



# Detection of COVID-19 from X-rays using hybrid deep learning models

Ritika Nandi<sup>1</sup> · Manjunath Mulimani<sup>1</sup>

Received: 5 May 2021 / Accepted: 4 September 2021 / Published online: 21 September 2021  
© Sociedade Brasileira de Engenharia Biomedica 2021

## Abstract

**Purpose** To propose a model that can detect the presence of Covid-19 from chest X-rays and can be used with low hardware resource-based personal digital assistants (PDA).

**Methods** In this paper, a hybrid deep learning model is proposed for the detection of coronavirus from chest X-ray images. The hybrid deep learning model is a combination of ResNet50 and MobileNet. Both ResNet50 and MobileNet are light deep neural networks (DNNs) and can be used with low hardware resource-based personal digital assistants (PDA) for quick detection of COVID-19 infection.

**Results** The performance of the proposed hybrid model is evaluated on two publicly available COVID-19 chest X-ray datasets. Both datasets include normal, pneumonia, and coronavirus-infected chest X-rays and we achieve 84.35% and 94.43% accuracy on Dataset 1 and Dataset 2 respectively.

**Conclusion** Results show that the proposed hybrid model is better suited for COVID-19 detection.

**Keywords** COVID-19 detection · MobileNet · ResNet50 · Hybrid model · Pneumonia · X-rays

## Introduction

The unprecedented increase in the number of cases of the new respiratory disease COVID-19 has left the healthcare sector frantic. Due to this sudden increase, many countries across the world are facing a shortage of resources to battle this pandemic. An important step towards mitigating this crisis is the quick and accurate detection of infected patients in order to utilize these limited resources wisely. The most effective tool for successfully identifying the infected patients is polymerase chain reaction (PCR) testing. PCR tests are utilized to directly sight the presence of the COVID-19 antigen, instead of the body's invulnerable reaction, or antibodies. The tests can tell whether or not

someone has the virus very early on, as they detect the viral RNA, which will be present in the body before symptoms start to appear or antibodies are formed (Kent 2020). However, this is a very complicated manual process that is extensively laborious, and time-consuming. Due to the sheer volume of reported cases, this process is in short supply.

It was established by early research (Huang et al. 2020; Ng et al. 2020) that patients infected with COVID-19 display abnormalities in the radiography scans. As a result, detection of visual indicators associated with COVID-19 infection through chest radiography images (e.g., X-ray and CT scans) could be used as an alternative to PCR testing successfully. One of the biggest challenges in this avenue is the ability of radiologists to successfully identify and interpret the minute indicators of COVID-19.

In an attempt to overcome this restriction, we suggest a novel model based on deep learning which can recognize the indicators of these infections in the chest X-ray images with almost perfect accuracy. However, we wish to stress the fact that we are not advocating for the employment of this model as a substitute to standard testing procedures for the detection of COVID-19, but merely as a sorting tool which can aid clinicians in improved screening (Fig. 1).

✉ Manjunath Mulimani  
manjunath.gec@gmail.com

Ritika Nandi  
ritika.nandi77@gmail.com

<sup>1</sup> Department of Computer Science and Engineering,  
Manipal Institute of Technology, Manipal Academy of Higher  
Education, Manipal, 576 104, India

## Related work

Over the last few years, the quantity of medical information in machine-readable form has increased tremendously, notably in medical imaging. New forms of neural networks have excelled at the task of identifying abnormalities from high-dimensional information, where previous generations of algorithms failed. Tasks like segmentation, detection, and classification from X-ray, CT scans, and MRIs have quickly become the “de facto” standard (Lundervold and Lundervold 2019). As additional proof, models dependent on neural networks command the leaderboard in most clinical imaging difficulties (Li et al. 2017).

Various deep learning models have been deployed to detect the presence of respiratory diseases, particularly pneumonia from chest X-ray scans (Rajpurkar et al. 2017; Ekagra et al. 2018; Wang et al. 2017; Weng et al. 2017) in recent years. These models typically range from basic convolutional neural networks (Jadhav and Yadav 2019) to complex network architectures like DenseNet121 (Rajpurkar et al. 2017; Li et al. 2017). A benchmark model for the detection of pneumonia from chest radiography images is CheXNet (Rajpurkar et al. 2017) which is built upon DenseNet121.

Keeping pace with the recent pandemic situation, several researchers have worked on machine learning and deep learning models. Machine learning algorithms like support vector machines (SVM) and random forest have been used by Alqudah et al. (2020), to detect COVID-19 cases with an accuracy of 90.5% and 81%, respectively. Ghoshal and Tucker (2020) suggested a model based on Bayesian CNN to detect COVID-19 from chest X-ray images and obtained an accuracy of 90%. CNN-based architectures have been proposed by various works (Salman et al. 2020; Narin et al. 2020; Wang and Wong 2020; Mangal et al. 2020) that can be used effectively to identify the presence of COVID-19 infection from Chest Radiology images. The COVID-Net architecture, suggested by (Wang and Wong 2020), is one of the first CNN-based models which identifies COVID-19 as well as bacterial and viral pneumonia while demonstrating significant sensitivity for COVID-19 detection. One disadvantage of this model is that it is an extremely resource-intensive model. COVIDAid model (Mangal et al. 2020), which is built upon the CheXNet (Rajpurkar et al. 2017), performs great on the task of detecting COVID-19 from images, achieving 90.5% accuracy. Abbas et al. (2020) proposed a DeTraC (Decompose, Transfer, and Compose) deep convolutional neural network for the classification of COVID-19 chest X-ray images and obtained an accuracy of 93.1% using this architecture. The Covidx-net model (Hemdan et al. 2020), proposes automatic detection of coronavirus infected persons using chest X-ray images

using seven different transfer learning models. They obtained an average accuracy of 91% for COVID cases using VGG and DenseNet models.

Unsurprisingly, deep learning models (Salman et al. 2020; Narin et al. 2020; Wang and Wong 2020; Mangal et al. 2020) perform better than machine learning models (Alqudah et al. 2020; Ghoshal and Tucker 2020) for the classification task. However, these DNN architectures are primarily built on heavy transfer learning models like DenseNet and VGG. They are deeper, slower, more complicated, and may not be suitable for personal digital assistants. In this paper, we utilize two light frameworks - MobileNet (Howard et al. 2017), and ResNet50 (He et al. 2016) and mould it according to our problem statement, for quick identification of COVID-19 infection. We have chosen two models - COVID-Net (Wang and Wong 2020) and COVIDAid (Mangal et al. 2020) as our baseline for an accurate comparison, as we have performed our experiments on the same data split as Covid-Net and COVIDAid models.

## Methods

### Datasets

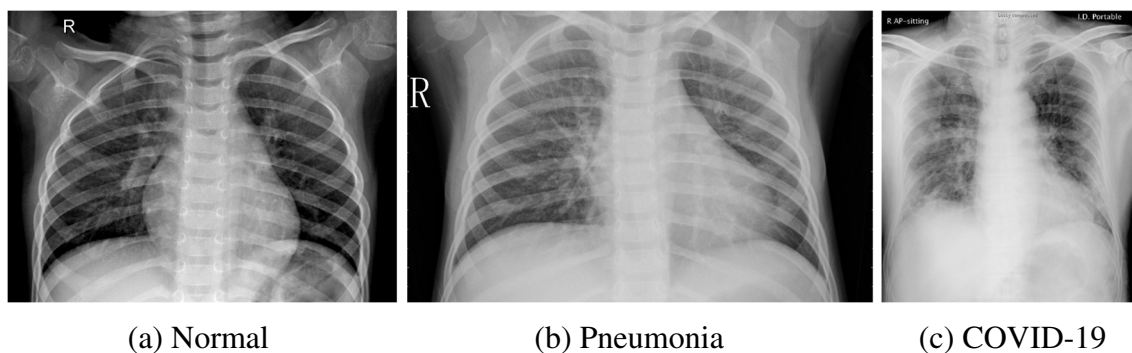
Two versions of the COVIDx data proposed by Wang and Wong (2020) are used in this paper. The sample-wise distribution of both the datasets is mentioned in Table 1. For both datasets, we have applied 5-fold cross-validation to the training set. 10% of the training data has been used for validation in each fold. We have used the average accuracy obtained from each fold as the final accuracy.

**Dataset 1** The first version of the dataset consists of 5949 posterior chest radiography images divided into 4 classes—normal, bacterial pneumonia, viral pneumonia, and COVID-19. Two publicly available datasets, (1) COVID chest X-ray dataset by Cohen et al. (2020) and (2) Kaggle chest X-ray images (pneumonia) dataset by Kermamy et al. (2018), are combined to create this dataset. The train/test distribution of the classes is depicted in Fig. 2a.

**Dataset 2** The second version contains 15,476 posterior chest radiography images divided into 3 classes—normal,

**Table 1** Sample wise dataset distribution

Dataset	Normal	Pneumonia		COVID-19
		Bacterial	Viral	
Dataset 1	2786	1589	1498	76
Dataset 2	8851	6052		573

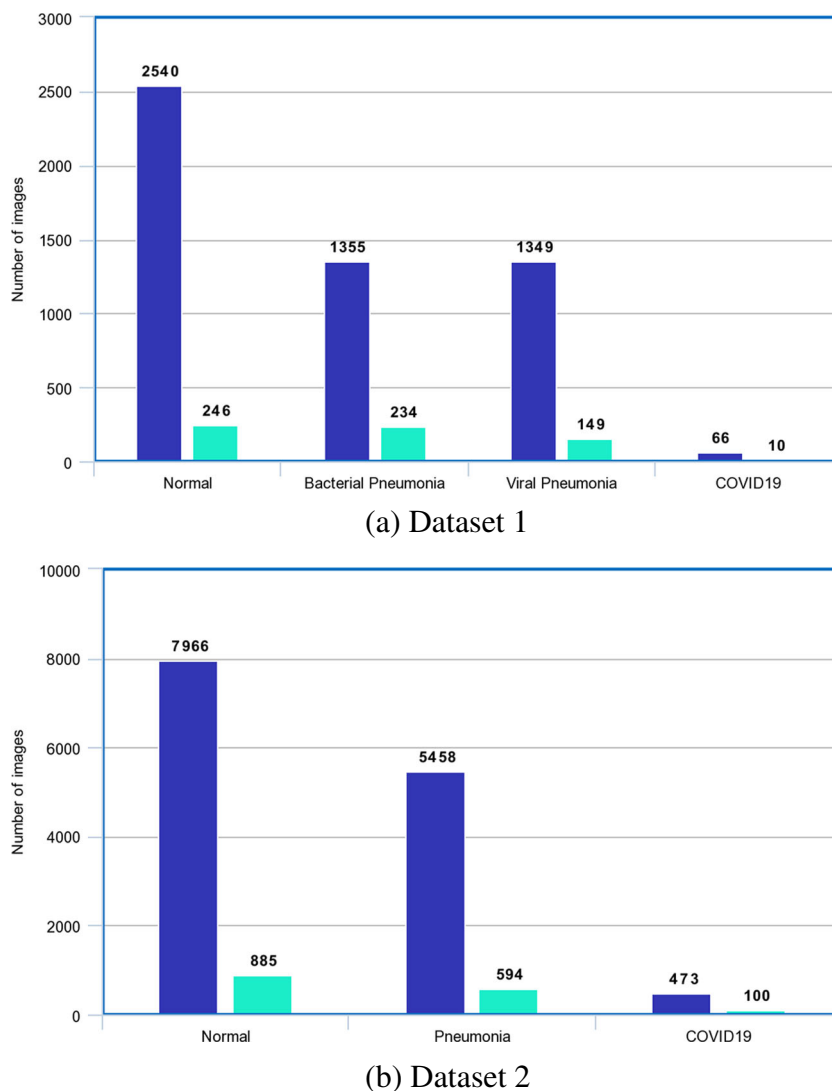


**Fig. 1** Examples of normal, pneumonia-infected, and COVID-19-infected chest X-ray images

pneumonia, and COVID-19 (see Fig. 1). Five different open-source repositories have been combined to create the second version of the dataset—(1) COVID-19 image data collection by Cohen et al. (2020), (2) Fig. 1 COVID-19 chest X-ray dataset initiative by (Chung 2020b), (3) ActualMed COVID-19 chest X-ray dataset initiative by

Chung (2020a), (4) RSNA pneumonia detection challenge dataset by Rsna pneumonia detection challenge (2018), which used publicly available CXR data from (Wang et al. 2017), and (5) COVID-19 radiography database by Chowdhury et al. (2020). The train/test distribution of the classes are shown in Fig. 2b.

**Fig. 2** Train-test class-wise distribution of the chest X-ray images



## Deep learning architectures

The proposed hybrid model is a combination of ResNet50 and MobileNet, which are explained below in brief.

**ResNet50** Proposed by He et al. (2016), ResNet-50 is a light architecture of a deep convolutional neural network. It contains 50 layers, which, instead of learning unreferenced functions, reformulate as learning residual functions with reference to the layer inputs (He et al. 2016). The ResNet model consists of a stack of comparable or “residual” blocks. This block acts as a stack of convolutional layers. The output of a block is additionally connected with its own input through an identity mapping path. To preserve the time complexity per layer, the feature mapping is repeatedly down-sampled by strided convolution along with the increase in channel depth (He et al. 2016). In the ResNet50 model, to allow for a reduced computational load while calculating the  $3 \times 3$  convolution, we have a 3-layer bottleneck block that uses  $3 \times 3$  convolutions to lower and consequently reestablish the channel depth.

**MobileNet** Proposed by Howard et al. (2017), MobileNet is a deep learning model that is lightweight in nature. The base paper (Howard et al. 2017) presents comprehensive tests on the resource and accuracy tradeoffs and demonstrating superior performance compared to other popular models. MobileNet uses depthwise separable convolutions, i.e., it performs a single convolution on each color channel rather than combining all three and flattening it as in standard 2D CNN. This has the effect of filtering the input channels. The model performance increases due to these depthwise

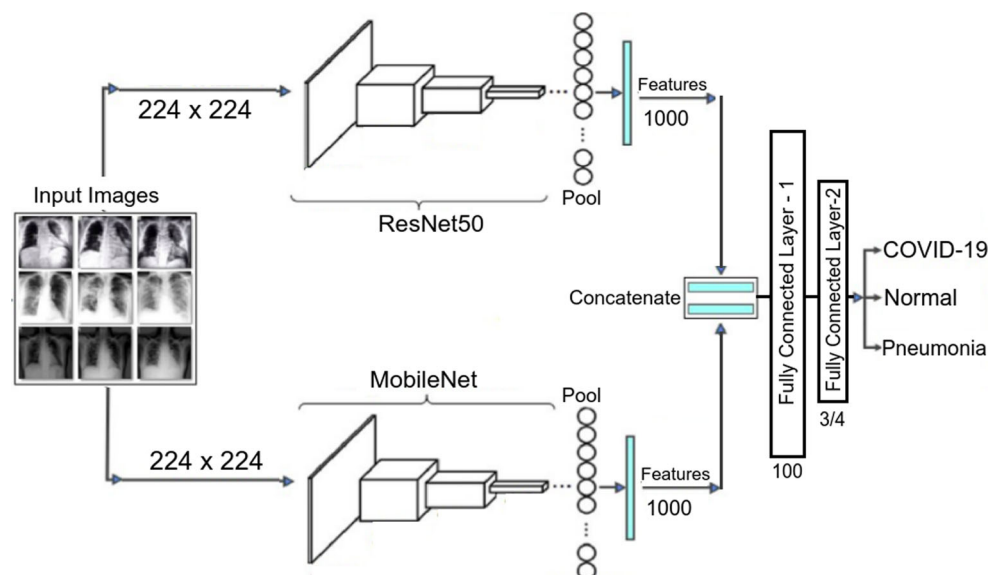
convolutions, and the input features are divided into two layers. Each layer is partitioned into the following layer by joining it with the output features until the procedure is finished. ReLU activation function is used between layers in MobileNet architecture (Howard et al. 2017). Thus, it allows the non-linear outputs from the previous layer to be flattened and fed as input to the subsequent layer.

We believe that our architecture works well as it combines *identity mapping* of ResNet50, diminishing the vanishing gradient problem and boosting the gradient backward flow in the network, and *depthwise separable convolutions* of MobileNets, which increases model performance.

## Implementation details

Large volumes of data are required to train a neural network from scratch. As a result, due to the scarcity of openly available Chest X-Ray samples (especially of COVID-19), we conduct our experiments using two deep learning models, ResNet50 and MobileNet, pre-trained on the ImageNet (Deng et al. 2009) dataset. These models are used separately and their individual results are concatenated to produce results. This is followed by two fully connected layers. Our final classification layer consists of 3 classes in the case of normal, pneumonia, COVID-19 configuration, and 4 classes where pneumonia is split into bacterial or viral. The architecture of the model is detailed in Fig. 3. This model is then trained on the COVID-19 dataset. Adam optimizer is used to update the neural network weights. We use early stopping on the lowest validation loss. The hyperparameters for training are: learning rate=1e-3, batch size=32, patience=15. The problem is treated as a multi-class classification problem, where each sample belongs to

**Fig. 3** Architecture of the proposed hybrid model



**Table 2** Model parameters and the time taken to train ResNet50, MobileNet, and our proposed hybrid model

Model	Trainable parameters	Non-trainable parameters	Total parameters	Time taken per epoch	
				Dataset 1	Dataset 2
ResNet50	23,739,896	53,120	23,793,016	~ 39s 7ms	~ 101s 7ms
MobileNet	3,309,880	21,888	3,331,768	~ 12s 45ms	~ 49s 12ms
Hybrid model	27,049,272	75,008	27,124,280	~ 46s 33ms	~ 126s 3ms

**Table 3** 5-fold cross validation performances of the proposed model with other baseline models on datasets 1 and 2

Model	Dataset	Accuracy(%)				
		Fold-1	Fold-2	Fold-3	Fold-4	Fold-5
ResNet50	Dataset 1	75.63	78.25	80.31	81.43	81.88
	Dataset 2	90.80	92.52	89.70	91.67	90.97
MobileNet	Dataset 1	82.50	85.35	81.23	83.08	84.14
	Dataset 2	95.28	92.90	93.56	93.89	93.02
ResNet50 + MobileNet	Dataset 1	82.86	83.52	85.89	84.73	84.76
	Dataset 2	94.80	95.65	93.27	94.78	93.72

**Table 4** Comparison of overall recognition Accuracy(%) and AUROC of the proposed model with other Baseline models on datasets 1 and 2

Model	Dataset	Class	AUROC	Overall accuracy(%)
ResNet50	Dataset 1	Normal	0.96	79.50
		Bacterial pneumonia	0.94	
		Viral pneumonia	0.87	
		COVID-19	0.98	
	Dataset 2	Normal	0.98	91.13
		Pneumonia	0.97	
		COVID-19	0.96	
MobileNet	Dataset 1	Normal	0.97	83.26
		Bacterial pneumonia	0.95	
		Viral pneumonia	0.93	
		COVID-19	<u>1.00</u>	
	Dataset2	Normal	0.98	93.73
		Pneumonia	<u>0.98</u>	
		COVID-19	<u>0.99</u>	
ResNet50 + MobileNet	Dataset 1	Normal	0.98	84.35
		Bacterial pneumonia	0.97	
		Viral pneumonia	0.95	
		COVID-19	<u>1.00</u>	
	Dataset 2	Normal	0.99	94.43
		Pneumonia	<u>0.98</u>	
		COVID-19	<u>0.99</u>	

one of the 3 or 4 classes. We use the categorical cross-entropy loss as given in (1).

$$L(y, \hat{y}) = - \sum_{j=0}^M \sum_{i=0}^N (y_{ij} \log(\hat{y}_{ij})) \quad (1)$$

The parameters chosen for training the models are the recommended values given by previous research which ensure an optimum performance. These recommended values have been taken as the baseline and subsequently, fine tuning of the hyperparameters has been performed to yield better results. The architecture was built and evaluated using the Keras, a python library for deep learning tasks which uses TensorFlow as a backend. The complexity of each model depends on various factors like the number of parameters, time taken to train the model, and optimizers used as shown in Table 2.

## Results

Performance metrics like accuracy and area under ROC (AUROC) curves are used to carry out an experimental analysis. We have presented the fold-wise accuracy of two baseline models—ResNet50, and MobileNet, and our proposed hybrid model in Table 3. The overall accuracy in Table 4 is calculated as the average of accuracies obtained from each fold. Table 4 also presents the class-wise AUROC of each model. The class-wise ROC curves are presented in Fig. 5.

Using the pre-trained ResNet50 model, we obtained an average accuracy of 79.50% and 91.13% on Datasets 1 and 2 respectively. The pre-trained MobileNet model classifies the images with an average accuracy of 83.26% and 93.73% on Datasets 1 and 2 respectively. The best result is obtained by concatenating the individual outputs of ResNet50 and

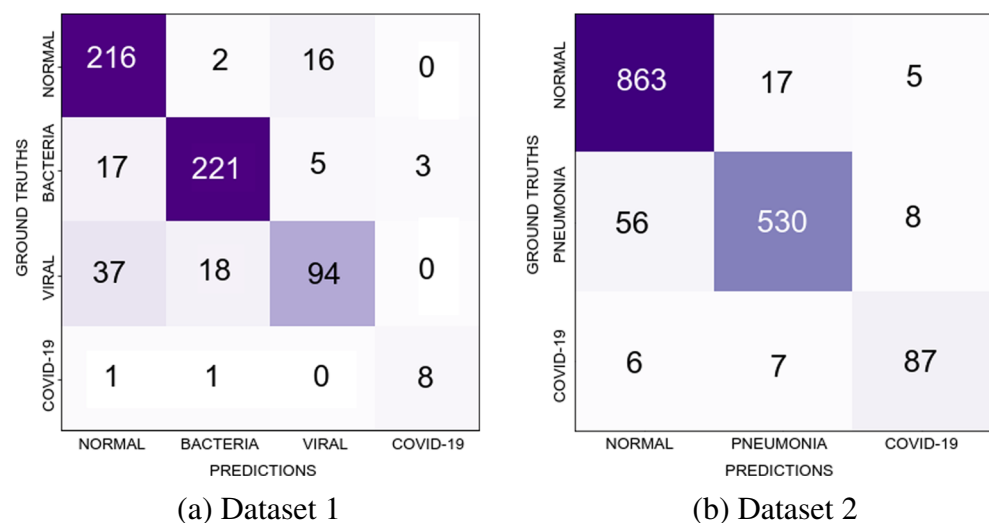
MobileNet, before passing them through the fully connected layer. Using this hybrid approach, we obtain an average accuracy of 84.35% and 94.43% on Datasets 1 and 2 respectively. This architecture leads to the detection of COVID-19 from the radiography images with an AUROC (Area under ROC curve) of 1.00 for the 4-class classification and 0.99 for the 3-class classification. The confusion matrix obtained from the hybrid model, for both the versions of the datasets, is shown in Fig. 4. In Dataset 1, it can be seen that some pneumonia cases have been misclassified as normal cases (Fig. 4a). We believe that the misclassified images may have a lower viral or bacterial infection load, due to which they are being misclassified as normal cases. However, this is just a hypothesis on our part (Fig. 5).

Table 5 shows the difference in the performance of the ResNet50 and MobileNet models, and the ResNet50 + MobileNet hybrid model proposed in this paper by *T*-test. There were statistically significant differences between ResNet50 and MobileNet models, when compared with the hybrid model ( $p < 0.5$ ).

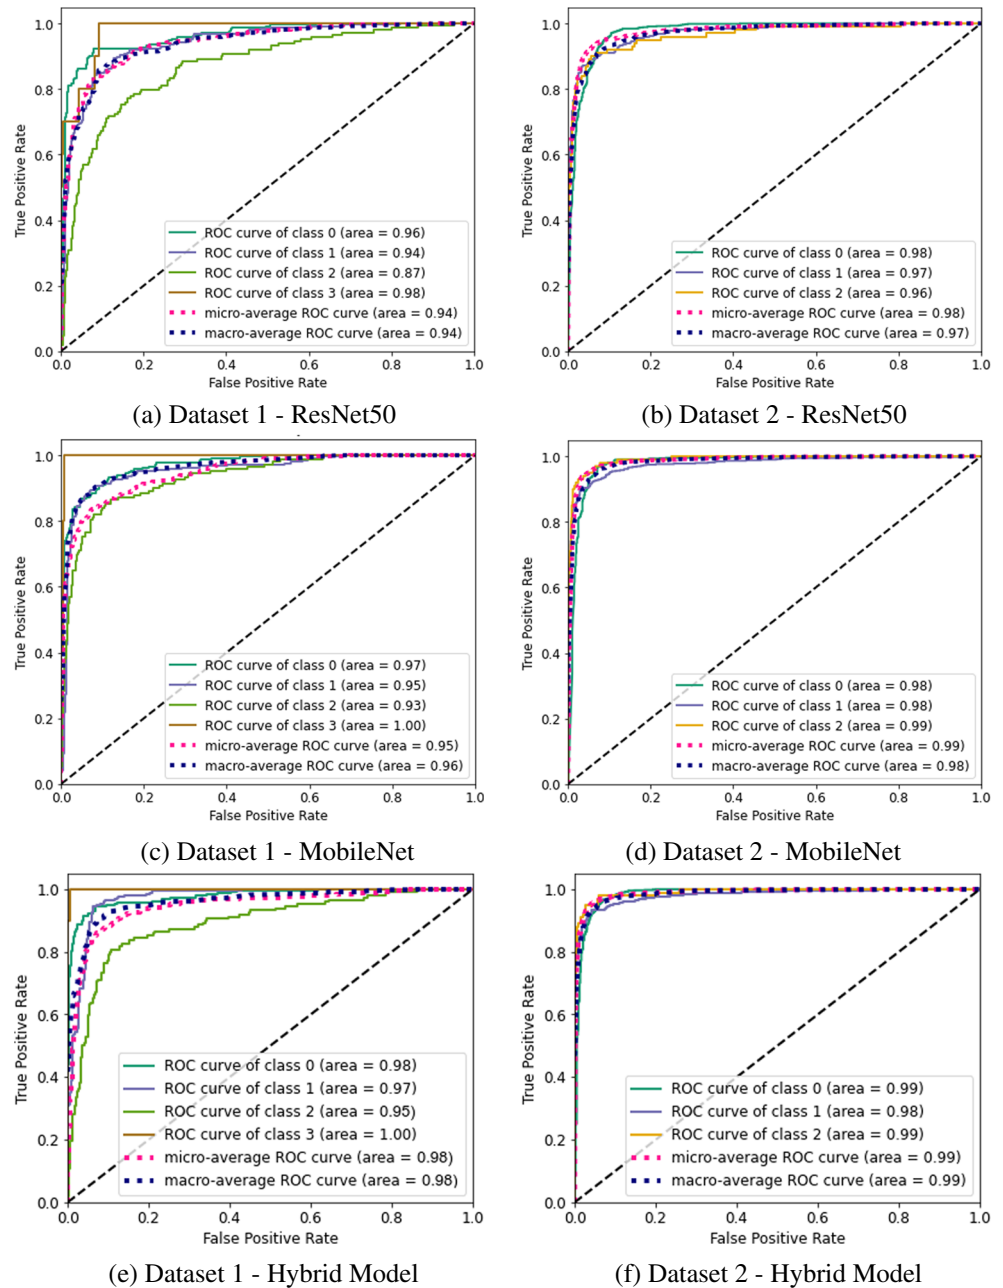
## Discussion

In order to obtain a general picture of the performance of our proposed model on the two case studies, we compare our model with COVIDAid (Mangal et al. 2020) and Covid-Net (Wang and Wong 2020) by evaluating on two separate datasets to eliminate intrinsic biases and impart robustness to the findings. Since we have used the same dataset and data split as the COVID-Net model, we use the values obtained from (Wang and Wong 2020) as it is. In the same way, for COVIDAid, we evaluate their approach on the same data split as ours and obtain the accuracy, ROC curve, and Confusion Matrix. As shown in Table 6, we obtain higher accuracy than the COVID-Net and COVIDAid models on

**Fig. 4** Confusion matrices of the hybrid model on Dataset 1 and Dataset 2



**Fig. 5** ROC curves obtained from the ResNet50, MobileNet, and hybrid models. In Dataset 1, classes 0, 1, 2, 3 refer to normal, bacterial pneumonia, viral pneumonia, and COVID-19 respectively; and in Dataset 2, classes 0, 1, 2 refer to normal, pneumonia, and COVID-19 respectively



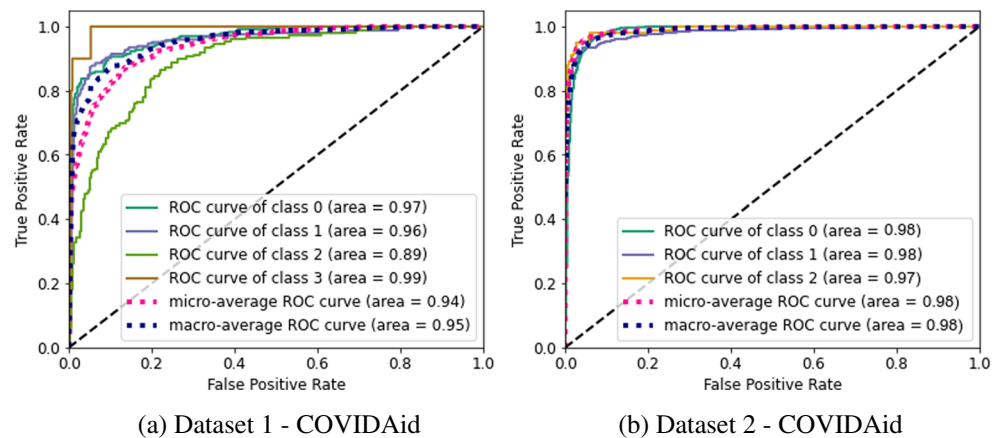
**Table 5** Comparison of performance ResNet50 and MobileNet models with the hybrid model using *T*-test

Model	Dataset	<i>T</i> -test <i>p</i> value
ResNet50	Dataset 1	0.0051
	Dataset 2	0.0008
MobileNet	Dataset 1	0.2495
	Dataset 2	0.2696

**Table 6** Comparison of accuracy (%) and statistical significance of baseline models with our proposed model

Model	Accuracy		<i>T</i> -test <i>p</i> value	
	Dataset 1	Dataset 2	Dataset 1	Dataset 2
COVID-Net	83.5	93.3	0.0033	0.0529
COVIDAid	82.47	93.41	0.0069	0.0084
Our model	84.35	94.43		

**Fig. 6** ROC curves of our model compared to COVIDAid model. In Dataset 1, classes 0, 1, 2, 3 refer to normal, bacterial pneumonia, viral pneumonia, and COVID-19 respectively; and in Dataset 2, classes 0, 1, 2 refer to normal, pneumonia, and COVID-19 respectively



the same data split. It is worth pointing out that lightweight MobileNet slightly outperforms both complex COVID-Net and COVIDAid architectures. The combination of ResNet50 and MobileNet reasonably outperform COVID-Net and COVIDAid architectures. Table 6 also shows that there are statistically significant differences between both baseline models, and the hybrid model ( $p < 0.5$ ). Hence, the hybrid model is more effective in coronavirus detection.

Fig. 6a and b presents the AUROC obtained from the COVIDAid model. Comparing with the ROC with our proposed hybrid approach in Fig. 5e and f, in 4 class classification, our model performs comparable to the COVIDAid model in normal, bacterial pneumonia, and COVID-19 classes, but performs significantly better in the viral pneumonia class. In 3 class classification, our model performs better in the normal and COVID-19 classes than the COVIDAid model. This hybrid approach clearly establishes a higher baseline for detection of pneumonia or COVID-19.

## Conclusion

In this paper, ResNet50, MobileNet, and a hybrid approach combining ResNet50 and MobileNet are used for the detection of COVID-19 from chest X-rays. The performance of the proposed model is evaluated on two publicly available datasets. Results show that our approach outperforms competitive COVID-Net and COVIDAid. It indicates that the proposed model has a significant contribution to the detection of COVID-19 infected chest X-rays. Proposed models for COVID-19 detection are light DNN as compared to other DNNs and can be used with low resource devices for quick identification of coronavirus.

**Data availability** The data that support the findings of this study are openly available and details regarding the same are mentioned in Section 1.

**Code availability** The code that supports the findings of this study is not openly available currently and will be made available after the publication of the work.

## Declarations

**Conflict of interest** The authors declare no competing interests.

## References

- Abbas A, Abdelsamea MM, Gaber MM. Classification of covid-19 in chest x-ray images using detrac deep convolutional neural network. 2020.
- Alqudah A, Qazan S, Alquran H, Qasmieh I, Alqudah A. Covid-2019 detection using X-ray images and artificial intelligence hybrid models. 2020;6:168. <https://doi.org/10.5455/jjee.204-1585312246>.
- Chowdhury MEH, Rahman T, Khandakar A, Mazhar R, Kadir MA, Mahbub ZB, Islam KR, Khan MS, Iqbal A, Al-Emadi N, Reaz MBI, Islam TI. Can ai help in screening viral and covid-19 pneumonia?. 2020.
- Chung A. Actualmed covid-19 chest x-ray dataset initiative. <https://github.com/agchung/Actualmed-COVID-chestxray-dataset>. 2020a.
- Chung A. Figure 1 covid-19 chest x-ray dataset initiative. <https://github.com/agchung/Figure1-COVID-chestxray-dataset>. 2020b.
- Cohen JP, Morrison P, Dao L. Covid-19 image data collection. arXiv:2003.11597. <https://github.com/ieee8023/covid-chestxray-dataset>. 2020.
- Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L. ImageNet: A large-scale hierarchical image database. In: CVPR09; 2009.
- Ekagra R, Soumava P, Siddharth K, Aupendu K, Ramanathan S, Debdoot S. Jointly learning convolutional representations to compress radiological images and classify thoracic diseases in the compressed domain. 12. 2018.
- Ghoshal B, Tucker A. Estimating uncertainty and interpretability in deep learning for coronavirus (covid-19) detection. arXiv:2003.10769. 2020.
- He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: 2016 IEEE Conference on computer vision and pattern recognition (CVPR). <https://doi.org/10.1109/cvpr.2016.90>; 2016.

- Hemdan EE-D, Shouman MA, Karar ME. Covidx-net: a framework of deep learning classifiers to diagnose covid-19 in x-ray images. arXiv:2003.11055. 2020.
- Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, Andreetto M, Adam H. Mobilenets: efficient convolutional neural networks for mobile vision applications. 2017.
- Huang C, Wang Y, Li X, Ren L, Zhao J, Hu Y, Zhang L, Fan G, Xu J, Gu X, Cheng Z, Yu T, Xia J, Wei Y, Wu W, Xie X, Yin W, Li H, Liu M, Cao B. Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. *The Lancet*. 2020;395:01.
- Jadhav SM, Yadav SS. Deep convolutional neural network based medical image classification for disease diagnosis. *J Big Data* 6. <https://doi.org/10.1186/s40537-019-0276-2>. 2019.
- Kent C. Different paths to the same destination: screening for covid-19. <https://www.medicaldevice-network.com/features/types-of-covid-19-test-antibody-pcr-antigen/>. 2020.
- Kermany DS, Zhang K, Goldbaum MH. Labeled optical coherence tomography (oct) and chest x-ray images for classification. 2018.
- Li Y, Poblens E, Dagunts D, Covington B, Devon B, Lyman K. Learning to diagnose from scratch by exploiting dependencies among labels. 2017.
- Lundervold AS, Lundervold A. An overview of deep learning in medical imaging focusing on mri. *Zeitschrift für Medizinische Physik*. 2019;29(2):102–127. ISSN 0939-3889. <http://www.sciencedirect.com/science/article/pii/S0939388918301181>. Special Issue: Deep Learning in Medical Physics.
- Mangal A, Kalia S, Rajgopal H, Rangarajan K, Namboodiri V, Banerjee S, Arora C. Covidaid: Covid-19 detection using chest x-ray. 2020.
- Narin A, Kaya C, Pamuk Z. Automatic detection of coronavirus disease (covid-19) using x-ray images and deep convolutional neural networks. arXiv:2003.10849. 2020.
- Ng M-Y, Lee EYP, Yang J, Yang F, Li X, Wang H, Lui MM, Lo CS-Y, Leung B, Khong P-L, Hui CK-M, Yuen K-Y, Kuo MD. Imaging profile of the covid-19 infection radiologic findings and literature review. *Radiol Cardiothoracic Imaging*. 2020;2(1):e200034.
- Rajpurkar P, Irvin J, Zhu K, Yang B, Mehta H, Duan T, Ding D, Bagul A, Langlotz C, Shpanskaya K, Lungren MP, Ng AY. Radiologist-level pneumonia detection on chest x-rays with deep learning. 2017.
- Rsna pneumonia detection challenge. <https://www.kaggle.com/rsna-pneumonia-detection-challenge/data>. 2018.
- Salman FM, Abu-Naser SS, Alajrami E, Abu-Nasser BS, Alashqar BAM. Covid-19 detection using artificial intelligence. 2020.
- Wang L, Wong A. Covid-net: a tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. 2020.
- Wang X, Peng Y, Lu L, Lu Z, Bagheri M, Summers RM. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: 2017 IEEE Conference on computer vision and pattern recognition (CVPR), pp 3462–3471; 2017.
- Wang X, Peng Y, Le L, Zhiyong L, Bagheri M, Summers RM. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: 2017 IEEE Conference on computer vision and pattern recognition (CVPR). <https://doi.org/10.1109/CVPR.2017.369>; 2017.
- Weng X, Zhuang N, Tian J, Liu Y. Chexnet for classification and localization of thoracic diseases. <https://github.com/arnoweng/CheXNet/>. 2017.

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.