U-Net Based Chest X-ray Segmentation with Ensemble Classification for Covid-19 and Pneumonia

https://doi.org/10.3991/ijoe.v18i07.30807

K. A. S. H. Kumarasinghe, S. L. Kolonne, K. C. M. Fernando, D. Meedeniya(✉)
Department of Computer Science and Engineering, University of Moratuwa, Moratuwa, Sri Lanka
dulanim@cse.mrt.ac.lk

Abstract—Respiratory diseases have been known to be a main cause of death worldwide. Pneumonia and Covid-19 are two of the dominant diseases. Several deep learning based studies are available in the literature that classifies infection conditions in chest X-ray images. In addition, image segmentation has been also applied to obtain promising results in deep learning approaches. This paper focuses on using a modified version of the U-Net architecture to conduct segmentation on chest X-rays and then use segmented images for classification to assess the impact on the performance. We achieved an Intersection over Union of 93.53% with the proposed modified U-Net architecture and achieved 99.36% accuracy on segmentation aided ensemble classification.

Keywords—segmentation, U-Net, classification, CNN, chest X-rays

1 Introduction

Pneumonia is an acute respiratory disease that causes inflammation in the lungs and is a prominent cause of death, especially among children. COVID-19, declared a pandemic in early 2020, is also a severe acute respiratory syndrome that can lead to fatal Pneumonia conditions [1][2]. Chest radiographs are widely used in the medical domain to diagnose both these conditions, due to their cost-effectiveness and non-invasive nature. However, identifying abnormalities by reading chest X-rays can be challenging for a non-expert. Moreover, distinguishing between normal and COVID-19 caused Pneumonia can also be difficult due to the identicalness between them.

Many studies have been conducted on using different classification techniques for diagnosis in the medical domain [3]. Our previous work has proposed Convolutional Neural Network (CNN) architecture based classification of chest X-ray images to diagnose COVID-19 Pneumonia and normal Pneumonia [4][5]. At the same time, image segmentation is also a widely discussed area in this domain [6]. Lung segmentation is a process that can identify the regions as well as the boundaries of the lung area from the surrounding thoracic tissue and is regarded as an important initial step in pulmonary image processing to perform further computer analysis [7]. This paper proposes a deep learning-based lung segmentation model together with an ensemble for classification to detect COVID-19 and normal Pneumonia in chest X-rays. We present the positive

impact of chest X-ray segmentation for the performance improvements of the classification model based on CNN architectures.

This paper is structured as follows. Section 2 explores the background of this study. Section 3 and Section 4 describe the model design and methodology. Section 5 presents the results and Section 6 provides the discussion. Finally, Section 7 concludes the paper.

2 Background

2.1 U-Net based medical image segmentation

In computer-aided diagnostics, the segmentation of medical images such as chest X-rays plays an important role [6]. Medical image segmentation becomes quite handy in extracting areas of interest from radiographs and CT images to increase the accuracies and performance of the classification models [7]. One such occasion is to extract the lung regions from chest X-rays and use them to train CNNs to gain better results performance and accuracy wise for classification tasks. In the field of biomedical imaging, U-Net is heavily employed in such image segmentation tasks [8][9]. It is a deep learning-based model that uses the auto-encoder structure with skip connections and provides solutions for issues in traditional segmentation techniques like the sliding window approach. In the sliding window approach, each pixel was examined separately which led to a lot of redundancy due to patches overlapping and it also made the training process more time and resource-intensive. U-Net architecture addresses these issues with fully convolutional networks. The objective of this architecture is to capture both context and localization features. It uses successive contracting layers followed by up-sampling layers that provide outputs with better resolution on the input images. Hence U-Net is much preferred in the medical imaging domain when it comes to making correct decisions.

2.2 U-Net architecture

The U-Net architecture consists of a contracting path and an expansive path [7][8]. In the contracting or the encoder path, which is designed to extract features, the classic architecture of a convolutional network is used for down-sampling. The expansive path or the decoder path on the other hand, which is designed to construct the segmentation map using the features extracted, uses an up-sampling of the feature map concatenated with the feature map from the contracting path cropped correspondingly. Since spatial information can be lost during down-sampling, long skip connections are used to retrieve them by transporting them from the encoder path to the decoder path. The spatial dimensions are reduced, and the channels are increased by the encoder at each layer the opposite happens at the decoder. In the end, the spatial dimensions are restored to make predictions for each pixel in the input image. Fine-grained information is recovered through the prediction by incorporating skip connections in the encoder-decoder architecture.

2.3 Related work with U-Net based chest X-ray segmentation & classification

Several studies have been carried out for the segmentation of chest X-ray images using U-Net based models and have incorporated chest X-ray classification combined with other deep learning architectures and models [3][7]. A study carried out by Teixeira et al. has used the U-Net model for lung segmentation and classification of the chest radiographs was performed using full and segmented CXR images independently employing VGG16, ResNet50V2, and InceptionV3 models for COVID-19 identification [10]. They have modified the U-Net architecture by adding dropout and batch normalization layers to reduce overfitting and the training time. Another study by Narayanan et al. has used a U-Net model with no data augmentation for computationally efficient lung segmentation using chest radiographs [11]. A two-stage framework, based on an adapted U-Net architecture was performed by Rahman et al. in their study for automatic lung segmentation. In the first stage, they have extracted patches from CXR and have trained a modified U-Net architecture to generate an initial segmentation and in the second stage, have carried out a post-processing step to deploy image processing techniques to obtain a clearer final segmentation [12]. A study by Hasan et al. has performed classification of COVID-19, pneumonia, and normal classes using chest X-ray images with the help of an ensemble of DenseNet121, EfficientNetB0, and VGG19 models, as well as segmentation of the affected part in the X-ray images using U-Net architecture based DenseNet103 model [13]. Moreover, they have used the Grad-CAM algorithm to output a heatmap visualization for a given class label. A study comparison is given in Section 6.

3 Model design

3.1 Dataset details

We used the V7-labs COVID-19 X-ray dataset to perform the segmentation of chest X-ray images [14]. This dataset contains chest X-ray images with relevant masks. A sample image from the dataset and its corresponding mask is shown in Figure 1. The dataset originally contained 1602 normal images, 4250 Pneumonia images and 439 Covid-19 images. We applied augmentations to balance the dataset.

Fig. 1. A sample image and its mask

3.2 Pre-processing

Pre-processing methods were applied to the input dataset before proceeding with the training. Both images and masks were resized to 224x224. Then CLAHE (Contrast Limited Adaptive Histogram Equalization) was applied to images to improve the contrast of the chest X-ray images. CLAHE applies histogram equalization in small tiles with contrast limiting. As the final pre-processing step, both images and masks were normalized to a range of [0,1]. Image normalization adjusts the values of the data to a common scale without changing the range. Having a similar data distribution makes convergence faster while training the network.

Although a generalized deep learning model needed a large amount of data, labelled medical image datasets are expensive to produce. Since that data generation methods are needed. Augmentation, Generative Adversarial Network (GAN), random image generation with changing the illumination, are some of the techniques that are used to increase the size of a dataset. Pix2pix and CycleGAN are some Generative Adversarial Network (GAN) based methods that perform synthetic data generation. Even though GAN creates sharp images, it is required to build and train a separate deep learning model for it. Hence instead of using GAN based approaches, we have chosen traditional data augmentation methods like flip, blur, and crop.

3.3 Augmentation

Augmentation is a technique of expanding and diversifying the dataset [15]. In the V7-labs COVID-19 X-ray dataset, although the pneumonia class had 4520 images there were only 439 COVID-19 images and 1602 normal images. For balancing the classes in the dataset augmentation was applied. For the COVID class, nine types of augmentation methods were applied for every image, while two types of augmentation methods were randomly applied for the Normal class. These steps resulted in the database values recorded in Table 1.

Table 1. Dataset classes

	Pneumonia	**Covid-19**	**Normal**	**Total**
Original	4250	439	1602	6291
After class balancing	4250	4390	4806	13446

We have mainly used median blurring, horizontal flip, and CenterCrop and combinations of those techniques. Both horizontal flipping and center crop techniques are geometric transformations that preserve the labels of the images post-transformation. Flipping is a common technique that has proved to be useful with many datasets, while center cropping is a practical technique for image data that have mixed dimensions. By using blurring on images, the model can be trained to work well even with real-world images with imperfections. In addition, to overcome the class imbalance problem, these augmentation techniques have increased the size of the dataset. This diversification method helps to improve the generalizability of the model while reducing overfitting.

3.4 Process flow

The proposed methodology, as shown in Figure 2, mainly consists of segmentation and classification. In the proposed methodology, first, lung segmentation is performed by the U-Net architecture. As the next step, morphological image dilation and erosion techniques are applied on the segmented masks and afterwards the lung area in the masks is extracted. Then the extracted lung images are used to train MobileNetV2, InceptionV3, ResNet50 and Xception models with added top layers. To achieve even better results an ensemble is created by combining the four models.

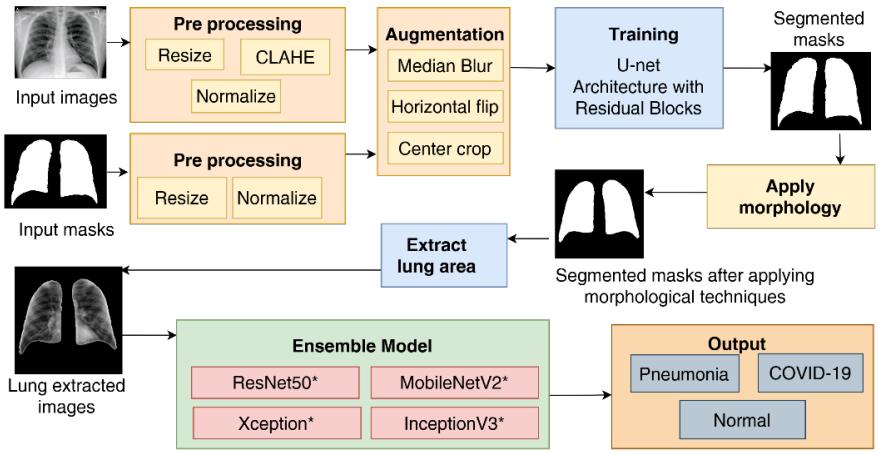

Fig. 2. Process flow

The dataset used to train the U-Net model consisted of images and the corresponding masks. Both the images and the masks were resized to 224 x 224. For the images, CLAHE (Contrast Limited Adaptive Histogram Equalization) was applied. As the final preprocessing step, both images and masks were normalized to a range of [0,1]. After the preprocessing steps, data augmentation was applied to balance the 3 classes. Then the dataset was split to the ratio of 60:20:20 for training, testing and validation, respectively. The convolutional block of the original U-Net model was replaced with residual convolutional blocks as shown in the architecture diagram, and these residual blocks were applied to both the encoder and the decoder of the U-Net architecture. A dropout layer was added to each block to avoid overfitting. ADAM was used as the optimizer and the dice loss as the loss function. The initial learning rate was set to 1e-5 and learning rate reduction was used with a patience of 5 and a factor of 0.1. The batch size was set to 5 and training was done for 20 epochs. The evaluation of the results was done using the Intersection over Union (IoU) score, dice coefficient, precision, and recall evaluation metrics.

3.5 U-Net architecture

We replaced the convolutional block with two convolutional layers in the original U-Net architecture with a residual convolutional block that consists of three convolutional layers and a residual skip connection as shown in Figure 3. Having the number of layers increased in each block means the models will be learning more features while the residual skip connections can help train the network with deep layers to extract low-level features with high precision, while avoiding degradation problem. In addition, a dropout layer with a 10% dropout rate was also added to avoid overfitting.

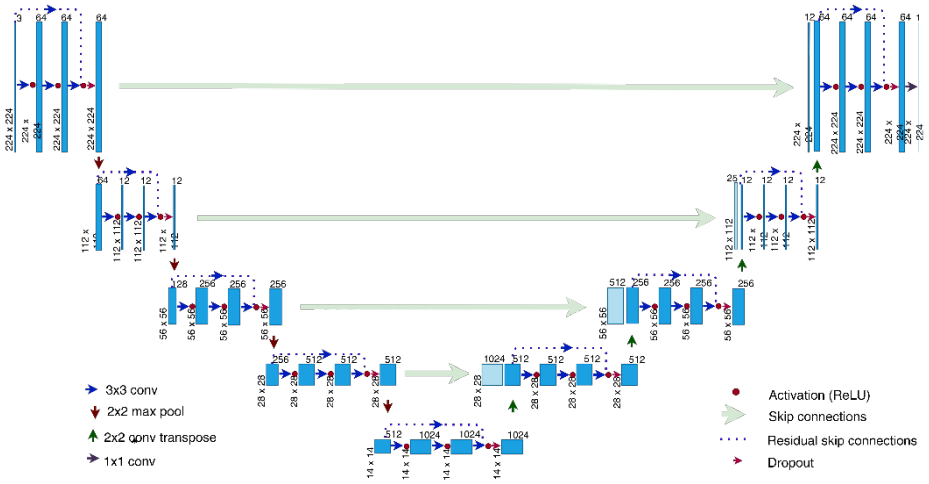

Fig. 3. Modified U-Net architecture

4 Implementation aspects

4.1 Segmentation using original U-Net model

We first performed segmentation for the V7-labs chest X-ray dataset using the original U-Net architecture by training it for 10 epochs, with a batch size of 5, ADAM as the optimizer and dice loss as the loss function. Initially, we have selected 10 epochs because when the number of epochs is higher than necessary, the training model learns patterns that are specific to the sample data to a great extent and makes the model incapable of performing well on a new dataset. The initial learning rate was set to 1e-5 and for better convergence learning rate reduction was used with a patience of 5, a factor of 0.1 and a minimum learning rate of 1e-7. Since the original dataset consists of a total of 6291 images, the dataset was split to train, validation and test datasets using a split ratio of 60:20:20 with 3775 training images, 1258 validation images and 1258 test images.

4.2 Segmentation using modified U-Net model

The modifications made to the architecture included replacing original convolutional blocks with residual convolutional blocks to improve performance and adding a dropout layer with a 0.1 dropout rate to the convolutional blocks to avoid overfitting of the model. Additionally, augmentation was applied to the Normal and COVID-19 classes which had fewer numbers of data compared to the Pneumonia class. The dataset with augmented images was split to train, validation and test datasets using a split ratio of 60:20:20. Thus, a total of 13446 images were divided into 8067 training, 2689 validation and 2689 test images. This modified architecture was trained for 10 epochs, with the same hyperparameters used for training the original U-Net model.

4.3 Classification with ensemble model

For classification, we made used the models trained in our previous studies [4][5]. The modified U-Net model, trained as described in earlier sections, was used to obtain segmentation masks of the datasets that were used for classification. We applied morphological dilation and erosion to the obtained masks to improve their quality before extracting the lung area by replacing it with the original pixels to obtain the area of interest. Then we separately trained the four modified CNN architectures MobileNetV2*, Resnet50*, Xception*, and InceptionV3* in our previous work [5], using the newly created dataset. Next, we combined these models to form an ensemble where the four models MobileNetV2*, Resnet50*, Xception*, InceptionV3* were given the weights 0.5, 0.1, 0.1, 0.4, respectively. The weights to be given for each model were determined by experimenting with various weight combinations until the best results possible were obtained.

5 Results and analysis

The dice coefficient, intersection over union (IoU), precision and recall were used as the performance metrics to assess the performance of the U-Net model while accuracy, recall, precision, and F1 score were used to assess the performance of the classification models and the ensemble. It was observed that the modified U-Net model could obtain better results than the original U-Net model.

5.1 Results of U-Net based segmentation

Figure 4 shows the results obtained from the modified U-Net architecture with residual convolutional blocks and dropout layers. Here, we have used an augmented dataset addressing the class imbalance problem. The training was performed for 10 epochs. The modified U-Net architecture has achieved higher results for the dice coefficient, IoU score, precision, and recall. The dice coefficient and IoU score measure the similarity of the output mask to the ground truth mask.

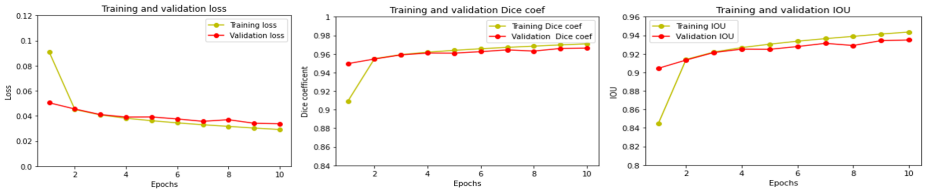

Fig. 4. Modified U-Net results: (a) loss, (b) dice coefficient, (c) IoU

Accordingly, 96.65% of a dice coefficient and 93.53% of an IoU score verify the model's ability to predict the segmentation masks accurately. The dice coefficient and IoU score measure the similarity of the output mask to the ground truth mask. 96.65% of a dice coefficient and 93.53% of an IoU score verify the model's ability to predict the segmentation masks accurately. Moreover, a precision value of 97.35% and a recall of 96.00% is achieved by the model. The recall is the proportion of the boundary pixels in the ground truth mask that were successfully detected by the automatic segmentation while precision is the proportion of boundary pixels in the predicted mask that corresponds to the boundary pixels in the ground truth mask. The precision and recall can measure the segmentation quality because they are sensitive to over and under-segmentation. Over-segmentation can lead to low precision scores, while under-segmentation can lead to low recall scores. By how the model has achieved higher values for both precision and recall, it can be verified that the boundaries are identified accurately and in turn has led to the accurate identification of regions themselves.

The training and validation curves for the dice coefficient, IoU score, and dice loss are shown in the graphs in Figure 4. With each training epoch, training, and validation curves for both dice coefficient and IoU have increased while the loss curves have decreased. It is visible in all the graphs that training and validation curves have a lesser gap during the last epochs which indicates that the overfitting is minimized in the model. This can be due to the measures taken for mitigating overfitting such as augmentation and addition of dropout layers. There was a class imbalance in the original dataset, and it was solved by oversampling the classes with a lesser number of images by applying several augmentation techniques like median blurring, horizontal flip, and Center Crop. In addition to that, the dataset was further generalized by the applied augmentation methods.

Dropout is a technique used to prevent a model from overfitting. Dropout works by randomly setting the outgoing edges of hidden units to 0 at each update of the training phase. We have added dropout layers with a 10% dropout rate since higher dropout rates tend to result in a higher variance to some of the layers resulting in the model not being able to fit properly. These dropout layers are added to the convolution blocks in both the encoder and the decoder of the U-Net model to minimize the overfitting.

5.2 Results of classification with U-Net based segmentation

We have obtained the performance metrics for the classification of different CNN architectures, where the input is the segmented images from the proposed U-Net model. Figure 5 shows the training and validation curves of accuracy and loss for the modified

models MobileNetV2*, Resnet50*, Xception*, and InceptionV3* that were trained using the lung extracted images. The dice coefficient and IoU curves for both training and validation have increased while the loss curves have decreased with each epoch. In addition to that, it is visible in all the graphs that training and validation curves have a lesser gap during the last epochs which indicates that the overfitting is minimized in the model.

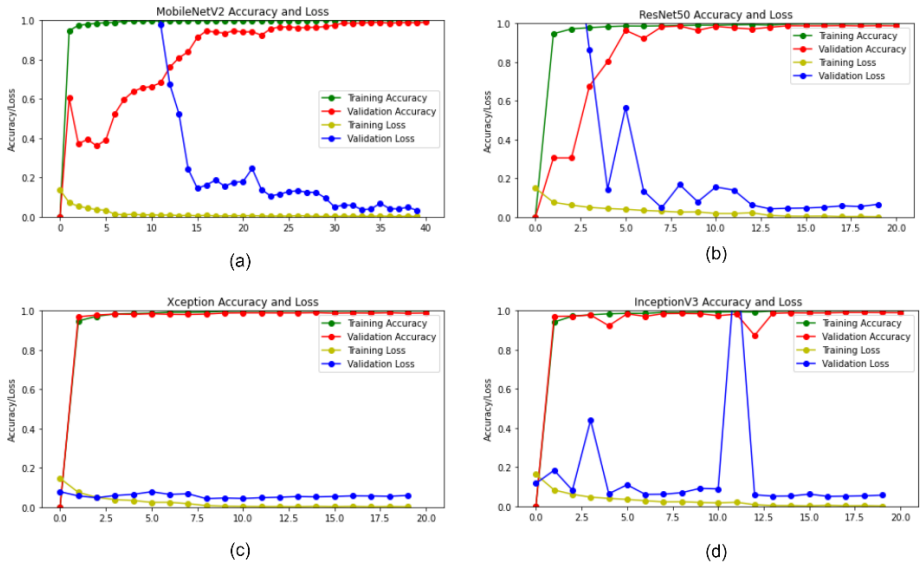

Fig. 5. Results of classification with modified U-Net (a) MobileNetV2*, (b) ResNet50*, (c) Xception *, (d) InceptionV3*

5.3 Analysis of results

Dice coefficient and IoU score are widely considered as the best evaluation matrices for segmentation tasks in imbalances datasets. As stated before, they check the similarity between the output mask and the ground truth mask. IoU is the area of overlap between the predicted mask and the ground truth mask divided by the area of union between the predicted mask and the ground truth mask while the dice coefficient is two times the area of overlap divided by the total number of pixels.

Table 2 presents a comparison between the results obtained for segmentation by the original U-Net and the modified U-Net proposed in this work. The original U-Net model, trained for 10 epochs without applying augmentation methods obtained a dice coefficient of 94.97% and an IoU score of 90.45% along with a 95.68% precision value and 95.05% recall value. As can be seen from the table, the overall results of the modified U-Net architecture with residual convolutional blocks and dropout layers trained for 10 epochs are higher than those of the original U-Net architecture. Compared to the original U-Net architecture, the IoU score, recall and dice coefficient has increased while the loss has reduced.

Table 2. Analysis of the results: Modified vs. original U-Net segmentation

Metric	Modified U-Net	Original U-Net
Dice coefficient	96.65%	94.97%
IoU score	93.53%	90.45%
Precision	97.35%	95.68%
Recall	96.00%	95.05%

In the residual convolutional blocks, residual skip connections are used to make it easy to learn an identity function. Since the dimensions of the input and output layers are different, a convolutional layer was added to the shortcut path. The ease of learning an identity function has caused an increase in performance. To avoid overfitting, a dropout layer was added to the convolutional block with a drop rate of 0.1.

Furthermore, Table 3 shows a comparison between the modified CNN architectures with direct classification [5], and classification results with the proposed U-Net segmentation. Also, it shows the results obtained from the ensemble. Here, Accuracy, precision, recall and F1-score are denoted as Acc, RC, PR and F1, respectively, while CLF denotes the classification results obtained in our previous work [5].

Table 3. Analysis of the results: different models

Metric	MobileNetV2*		Resnet50*		Xception*		InceptionV3*		U-Net + Ensemble
	CLF	*with U-Net*	*CLF*	*with U-Net*	*CLF*	*with U-Net*	*CLF l*	*with U-Net*	
Acc	98.65	99.24	98.87	98.86	98.61	98.65	98.61	98.86	99.36
RC	98.15	99.26	98.54	98.89	98.61	98.70	98.22	98.89	99.38
PR	98.20	99.27	98.37	98.91	97.96	98.69	98.02	98.90	99.40
F1	98.17	99.26	98.45	98.90	98.14	98.69	98.12	98.89	99.39

When comparing the results of the four models, it can be seen that all models have achieved better results after segmentation was applied to the dataset. When the dataset was used as it is, the models make use of the features in the background as well, which can result in lower accuracies. The segmented images will only contain the area of interest, which is the lung area. Hence, the models will only make use of the relevant features to make decisions, thus achieving higher accuracy. The ensemble has achieved even better results than the individual models. Different models focus on different features of the data and combining all of them can bring better results.

Figure 6 shows the confusion matrix obtained for the classification with the modified CNN architectures in our previous work [5], for the comparison purpose. Figure 7 depicts the confusion matrix obtained for the classification with the modified architectures trained using the lung area-extracted dataset.

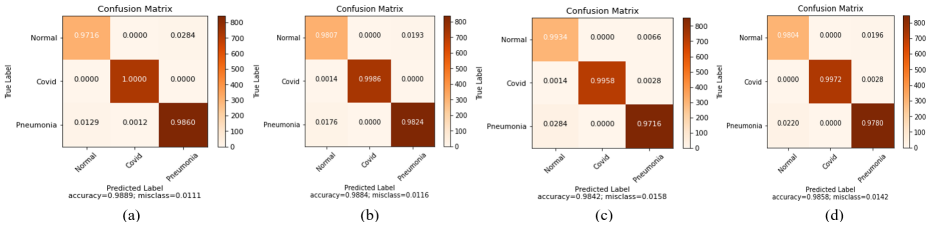

Fig. 6. Confusion matrix for direct classification (a) MobileNetV2*, (b) ResNet50*, (c) InceptionV3*, (d) Xception*

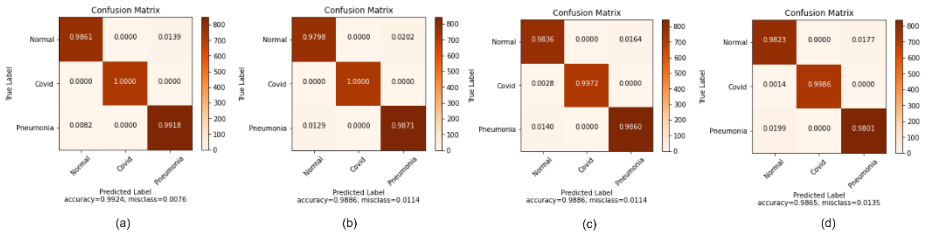

Fig. 7. Confusion matrix for classification with U-Net (a) MobileNetV2*, (b) ResNet50*, (c) InceptionV3*, (d) Xception*

When comparing the matrices of the different models in both cases, it can be seen that the percentages of true positives detected for each of the three classes vary among the models. Although all models have obtained somewhat similar results in each case, at the class level, their results have varied a lot. But overall, it can be said that results obtained during the case of using segmentation are considerably higher compared to the case without.

Figure 8 shows the confusion matrix obtained for the ensemble model with U-Net segmentation. The confusion matrix shown here provides an evaluation of the classification of chest X-ray images using an ensemble model combined with the U-Net based segmentation. Unlike in the previous cases, in this matrix, it can be seen that for all three classes which are, the normal class, COVID-19 class, and the Pneumonia class, the percentages for the detection of true positives have all exceeded 98%.

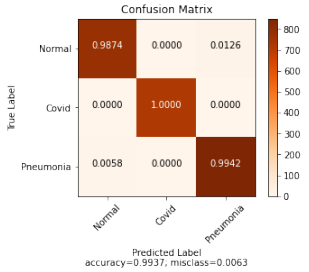

Fig. 8. Confusion matrix for the ensemble model with U-Net segmentation

6 Discussion

Generally, different machine learning and deep learning classification techniques are widely used in the medical domain [3][16][17][18]. On the other hand, U-Net is specifically used for segmentation tasks. Segmentation use cases are low compared to the classification and regression use cases, but segmentation can be of help in improving the performance of classification tasks [7]. In the medical domain, for chest X-ray images, CNN models like MobileNet, ResNet and Inception are used for classifying images into different classes like Pneumonia, COVID-19 and other varIoUs lung diseases while semantic segmentation tasks are performed using models like the U-Net model, a deep learning model highly used for image segmentation. Table 4 shows the comparison of the proposed approach with the existing U-Net based implementations for chest X-ray analysis.

Table 4. Comparison with the existing studies

Study	Dataset Name	Used DL models	Evaluation metrics
[11]	Private dataset JRST dataset Shenzhen Dataset	U-Net	Segmentation: IoU = 95% Accuracy = 98.3%
[12]	Montgomery County X-ray Set	U-Net	Segmentation: IoU - 91.37%. Dice Coefficient = 94.21%
[13]	Covid19 - cohen Normal and pneumonia - Kaggle chest X-ray dataset	Segmentation - U-Net model (DenseNet103 as backend). Classification - Ensemble (DenseNet121, EfficientNetB0, VGG19)	Segmentation: Accuracy = 92% IoU = 90% Dice coefficient = 92% Classification :- Accuracy = 99.2%
[10]	V7-labs COVID-19 X-ray dataset, Cohen, RSNA pneumonia detection challenge, RYDLS-20-v2 database	Semantic segmentation using U-Net. Classification using VGG, ResNet and Inception separately.	Segmentation: Jaccard distance = 3.4% Dice coefficient = 98.2% Classification :- Multi-class: F1 = 88% COVID-19: F1 = 83%
Proposed U-Net + Ensemble	Segmentation - V7-labs COVID-19 X-ray dataset. Classification - Kaggle pneumonia dataset, Kaggle covid-19 radiography.	Segmentation - Modified U-Net with residual convolution blocks and dropout layers. Classification - Ensemble model (mobilenetv2, Inceptionv3, Xception, ResNet50)	Segmentation: Dice coefficient = 96.65% IoU = 93.53% Precision = 97.35% Recall = 96.00% Classification :- Accuracy = 99.36% F1-score = 99.39% Precision = 99.40% Recall = 99.38%

Most of the related studies that have used the U-Net model for the chest X-ray analysis have performed segmentation and classification tasks separately using different models. Some of the studies have used U-Net architecture with different backbones like DenseNet103 for the segmentation task while for the classification task, widely used models like VGG, ResNet, and Inception have been used separately or as an ensemble

of the models. In many studies, higher results are achieved for either segmentation with higher values for metrics like dice coefficient and IoU score, or the classification task with higher values for metrics like classification accuracy. Accordingly, we were able to achieve higher results for the metrics for both segmentation and classification tasks.

From a medical practitioners' perspective, the main requirement would be to correctly identify the disease of the patient. Hence, automated tool support for the classification of chest X-rays is important. However, extracting the lung regions by applying segmentation to the chest X-rays and feeding those to the CNNs will allow the models to achieve better results for classification in terms of accuracy and performance. The CNN would perform better because it will only consider the area of interest and not the background when making the decisions. In addition, since the model is not considering any unwanted areas, performance would be improved. Apart from that, these segmented images can also be used to incorporate explainability to the models to help doctors understand the workings of the model better [19].

7 Conclusion

This paper presented a deep learning based approach to support the chest X-ray analysis in terms of Pneumonia and COVID-19 diseases. We have shown that, promising results can be obtained by incorporating both segmentation and classification techniques during the deep learning approach. We proposed a U-Net based architecture in which the convolutional blocks in the original architecture were replaced with residual blocks and dropout layers. The proposed model was able to perform segmentation on chest X-rays achieving higher accuracy than the original U-Net was capable of. Employing augmentation techniques to solve the class unbalance problems in the dataset further helped in obtaining these results. We extracted the lung areas of the chest X-rays using the said segmentation model and fed them to CNN models proposed in a previous study of ours and demonstrated how the segmentation was able to improve the performance of the classification. Finally, the separate models, together with the segmentation model, were combined to form an ensemble model which was able to achieve even better results.

The proposed model can be extended to develop as a tool that can be used in real-world clinical practice. To overcome the black-box nature of the model and make it more understandable to users, an explainability model can be integrated with the proposed model. A model such as LIME can be used to generate an image output showing which regions of the lungs were used by the model in making the decisions.

8 Acknowledgement

The publication of this paper was made possible by a sponsorship from CTI Global, Frankfurt, Germany.

9 References

[1] World Health Organization, Pneumonia, [Online; accessed 10-March-2022] (2021). https://www.who.int/news-room/fact-sheets/detail/pneumonia

[2] Kong, W., & Agarwal, P. P. (2020). Chest imaging appearance of COVID-19 infection. Radiology: Cardiothoracic Imaging, 2(1): e200028. https://doi.org/10.1148/ryct.2020200028

[3] Kieu, S. T. H., Bade, A., Hijazi, M. H. A., & Kolivand, H. (2020). A survey of deep learning for lung disease detection on medical images: state-of-the-art, taxonomy, issues and future directions. Journal of Imaging, 6(12), 131. https://doi.org/10.3390/jimaging6120131

[4] Kolonne, S., Fernando, C., Kumarasinghe, H., & Meedeniya, D. (2021). MobileNetV2 Based Chest X-Rays Classification. In 2021 International Conference on Decision Aid Sciences and Application (DASA). pp. 57-61. https://doi.org/10.1109/DASA53625.2021.9682248

[5] Fernando, C., Kolonne, S., Kumarasinghe, H., & Meedeniya, D. (2022). Chest Radiographs Classification Using Multi-model Deep Learning: A Comparative Study. In 2022 2nd International Conference on Advanced Research in Computing (ICARC). pp. 165-170 http://dx.doi.org/10.1109/ICARC54489.2022.9753811

[6] Mamdouh, R., El-Khamisy, N., Amer, K., Riad, A., & El-Bakry, H. M. (2021). A New Model for Image Segmentation Based on Deep Learning. International Journal of Online & Biomedical Engineering (iJOE). 17(7). pp. 28-47. https://doi.org/10.3991/ijoe.v17i07.21241

[7] Ayad, H., Ghindawi, I., & Kadhm, M. (2020). Lung Segmentation Using Proposed Deep Learning Architecture. International Journal of Online & Biomedical Engineering (iJOE). 16(15). pp. 141-147. https://doi.org/10.3991/ijoe.v16i15.17115

[8] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention. pp. 234-241. Springer, Cham.

[9] Rubasinghe, I., & Meedeniya, D. (2019). Ultrasound nerve segmentation using deep probabilistic programming. Journal of ICT Research and Applications, 13(3), 241-256. https://doi.org/10.5614/itbj.ict.res.appl.2019.13.3.5

[10] Teixeira, L. O., Pereira, R. M., Bertolini, D., Oliveira, L. S., Nanni, L., Cavalcanti, G. D., & Costa, Y. M. (2021). Impact of lung segmentation on the diagnosis and explanation of COVID-19 in chest X-ray images. Sensors, 21(21), 7116. https://doi.org/10.3390/s21217116

[11] Narayanan, B. N., & Hardie, R. C. (2019). A computationally efficient u-net architecture for lung segmentation in chest radiographs. In IEEE National Aerospace and Electronics Conference (NAECON). pp. 279-284. https://doi.org/10.1109/NAECON46414.2019.9058086

[12] Rahman, M. F., Tseng, T. L. B., Pokojovy, M., Qian, W., Totada, B., & Xu, H. (2021, February). An automatic approach to lung region segmentation in chest x-ray images using adapted U-Net architecture. In Medical Imaging 2021: Physics of Medical Imaging. 11595, pp. 115953I. International Society for Optics and Photonics. https://doi.org/10.1117/12.2581882

[13] Hasan, M. J., Alom, M. S., & Ali, M. S. (2021). Deep learning based detection and segmentation of COVID-19 & pneumonia on chest X-ray image. In International Conference on Information and Communication Technology for Sustainable Development (ICICT4SD). pp. 210-214. https://doi.org/10.1109/ICICT4SD50815.2021.9396878

[14] v7labs. (2020). Covid-19-xray-dataset. https://github.com/v7labs/COVID-19-xray-dataset

[15] Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. Journal of big data, 6(1), pp. 1-48. https://doi.org/10.1186/s40537-019-0197-0

[16] Li, Y., Zhang, Z., Dai, C., Dong, Q., & Badrigilan, S. (2020). Accuracy of Deep Learning for Automated Detection of Pneumonia Using Chest X-Ray Images: A Systematic Review and Meta-Analysis. Computers in Biology and Medicine, 123, 103898. https://doi.org/10.1016/j.compbiomed.2020.103898

[17] Wijethilake, N., Meedeniya, D., Chitraranjan, C., Perera, I., Islam, M., & Ren, H. (2021). Glioma Survival Analysis Empowered with Data Engineering - A Survey. IEEE Access, 9, 43168-43191. https://doi.org/10.1109/ACCESS.2021.3065965

[18] De Silva, S., Dayarathna, S., Ariyarathne, G., Meedeniya, D., & Jayarathna, S. (2019). A Survey of Attention Deficit Hyperactivity Disorder Identification using Psychophysiological Data. International Journal of Online and Biomedical Engineering (iJOE). 15 (13), pp. 61-76. https://doi.org/10.3991/ijoe.v15i13.10744

[19] Singh, A., Sengupta, S., & Lakshminarayanan, V. (2020). Explainable deep learning models in medical image analysis. Journal of Imaging, 6(6), 52. https://doi.org/10.3390/jimaging6060052

10 Authors

Ms. K. A. S. H. Kumarasinghe is a B.Sc. Eng. undergraduate in the Department of Computer Science and Engineering at the University of Moratuwa, Sri Lanka. Her main research interests are computer vision and deep learning (email: kashkumarasinghe.17@cse.mrt.ac.lk).

Ms. S. L. Kolonne is a B.Sc. Eng. undergraduate in the Department of Computer Science and Engineering at the University of Moratuwa, Sri Lanka. Her main research interests are computer vision and deep learning (email: shammi.17@cse.mrt.ac.lk).

Ms. K. C. M. Fernando is a B.Sc. Eng. undergraduate in the Department of Computer Science and Engineering at the University of Moratuwa, Sri Lanka. Her main research interests are computer vision and deep learning (email: kmadhushani.17@cse.mrt.ac.lk).

Prof. D. Meedeniya is a Professor in Computer Science and Engineering at the University of Moratuwa, Sri Lanka. She holds a PhD in Computer Science from the University of St Andrews, United Kingdom. She is the director of the Bio-Health Informatics group at her department and engages in much collaborative research. She is a co-author of 100+ publications in indexed journals, peer-reviewed conferences, and international book chapters. She serves as a reviewer, program committee and editorial team member in many international conferences and journals. Her main research interests are Software modelling and design, Bio-Health Informatics, Deep Learning and Technology-enhanced learning. She is a Fellow of HEA (UK), MIET, MIEEE, Member of ACM and a Chartered Engineer registered at EC (UK) (email: dulanim@cse.mrt.ac.lk).

Article submitted 2022-03-12. Resubmitted 2022-04-23. Final acceptance 2022-04-23. Final version published as submitted by the authors.