



# CovidDiagnosis: Deep Diagnosis of COVID-19 Patients Using Chest X-Rays

Kushagra Mahajan<sup>1</sup>(✉), Monika Sharma<sup>1</sup>, Lovekesh Vig<sup>1</sup>, Rishab Khincha<sup>2</sup>,  
Soundarya Krishnan<sup>2</sup>, Adithya Niranjana<sup>2</sup>, Tirtharaj Dash<sup>2</sup>,  
Ashwin Srinivasan<sup>2</sup>, and Gautam Shroff<sup>1</sup>

<sup>1</sup> TCS Research, New Delhi, India

{kushagra.mahajan,monika.sharma1,lovekesh.vig,gautam.shroff}@tcs.com

<sup>2</sup> BITS Pilani, Sancoale, Goa, India

{f20160517,f20160472,f20160444,tirtharaj,ashwin}@goa.bits-pilani.ac.in

**Abstract.** As the COVID-19 pandemic threatens to overwhelm health-care systems across the world, there is a need for reducing the burden on medical staff via automated systems for patient screening. Given the limited availability of testing kits with long turn-around test times and the exponentially increasing number of COVID-19 positive cases, X-rays offer an additional cheap and fast modality for screening COVID-19 positive patients, especially for patients exhibiting respiratory symptoms. In this paper, we propose a solution based on a combination of deep learning and radiomic features for assisting radiologists during the diagnosis of COVID-19. The proposed system of *CovidDiagnosis* takes a chest X-ray image and passes it through a pipeline comprising of a model for lung isolation, followed by classification of the lung regions into four disease classes, namely Healthy, Pneumonia, Tuberculosis and COVID-19. To assist our classification framework, we employ embeddings of disease symptoms produced by the CheXNet network by creating an ensemble. The proposed approach gives remarkable classification results on publicly available datasets of chest X-rays. Additionally, the system produces visualization maps which highlight the symptoms responsible for producing the classification decisions. This provides trustworthy and interpretable decisions to radiologists for the clinical deployment of CovidDiagnosis. Further, we calibrate our network using temperature scaling to give confidence scores which are representative of true correctness likelihood.

**Keywords:** COVID-19 · X-rays · Coronavirus · CovidDiagnosis · Radiomics

## 1 Introduction

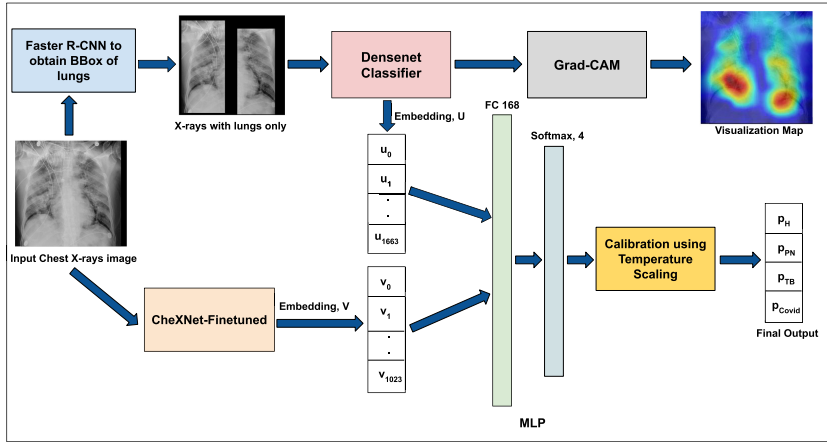
The novel Coronavirus 2019 (COVID-19) originated in December, 2019 in China and soon spread across the entire world as a pandemic. Its spread has left a devastating mark on public health and global economy. The need of the hour is to

detect the COVID-19 positive cases as quickly as possible to stop further spread of the infection. The gold standard diagnostic criteria for COVID-19 is a real time reverse transcript polymerase chain reaction (RT-PCR) [7] test kit which is expensive, time-consuming and in short supply. In addition, these tests do suffer to produce false-negative and partial positive results at the onset of COVID-19 infection because they are unable to capture the minor viral infections in early stages [15, 29]. Such patients when go undetected are a risk to other people. Several studies [4, 17, 20, 31] published online claimed that Chest CT / X-rays can help in early screening of such suspected cases as they are cheaper and faster to test. While CT-Scans are more accurate, they are also more expensive, require more time and a more elaborate set up to ensure a contact free environment. X-rays, while not as sensitive as CT-Scans are able to perform with high accuracy in patients suffering from respiratory symptoms, which constitute the majority of patients visiting hospitals. Moreover, the X-ray technology is readily available in almost all parts of the world, and X-ray machines are already installed at public places like airports and railways stations and can be re-purposed to screen for characteristic COVID-19 radiomic features such as ground glass opacities and consolidations which are likely to have a peripheral distribution with bilateral, multi-focal lower lung involvement [13, 16].

Radiologists are able to read X-rays and identify symptoms specific to COVID-19. However, manual diagnosis using X-rays is a time-consuming and effort-intensive task and there are not enough experienced medical professionals to cater to the exponentially increasing numbers of suspected cases around the world. This necessitates the automation of the COVID-19 diagnosis from Chest X-ray images using AI-based deep learning approaches which can facilitate quick identification of high risk COVID-19 patients who can then be further prescribed for RT-PCR tests. Meanwhile, such patients can be put into isolation wards to help reduce the spread of the coronavirus disease.

While there are burgeoning efforts on automation of COVID-19 diagnosis using Chest X-rays [3, 8, 18, 24, 27], we have contextualized our method to distinguish COVID-19 from commonly occurring respiratory problems in South-East Asia, such as Tuberculosis and Pneumonia. In this paper, we utilize Chest X-rays for classifying patients into four pathology categories namely, healthy, pneumonia, tuberculosis and COVID-19 positive. The proposed method first extracts the lung-region bounding boxes from the chest X-ray images using a trained Faster-RCNN [22] network. This step is performed to focus the subsequent classifiers attention on lung areas. Subsequently, an X-ray image containing only lung-regions is fed as input to a DenseNet-169 [10] network to yield a probability distribution over the four classes. Since, Chest X-rays contain several visually discernible symptoms such as ground glass opacities, consolidations, fibrosis, and pneumothorax which are distinctive features for detecting COVID-19 from other pathologies, we finetune the existing CheXNet model [21] on a combination of classes from the NIH dataset [28] and the Stanford CheXpert [11] dataset. We combine the embeddings from the second last layer of DenseNet-169 and finetuned CheXNet and then, train an MLP classifier to obtain the final output. To

enhance trust in the model predictions, we isolate the lung regions and generate visualization maps using Grad-CAM [23] to highlight the affected lung regions in the X-ray image. We also calibrate our network using temperature scaling [9] to obtain accurate confidence scores for a prediction. A well calibrated model is critical for medical applications as it allows medical professionals identify and examine only those cases that are confusing for the model.



**Fig. 1.** Figure showing pipeline of our proposed method - CovidDiagnosis.

To summarize, we make following contributions in the paper:

- We propose a deep learning based automated method for COVID-19 diagnosis using Chest X-rays which is calibrated using temperature scaling [9] and outputs accurate and reliable confidence scores for predictions.
- We propose a diagnostic pipeline comprising first of lung isolation using Faster-RCNN network [22] allowing subsequent models to focus entirely on disease symptoms present in the lungs and not pay attention to other elements in the image.
- We finetune CheXNet [21] on classes from NIH and Stanford datasets, and further use this network’s embeddings to capture information about visual disease symptoms useful for discriminating COVID-19 from diseases such as tuberculosis and pneumonia for making the final diagnosis. The final prediction is made by using a combined embedding from ChexNet and a DenseNet model proposed for classification.
- We also illustrate the effectiveness of our proposed approach by creating activation maps using Grad-CAM [23] to highlight affected regions of the lungs which can prove helpful to the radiologist to validate the model.
- We evaluate CovidDiagnosis on a publicly available Chest X-rays dataset and compare the performance against state-of-the-art networks such as Covid-Net [27] and CovidAID [18]. We show empirically that our pipeline outperforms competing approaches.

The remainder of the paper is organized as follows: Sect. 2 will describe the work previously done on COVID-19 diagnosis using X-ray images and how our approach is different from prior work. Next, we provide a description of our proposed method in Sect. 3. This is followed by brief detail on a Chest X-ray images dataset used for training and evaluation purposes. Subsequently, Sect. 5 provide details of the training, experiments conducted, their results and discussions on them. In the end, Sect. 6 will conclude the work with future avenues in this field.

## 2 Related Work

In recent times, deep learning has made significant strides in medical image classification and segmentation [1, 5, 19, 30] in addition to being applied to standard image processing, computer vision and natural language processing tasks. Several existing deep learning networks in literature have been proposed to identify various thoracic diseases from Chest X-rays like pneumonia, pneumothorax, fibrosis etc. [14, 21, 25]. Since, patients suffering from COVID-19 develop pneumonia and certain other infectious symptoms in the lung areas of Chest X-rays, it may be convenient to use X-rays for automation of COVID-19 screening.

With the sudden spike in the number of COVID-19 infected patients, there have been a number of publications addressing the problem of automatic diagnosis so as to quickly and effectively isolate infected patients and curb the spread of the virus. Authors in [27] proposed a deep neural network based Covid-Net to distinguish viral and bacteria pneumonia from COVID-19 patients and have also released their evaluation dataset COVIDx. Another paper utilized pre-trained CheXNet to finetune the network to classify X-ray images and named the network CovidAID [18]. Tulin et al. [24] developed a DarkNet based model to classify X-ray images into COVID-19, No-Findings and Pneumonia. Sanhita et al. [3] employ transfer learning based convolutional neural network for classification into four classes, e.g, normal, other disease, pneumonia and COVID-19. A slightly different approach was adopted by authors in [26] where they trained deep learning models like MobileNetV2, SqueezeNet on the created stacked dataset, and the features obtained by the models were processed using the Social Mimic optimization method. In the next step, efficient features were combined and classified using an SVM. A comparative study of the use of the recent deep learning models (VGG16, VGG19, DenseNet201, InceptionResNetV2, InceptionV3, Resnet50, and MobileNetV2) to deal with detection and classification of COVID-19 pneumonia is presented in paper [2].

We found that prior-art mostly consist of deep learning methods trying to distinguish COVID-19 from healthy cases and pneumonia in Chest X-rays. However, in our paper, we also propose to contextualize the solution to distinguish COVID-19 from commonly occurring respiratory problem in South-East Asia such as Tuberculosis. Moreover, the methods described above simply feed the Chest X-ray images into deep networks for classification and hence, it is uncertain if the prior networks are making decisions based on relevant information from Chest X-ray images and focusing on the right visual features. Therefore, we

build a pipeline which first localizes the regions of interest (lungs) using a Faster-RCNN [22]. We also supplement our network with disease symptom embeddings obtained from CheXNet. Additionally, we have not seen any COVID-19 in the literature which is calibrated to give output probabilities with confidence scores representative of prediction correctness likelihood.

### 3 Proposed Approach: CovidDiagnosis

Figure 1 shows the overall architecture of the proposed CovidDiagnosis method for Chest X-rays. It consists of the following modules for processing X-ray images of patients to classify them into four classes - healthy, pneumonia, tuberculosis and COVID-19 positive:

- **Lung Isolation using Faster-RCNN:** As a first step, we train a Faster-RCNN [22] network for identifying bounding boxes (Bbox) of lungs present in Chest X-rays. The training data for the lungs Bbox is obtained by using Lung-Finder<sup>1</sup> which uses basic image processing features such as HAAR, LBP and HOG to detect the left and the right lungs from an X-ray image. The detections produced by HAAR features are better when compared with those produced by HOG determined through visual inspection, and hence we use bounding boxes from HAAR features for training the faster-rcnn. We manually verified the resultant Bboxes for having correct lung region and discarded those set of images for which lungs Bbox is not correct. We train a faster-rcnn on the lung Bboxes using VGG16 as the backbone network.
- **DenseNet Classifier:** DenseNet [10] is a convolutional network which has  $n$  layers and  $n \times (n + 1)/2$  direct connections. This means that each subsequent layer has feature maps of all preceding layers as input. DenseNets have fewer parameters to train and dense connections have a regularizing effect which reduces over-fitting on tasks with lesser training data. We use DenseNet-169 to train a classifier on a 3-channel input image comprising left and right isolated lung (channels 1 and 2) and a complete Chest X-ray image (channel 3). The underlying idea behind using lung-isolation is to explicitly enable the network to focus its attention on the lung regions for making a classification decision. The classifier is trained to classify input Chest X-ray images into one of four categories - healthy, pneumonia, tuberculosis and COVID-19 positive. We utilize the embeddings  $U = (u_0, u_1, \dots, u_{1663})$  from the second last layer of the trained DenseNet classifier for further processing.
- **CheXNet Finetuning for disease symptoms:** To enhance the discriminatory capabilities of our classifier, we propose to explicitly provide information about disease symptoms present in Chest X-rays such as opacities, consolidations, mass, fibrosis, pneumothorax etc. in the form of an embedding vector. This additional information assists the classifier by drawing correlations between our four disease classes and the disease symptoms displayed

---

<sup>1</sup> Lungs-Finder: <https://github.com/xiaoyongzhu/lungs-finder>.

by the X-ray. This, in turn, helps to ignore any bias present in the chest X-rays dataset obtained from different sources as is commonly the case when images for classes are gathered from different data sources. We use the NIH dataset consisting of 14 classes that represent visual disease symptoms present in chest X-rays. Since, the NIH dataset does not have ‘opacities’ as one of the classes which is an important COVID-19 positive feature, we include the samples of opacities class from Stanford CheXpert dataset and subsequently finetune the CheXNet [21] model for these 15 classes of disease symptoms. Further, the finetuned CheXNet is used to obtain second last layer embeddings  $V = (v_0, v_1, \dots, v_{1023})$  of Chest X-rays images which are used for processing in the next module.

- **MLP classifier:** The embeddings from DenseNet-169 and CheXNet are concatenated to train a 2-layer MLP for final classification of X-rays images into four classes - healthy, pneumonia, tuberculosis and COVID-19 positive. The MLP comprises of (168, 4) hidden units, and is fed an input of 2688 dimensional vector  $W = (u_0, u_1, \dots, u_{1663}, v_0, v_1, \dots, v_{023})$ .
- **Calibration of Network:** We perform calibration of our proposed network so that radiologists can rely on output predictions of our network and the solution can be deployed for clinical use. The simplest way to visualize calibration is to plot the accuracy as a function of confidence and for a well-calibrated network, the plot should be an identity function. The reliability diagram of our proposed architecture is shown in Fig. 2. In order to calibrate the network, we have used Temperature scaling [9] which divides the logit vector obtained from the classification network by a learned scalar parameter  $T$  as follows:

$$P(\hat{y}) = \frac{e^{\mathbf{z}/T}}{\sum_j e^{z_j/T}} \quad (1)$$

where  $\hat{y}$  and  $\mathbf{z}$  is the output prediction and logit, respectively. The parameter  $T$  is learned on a validation set where  $T$  is chosen to minimize negative log-likelihood. In essence, the temperature scaling technique softens the network outputs, thereby making the network less confident and in turn, making the confidence scores reflect the true correctness probabilities.

- **Activation Maps Generation:** An extensive qualitative analysis is essential for determining whether our diagnostics approach is looking at the right portions of the image for determining the four disease classes. These correct regions were determined through consultations with a team of doctors. To obtain the regions of interest where our network is focusing its attention, we make use of Grad-CAM [23]. The activation maps highlight the regions where abnormalities such as consolidations and opacities etc. are present which can aid radiologists in deeper investigation of suspected cases.

## 4 Dataset

We utilize Chest X-rays of four classes - healthy, pneumonia, tuberculosis and COVID-19 from publicly available datasets for diagnosis of COVID-19. X-ray

samples for COVID-19 positive patients are taken from the X-ray image database made available by Cohen JP et al. [6]. X-ray images for healthy and tuberculosis patients are obtained from Pulmonary Chest X-rays dataset [12]. The NIH dataset is used for pneumonia X-ray images. The dataset is divided into train (823), validation (181) and test sets (344) with number of samples for each class as follows: Train-set (COVID-19 - 153, Healthy - 240, Pneumonia - 190 and Tuberculosis - 240); Val-set ( COVID-19 - 27, Healthy - 54, Pneumonia - 44 and Tuberculosis - 55); Test-set (COVID-19 - 46, Healthy - 110, Pneumonia - 88 and Tuberculosis - 100).

## 5 Experimental Results and Discussions

### 5.1 Training Details

The input to our classifier consists of 3 channels: the first channel is the left lung image (with the rest of the image blacked out), the second channel is the right lung image (with the rest of the image blacked out), and the third channel is the full image. We have experimented with images of size  $224 \times 224$  and  $448 \times 448$  for training CovidDiagnosis. We apply augmentations to all the above datasets by applying random transformations like rotation ( $-20^\circ$  to  $+20^\circ$ ), scaling ( $-10\%$  to  $+10\%$ ), normalization, and horizontal and vertical translation. The chest region is always present at the centre of the image and hence, the small translations do not push the important regions out of the image. We utilize DenseNet-169 pre-trained on ImageNet for finetuning on the four disease classes with SGD optimizer having an initial learning rate of  $1e-3$  and momentum of 0.9. We apply softmax activation function in the final layer of our DenseNet classifier to obtain the probabilities of predictions.

### 5.2 Baseline Networks

To show the effectiveness of our proposed CovidDiagnosis method, we compare its performance with two prior deep networks namely, Covid-Net [27] and CovidAID [18].

- **Covid-Net**: The architecture of Covid-Net is based on a residual projection-expansion-projection-extension (PEPX) design pattern and comprises of a mixture of convolutional layers with diverse kernel sizes and grouping configurations. We have used the original architecture and training configuration, as given by authors [27].
- **CovidAID**: Mangal et al. [18] proposed CovidAID which simply finetunes the CheXNet network for classifying Chest X-ray images into four classes - Normal, Bacterial Pneumonia, Viral Pneumonia and COVID-19 using a sigmoid activation. We used the same training configuration for comparison.

We trained all the three networks Covid-Net, CovidAID and CovidDiagnosis for 3-way classification of X-rays into healthy, pneumonia and COVID-19 positive from our Chest X-rays dataset, as mentioned in Sect. 4 for performance evaluation and comparison.

**Table 1.** Performance impact of different modules of CovidDiagnosis on AUC and Accuracy values of X-rays pathology classification

	Method	AUC	Acc
I(224 * 224)	DenseNet	0.9699	0.8799
	DenseNet + LungIsolation(LI)	0.9823	0.9040
I(448 * 448) w/ aug	DenseNet	0.9798	0.8924
	DenseNet + LungIsolation(LI)	<b>0.9885</b>	<b>0.9069</b>
	DenseNet + LI + CheXNet + MLP	0.9694	<b>0.9157</b>
I(448 * 448) w/o aug	DenseNet + LungIsolation(LI)	0.9462	0.8488
	DenseNet + LI + CheXNet + MLP	0.9822	0.8895

### 5.3 Results and Discussion

First, we present the results of introducing different modules in our proposed method CovidDiagnosis in Table 1. Initially, we train the DenseNet-169 classifier for classifying the entire Chest X-rays image of size  $224 \times 224$  as input and as can be seen in Table 1, we achieve an accuracy of 87.99% and AUC value of 0.9699. In the next experiment, we use lung isolation to explicitly enable DenseNet to focus on right areas in an X-ray image i.e., lungs to learn relevant discriminating features for classification. We can clearly see that lung isolation gives a boost in accuracy from 87.99% to 90.40%. Further, we use images of size  $448 \times 448$  for training DenseNet with and without lung isolation and obtain an accuracy of 89.24% and 90.69%, and AUC of 0.9798 and 0.9885 respectively. As is evident from Table 1, lung isolation experiments for images of sizes  $224 \times 224$  and  $448 \times 448$  give almost equal performance. Thus, we can conclude that lung isolation is beneficial for getting good performance on low-resolution images as it helps the network narrow down its field of view, and allows it to focus its attention entirely on lung regions.

**Table 2.** Class-wise different performance measures for CovidDiagnosis (using DenseNet + LI + CheXNet + MLP) on image size  $448 \times 448$ . PN and TB represents pneumonia and tuberculosis, respectively

<i>Confusion matrix</i>	<i>COVID</i>	<i>Healthy</i>	<i>PN</i>	<i>TB</i>	<i>Precision</i>	<i>Recall</i>	<i>F1-Score</i>
<i>COVID</i>	<b>45</b>	0	1	0	0.98	0.98	0.98
<i>Healthy</i>	0	<b>102</b>	0	8	0.88	0.93	0.90
<i>PN</i>	1	0	<b>84</b>	3	0.97	0.95	0.96
<i>TB</i>	0	14	2	<b>84</b>	0.88	0.84	0.86

Subsequently, we concatenate embeddings from DenseNet and CheXNet to train a 2-layer MLP classifier and observe that it gives an improvement in classification accuracy (91.57%). This implies that adding explicit information




about various abnormalities present in Chest X-rays in the form of an embedding vector from CheXNet enhances the discriminatory power of CovidDiagnosis and improves the classification performance. Since, the network performance is already high in our case, there is very limited room for improvement on this small dataset. In order to show the effectiveness of CheXNet embeddings as an additional feature, we perform an experiment by training DenseNet on our dataset without using augmentation techniques during training. Here, we can see that CheXNet embeddings improved classification accuracy from 84.88% to 88.95%. We also present the confusion matrix, precision, recall and F1-score values of CovidDiagnosis for image sizes  $448 \times 448$  to illustrate the class-wise performance in Table 2. We observe that CovidDiagnosis identifies COVID-19 positive samples with 0.98 for all the measures - precision, recall and F1-score.

The comparison of the performance of CovidDiagnosis against two baseline networks - Covid-Net and CovidAID is presented in Table 3. It is evident from the accuracy, AUC values and the confusion matrices that CovidDiagnosis outperforms both of the other methods for the 3-way classification problem. There is a boost of approx. 3% and 2% in classification accuracy of CovidDiagnosis when compared with Covid-Net and CovidAID respectively.

**Table 3.** Performance comparison with state-of-the-art networks for diagnosis of COVID-19 from Chest X-rays for 3-way classification


<i>Confusion matrix</i>		COVID	Healthy	Pneumonia	<b>AUC</b>	<b>Acc</b>
<b>Covid-Net</b>	COVID	<b>37</b>	0	9	0.9633	0.8934
	Healthy	3	<b>107</b>	0		
	Pneumonia	4	10	<b>74</b>		
<b>CovidAID</b>	COVID	<b>27</b>	1	18	0.9721	0.9061
	Healthy	0	<b>110</b>	0		
	Pneumonia	5	0	<b>83</b>		
<b>Covid-Diagnosis</b>	COVID	<b>46</b>	0	0	<b>0.9794</b>	<b>0.9221</b>
	Healthy	0	<b>107</b>	3		
	Pneumonia	2	14	<b>72</b>		

Figure 2 shows the reliability diagram of CovidDiagnosis for classifying Chest X-rays into 4 classes. The plot in Fig. 2(a) shows the confidence scores before calibration and we can see that the network is not well-calibrated and is over-confident in predicting outputs. But after using temperature scaling for network calibration, the reliability diagram is close to an identity plot as shown in Fig. 2(b) indicating that the network is not over-confident in its predictions. This builds up trust with medical practitioners to adopt our system for its practical use in hospitals, clinics, and public places like airports etc.



**Fig. 2.** Plots showing reliability diagrams (a) before and (b) after calibration of Covid-Diagnosis using temperature scaling [9].

We create activation maps for COVID-19 positive Chest X-rays using GRAD-Cam [23], as shown in Fig. 3. The activation maps indicate that the portions of the lungs that are actively examined by the model are in agreement with where abnormalities such as consolidations and opacities are present, showing that our network is indeed looking at the right locations for learning discriminating features for classification.



**Fig. 3.** Visualization maps of COVID-19 positive X-rays of patients using GRAD-Cam [23] highlighting regions of abnormalities.

## 6 Conclusion

We propose a system named CovidDiagnosis for segregating suspected COVID-19 positive patients having high confidence scores, thereby helping to reduce the spread of COVID-19 disease. We demonstrate that utilizing lung isolation and combining embeddings from CheXNet of visual disease symptoms in Chest X-rays enhances classification accuracy and enables the network to learn relevant discriminating features by focusing on correct lung regions. Thus, it produces

improved classification performance compared to other existing COVID-19 X-ray classification approaches. The proposed method is well-calibrated via temperature scaling which can prove beneficial to radiologists as the predicted probabilities reflect the true correctness likelihood. We also aid radiologists by providing activation maps highlighting regions of disease symptoms for deeper investigation. Going forward, we wish to apply few-shot techniques such as meta-learning to train the network on limited available data and eventually, evaluate on larger and more diverse datasets.

## References

1. Alom, M.Z., Yakopcic, C., Taha, T.M., Asari, V.K.: Nuclei segmentation with recurrent residual convolutional neural networks based u-net (r2u-net). In: NAECON 2018 - IEEE National Aerospace and Electronics Conference, pp. 228–233, July 2018. <https://doi.org/10.1109/NAECON.2018.8556686>
2. Asnaoui, K.E., Chawki, Y.: Using x-ray images and deep learning for automated detection of coronavirus disease. *J. Biomolecular Structure Dyn.* 1–12 (2020). <https://doi.org/10.1080/07391102.2020.1767212>
3. Basu, S., Mitra, S., Saha, N.: Deep learning for screening covid-19 using chest x-ray images. arXiv (2020), <https://arxiv.org/abs/2004.10507>
4. Bernheim, A., et al.: Chest ct findings in coronavirus disease-19 (covid-19): relationship to duration of infection. *Radiology* (2020). <https://pubs.rsna.org/doi/10.1148/radiol.2020200463>
5. Chen, S., Ma, K., Zheng, Y.: Med3d: Transfer learning for 3d medical image analysis. ArXiv abs/1904.00625 (2019)
6. Cohen, J.P., Morrison, P., Dao, L., Roth, K., Duong, T.Q., Ghassemi, M.: Covid-19 image data collection: Prospective predictions are the future. arXiv 2006.11988 (2020). <https://github.com/ieee8023/covid-chestxray-dataset>
7. Emery, S.L., et al.: Real-time reverse transcription-polymerase chain reaction assay for sars-associated coronavirus. *Emerging Infectious Diseases* (2004). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3322901/>
8. Gozes, O., et al.: Rapid AI development cycle for the coronavirus (covid-19) pandemic: initial results for automated detection & patient monitoring using deep learning ct image analysis (2020)
9. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: Proceedings of the 34th International Conference on Machine Learning, vol. 70. pp. 1321–1330. ICML 2017, JMLR.org (2017)
10. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2261–2269 (2017)
11. Irvin, J., et al.: Chexpert: a large chest radiograph dataset with uncertainty labels and expert comparison. CoRR abs/1901.07031 (2019). <http://arxiv.org/abs/1901.07031>
12. Jaeger, S., Candemir, S., Antani, S., Wang, Y.X.J., Lu, P.X., Thoma, G.: Two public chest x-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative Imag. Med. Surgery* 4(6), 475 (2014)
13. Kong, W., Agarwal, P.P.: Chest imaging appearance of covid-19 infection. *Radiology* (2020). <https://pubs.rsna.org/doi/10.1148/ryct.2020200028>

14. Kumar, P., Grewal, M., Srivastava, M.M.: Boosted cascaded convnets for multilabel classification of thoracic diseases in chest radiographs. In: Campilho, A., Karray, F., ter Haar Romeny, B. (eds.) *Image Analysis and Recognition*, pp. 546–552. Springer, Cham (2018)
15. Lan, L., Xu, D., Ye, G., Xia, C., Wang, S., Li, Y., Xu, H.: Positive RT-PCR test results in patients recovered from COVID-19. *JAMA* **323**(15), 1502–1503 (2020). <https://doi.org/10.1001/jama.2020.2783>
16. Lee, E.Y.P., Ng, M.Y., Khong, P.L.: COVID-19 pneumonia: what has CT taught us? *The Lancet Infectious Diseases* **20**(4), 384–385 (2020). <http://www.sciencedirect.com/science/article/pii/S1473309920301341>
17. Long, C., et al.: Diagnosis of the coronavirus disease (covid-19): RRT-PCR or CT? *Euro. J. Radiol.* **126**, 108961 (2020). <http://www.sciencedirect.com/science/article/pii/S0720048X20301509>
18. Mangal, A., Kalia, S., Rajgopal, H., Rangarajan, K., Namboodiri, V., Banerjee, S., Arora, C.: Covidaid: COVID-19 detection using chest x-ray (2020)
19. Oktay, O., et al.: Attention u-net: learning where to look for the pancreas. *CoRR abs/1804.03999* (2018). <http://arxiv.org/abs/1804.03999>
20. Eng Pan, Ye, T., et al.: Time course of lung changes at chest ct during recovery from coronavirus disease 2019 (COVID-19). *Radiology* (2020). <https://doi.org/10.1148/radiol.2020200370>
21. Rajpurkar, P., et al.: Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *CoRR abs/1711.05225* (2017). <http://arxiv.org/abs/1711.05225>
22. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(6), 1137–1149 (2017). <https://doi.org/10.1109/TPAMI.2016.2577031>
23. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: visual explanations from deep networks via gradient-based localization. In: *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 618–626 (2017)
24. Ozturk, T., et al.: Automated detection of COVID-19 cases using deep neural networks with x-ray images. *Comput. Biol. Med.* (2020). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7187882/>
25. Tang, Y., Tang, Y., Xiao, J., Summers, R.M.: Xlsor: A robust and accurate lung segmentor on chest x-rays using criss-cross attention and customized radiorealistic abnormalities generation. In: *MIDL* (2019)
26. Toğaçar, M., Ergen, B., Cömert, Z.: Covid-19 detection using deep learning models to exploit social mimic optimization and structured chest x-ray images using fuzzy color and stacking approaches. *Comput. Biol. Med.* **121**, 103805 (2020). <http://www.sciencedirect.com/science/article/pii/S0010482520301736>
27. 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)
28. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: 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*, pp. 2097–2106 (2017)
29. Xie, X., Zhong, Z., Zhao, W., Zheng, C., Wang, F., Liu, J.: Chest ct for typical 2019-ncov pneumonia: relationship to negative rt-pcr testing. *Radiology* (2020). <https://pubs.rsna.org/doi/10.1148/radiol.2020200343>

30. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: a review. *Med. Image Anal.* **58**, 101552 (2019). <http://www.sciencedirect.com/science/article/pii/S1361841518308430>
31. Zu, Z.Y., Jiang, M.D., Xu, P.P., Chen, W., Ni, Q.Q., Lu, G.M., Zhang, L.J.: Coronavirus disease 2019 (covid-19): a perspective from china. *Radiology* (2020). <https://pubs.rsna.org/doi/10.1148/radiol.2020200490>