3):703–715, 2019.
- [16] Linda Wang and Alexander Wong. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest radiography images. *arXiv*, pages arXiv:2003, 2020.
- [17] Joseph Paul Cohen, Paul Morrison, Lan Dao, Karsten Roth, Tim Q Duong, and Marzyeh Ghassemi. Covid-19 image data collection: Prospective predictions are the future. *arXiv 2006.11988*, 2020.
- [18] Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131, 2018.
- [19] Feng Shi, Jun Wang, Jun Shi, Ziyang Wu, Qian Wang, Zhenyu Tang, Kelei He, Yinghuan Shi, and Dinggang Shen. Review of artificial intelligence techniques in imaging data acquisition, segmentation and diagnosis for covid-19. *IEEE Reviews in Biomedical Engineering*, 2020.
- [20] Biraja Ghoshal and Allan Tucker. Estimating uncertainty and interpretability in deep learning for coronavirus (covid-19) detection. *arXiv preprint arXiv:2003.10769*, 2020.
- [21] Ali Narin, Ceren Kaya, and Ziyne Pamuk. Automatic detection of coronavirus disease (covid-19) using x-ray images and deep convolutional neural networks. *arXiv preprint arXiv:2003.10849*, 2020.
- [22] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [24] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [25] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2097–2106, 2017.
- [26] Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, 2016.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019.
- [30] Christopher JC Burges. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery*, 2(2):121–167, 1998.
- [31] PS Bradley and OL Mangasarian. Massive data discrimination via linear support vector machines. *Optimization methods and software*, 13(1):1–10, 2000.
- [32] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [33] Vladimir Cherkassky and Filip M Mulier. *Learning from data: concepts, theory, and methods*. John Wiley & Sons, 2007.

## A Appendix

### A.1 Loss functions for the neural networks in twin setting for binary classification problem

The final layer in the Twin Neural Network (Twin NN) needs to use an activation function such as *tanh* activation, or any other origin symmetric activation function, though it is common to have all layers using the same. The TNN may be viewed as a set of layers that learn a map that transforms input samples into an appropriate feature space, followed by a final layer that acts as a hyperplane classifier. In a binary classification setting, the final layer has two neurons, one corresponding to each class. The decision boundary learnt by each neuron corresponds to a hyperplane that lies in the feature space determined by the previous layers. The first neuron's hyperplane is required to pass through samples of Class 1, while being far away from samples of the other class (Class -1). Let the net input to the neuron be denoted by  $net^i$  when sample  $x^i$  is presented at the input layer; the output is then given by  $f(net_1^i)$  (where  $f(\cdot)$  denotes the activation function). Assume that  $f(\cdot)$  lies between -1 and 1. For the Class 1 output neuron, we require

$$f(net_1^i) = \begin{cases} 0, & \text{if } y_i = 1 \\ -1, & \text{if } y_i = -1 \end{cases} \quad (7)$$

where  $y_i$  is the class label of sample  $x^i$ . Note that  $net^i = w_1^T \phi(x^i) + b_1$ , where  $\phi(x^i)$  is the image vector formed at the penultimate layer corresponding when sample  $x^i$  is presented at the input;  $w_1$  is the weight vector of the class 1 neuron, and  $b_1$  is its bias. Since  $f(0) = 0$ ,  $f(net^i) = 0$  implies that  $w_1^T \phi(x^i) + b_1 = 0$ , which means that the hyperplane  $w_1^T \phi(x^i) + b_1 = 0$  passes through the image vector  $\phi(x^i)$ . Similarly,  $f(net^i) = -1$  implies that  $w_1^T \phi(x^i) + b_1 = -\infty$ , which means that the hyperplane  $w_1^T \phi(x^i) + b_1 = 0$  lies far away from image vector  $\phi(x^i)$ . Let the weight vector and bias of the Class -1 output neuron be denoted by  $w_{-1}$  and  $b_{-1}$ , respectively. Then, for the Class -1 output neuron, we require

$$f(net_{-1}^i) = \begin{cases} 0, & \text{if } y_i = -1 \\ +1, & \text{if } y_i = +1 \end{cases} \quad (8)$$

where  $net_{-1}^i = w_{-1}^T \phi(x^i) + b_{-1}$ . Given a test sample  $x$  at the input, we determine the geometric distances of the image vector to both hyperplanes. The class label  $y$  of the sample  $x$  is determined from the closer hyperplane. That is,

$$y = \begin{cases} 1, & \text{if } d_1 < d_{-1} \\ -1, & \text{if } d_{-1} < d_1 \end{cases} \quad (9)$$

where,

$$d_1 = \frac{\|net_1\|}{\|w_1\|}, d_{-1} = \frac{\|net_{-1}\|}{\|w_2\|} \quad (10)$$

It may be noted that the loss functions for the two networks in the twin setup can be written as (5) and (6) in Section 4.

### A.2 Number of planes argument

The Twin NN uses a generalization of the above concept. Each class 1 and class -1 neural network is associated with  $k > 1$  hyperplanes. This additional argument allows each network to easily work with upto  $k$  disjoint clusters of the corresponding class in the dataset. During training, when a sample of Class 1 is presented, the closest class 1 hyperplane (out of the  $k$  hyperplanes) is determined and its weights are updated to be closer to the sample. Similarly the closest class -1 hyperplane is determined and its weights are updated to be as far as possible from the same sample.

During testing, the closest hyperplane from each class is used to determine the class of the input sample. The integer  $k$  is termed as the *number of planes* argument.

We define:

$$a_{m,j,l}^i = w_{m,j}^T \phi(x_l^i) + b_{m,j} \quad (11)$$

and

$$a_{j,l}^i = \min_{m \in \{0, \dots, k-1\}} |a_{m,j,l}^i| \quad (12)$$

Here,  $j$  is the class of the network (+1 or -1 for binary classification problem),  $m$  is the  $m^{th}$  hyperplane of the  $j^{th}$  class neural network and  $x_l^i$  corresponds to the  $i^{th}$  sample of  $l^{th}$  class. Hence the loss function in the binary setting using the *tanh*( $\cdot$ ) activation function is given by

$$E_{(1)} = \frac{1}{2N_{-1}} \sum_{i=1}^{N_{-1}} (y^i - \tanh(a_{1,-1}^i))^2 + \frac{C_1}{2N_1} \sum_{i=1}^{N_1} (a_{1,1}^i)^2 \quad (13)$$

$$E_{(-1)} = \frac{1}{2N_1} \sum_{i=1}^{N_1} (y^i - \tanh(a_{-1,1}^i))^2 + \frac{C_{-1}}{2N_{-1}} \sum_{i=1}^{N_{-1}} (a_{-1,-1}^i)^2 \quad (14)$$

### A.3 Ablation study for Covid CXR Dataset

There are two main hyperparameters in the twin architecture. One is the number of hidden layers in each network. This corresponds to the capacity of the network to perform implicit kernel optimization. The other argument is the number of planes argument. It should be equal to the maximum of the number of disjoint clusters in each class for the best performance (as shown in Fig. 8).



Figure 8: The effect of using the number of planes argument. For data having disjoint clusters, having more than one number planes helps in better generalization. Here blue is class +1 and orange is class -1.

We present our results in Fig. 9. While changing the number of hidden layers, we kept number of hyperplanes = 4. The first hidden layer was of 256 neurons and then we added a layer of 128 neurons each time. For number of planes test, we fixed the architecture to have two hidden layers (one of 256 neurons, other of 128 neurons) and varied the number of planes argument. We used Mobilenet v2 [22] encodings using unweighted cross-entropy loss. Interestingly we note that the twin is actually very stable to hyperparameter variation hence adding to its robustness. We also note that in each run, the F1 was better than the base classifier.

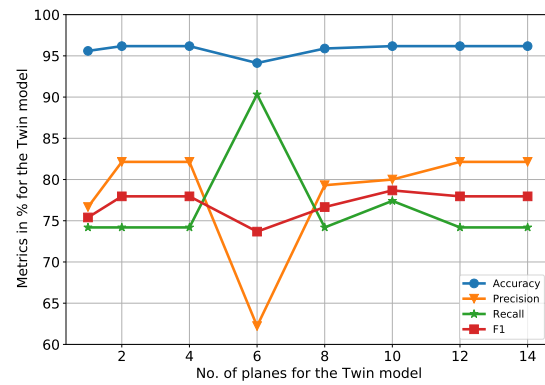
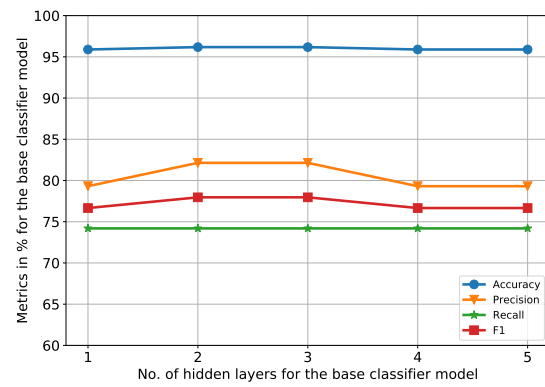


Figure 9: Results of Ablation study. (Top) No. of hidden layers was varied using the same encodings. (Bottom) No. of planes was varied using the same encodings