

Leveraging Deep Learning for Critical X-ray Classification in the Era of Respiratory Diseases

By Ahmet SOLAK

Leveraging Deep Learning for Critical X-ray Classification in the Era of Respiratory Diseases

Ahmet SOLAK^{1,*} 

¹ Konya Technical University, Faculty of Engineering and Natural Sciences, Department of Electrical and Electronics Engineering, Konya, Türkiye.

Abstract

This research addresses the challenge of categorising X-ray images into three categories: COVID-19, normal, and viral pneumonia. Effective classification in respiratory diseases is crucial for timely diagnosis and treatment planning, with direct implications for patient health. Deep learning, particularly Convolutional Neural Networks (CNNs), is becoming an essential tool in medical image classification due to its unique capabilities in pattern recognition and feature extraction. This paper emphasises the significance of using deep learning techniques to automate and enhance the accuracy of X-ray image classification, thereby assisting healthcare professionals in making informed decisions. Transfer learning is an important aspect of this work, as it can utilise pre-existing knowledge from large datasets. This approach is particularly useful in the medical field, where labelled data is often limited and difficult to obtain. The discussion emphasises the advantages of transfer learning in optimising model performance and reducing the need for extensive labelled data. The EfficientNetB0 model was chosen due to its superior efficiency in balancing model complexity and computational resources. The rationale behind this choice is explained by highlighting the model's robustness and generalisation capabilities in medical image classification tasks. The paper concludes by presenting results that prove the effectiveness of the proposed methodology and underline its contribution to the field. The research presented here is unique in its comprehensive approach, which combines X-ray classification, deep learning, and transfer learning to create an efficient model. The evidence provided in this research demonstrates the model's strong performance, contributing to the advancement of medical image analysis and holding promise for real-world applications in the rapid and accurate diagnosis of respiratory diseases.

Keywords: X-Ray Classification, Deep Learning, Transfer Learning, EfficientNetB0, Respiratory Diseases

1. Introduction

Respiratory diseases are a significant global health concern, placing a burden on individuals and healthcare systems worldwide. These conditions range from mild respiratory infections to severe chronic diseases and are influenced by environmental, genetic, and lifestyle factors [4]. In recent years, the world has experienced an increased focus on respiratory health due to the emergence of the novel coronavirus, SARS-CoV-2, and the disease it causes, COVID-19 [2]. The ongoing global pandemic has highlighted the vulnerability of the respiratory system and the significant impact a single respiratory virus can have on populations worldwide.

Advancements in medical imaging and artificial intelligence (AI) have revolutionised the diagnosis of respiratory diseases, particularly through the analysis of X-ray images. AI-based classification systems have emerged as powerful tools in assisting healthcare professionals to identify respiratory conditions early and accurately. These systems leverage deep learning algorithms to discern subtle patterns and anomalies in X-ray images that may elude the human eye. In the field of respiratory diseases, artificial intelligence (AI) models are being trained to accurately distinguish between various conditions, including pneumonia, tuberculosis, and COVID-19 [3-5]. This is achieved through the integration of convolutional neural networks (CNNs) and other machine learning techniques, resulting in the development of robust models capable of rapid and precise disease classification. AI-driven approaches can improve diagnostic speed and accuracy, and potentially alleviate the burden on healthcare systems, especially in regions with limited access to specialized medical expertise. However, ongoing research is necessary to refine and responsibly implement AI-based classification systems for respiratory diseases on X-ray images, due to challenges such as dataset biases, interpretability, and ethical considerations [6].

In recent years, the field of medical image analysis has undergone a paradigm shift with the introduction of deep learning architectures. These architectures have demonstrated remarkable capabilities in accurately classifying complex medical conditions [7, 8]. Among these architectures, the EfficientNetB0 [9] model has

* Corresponding author. e-mail address: asolak@ktun.edu.tr.

1

emerged as a powerful tool for image classification tasks. The EfficientNetB0 model leverages the principles of neural network efficiency and optimization to strike an optimal balance between model size and performance. This makes it particularly well-suited for resource-constrained environments. In the realm of respiratory disease classification, where precision and efficiency are paramount, the deployment of EfficientNetB0 has shown great promise. The study aims to classify X-ray images into distinct categories, including viral pneumonia, COVID-19, and normal cases, using the EfficientNetB0 model. Traditional classification methods face challenges due to the intricate patterns and subtle nuances present in X-ray images, highlighting the need for advanced deep learning models. Our initial findings suggest that the EfficientNetB0 model is highly accurate in distinguishing between respiratory diseases and is also able to generalize well across diverse datasets. This research adds to the existing literature on the potential of deep learning models, specifically EfficientNetB0, to transform the diagnosis and classification of respiratory diseases. This offers a promising opportunity to improve clinical decision support systems in modern healthcare.


2. Materials and Methods

2.1. Dataset

Accurate and timely diagnosis of respiratory disease is critical for effective patient care, especially in the context of the ongoing COVID-19 pandemic. To contribute to the advancement of medical imaging and diagnostic research, using a comprehensive dataset of X-ray images representative of three different categories: COVID-19, viral pneumonia and normal cases. The dataset, carefully curated from multiple sources, includes high-resolution digital radiographs that capture the complex pulmonary manifestations associated with these respiratory diseases. A total of 317 images are included in the dataset, with a balanced distribution across the three classes, ensuring robust training and evaluation of our classification model [10].

The COVID-19 subset of the dataset consists of radiographs of individuals with confirmed COVID-19 diagnoses, demonstrating the diverse manifestations of this viral infection. In addition, images of patients with viral pneumonia highlight the similarities and differences in radiological patterns between viral pneumonia and COVID-19. The normal class includes radiographs of individuals without respiratory abnormalities, providing a baseline for distinguishing pathological features.

The use of this carefully curated dataset forms the cornerstone of our investigation into the effectiveness of the EfficientNetB0 model for classifying COVID-19, viral pneumonia and normal cases. Our aim is to provide evidence to help healthcare professionals make more informed diagnostic decisions, ultimately improving patient outcomes in the management of respiratory disease. Figure 1 shows examples of the images available in the dataset.



5

Figure 1. Sample images of each class in the dataset.

2.2. Methods

The architecture known as EfficientNetB0, proposed by Tan et al. in 2019 [9], represents a milestone in the field of convolutional neural networks (CNNs), introducing a novel approach to balancing model size and performance. EfficientNetB0 is based on a compound scaling method that systematically optimises the depth, width and resolution of the model, resulting in a highly efficient and powerful neural network. The core idea behind EfficientNetB0 is to scale model dimensions uniformly to improve representation learning while maintaining computational efficiency. The architecture is characterised by a progressive increase in depth and width with a corresponding decrease in resolution, effectively addressing the trade-off between model complexity and computational cost.

At its core, EfficientNetB0 consists of a stack of repeated blocks, each containing depth-separable convolutions and inverted residuals. The depth-separable convolutions reduce computational complexity by decoupling spatial and channel-wise convolutions. The inverted residuals, which include skip connections, facilitate the flow of information through the network, allowing for more effective gradient propagation during training. This unique combination of architectural elements allows EfficientNetB0 to achieve state-of-the-art performance in various image classification tasks while maintaining a remarkably compact model size. This study utilized two distinct methodologies for the classification of brain tumor images: a custom-designed Convolutional Neural Network (CNN) architecture and EfficientNetB0, a widely adopted model within the domain of transfer learning.

The benefits of EfficientNetB0 go beyond its architectural elegance. In particular, the model demonstrates superior parameter efficiency, achieving competitive accuracy with significantly fewer parameters compared to other contemporary CNN architectures. This efficiency makes EfficientNetB0 particularly suitable for resource-constrained environments, a critical consideration in medical image analysis where computational resources may be limited. In the context of our respiratory disease classification study, the use of EfficientNetB0 serves as a strategic choice to optimise both computational efficiency and classification accuracy, laying the foundation for a robust and scalable diagnostic tool.

Figure 2 presents a summary of the model obtained by adapting the EfficientNetB0 to our classification problem. Then, optimization difficulties are addressed. Firstly, the trained blocks of the efficientnet model are taken for the problem by transfer learning in the functional block. Batch normalization, dense layers for classification, and dropout layers to prevent overfitting were used to promote consistent and efficient training and improve the overall performance and generalization ability of the models. Finally, the network was completed with a customized dense layer based on the number of classes.

Layer (type)	Output Shape	Param #
efficientnetb0 (Functional)	(None, 1280)	4049571
batch_normalization (Batch Normalization)	(None, 1280)	5120
dense (Dense)	(None, 256)	327936
dropout (Dropout)	(None, 256)	0
dense_1 (Dense)	(None, 3)	771
Total params: 4,383,398		
Trainable params: 4,338,815		
Non-trainable params: 44,583		

Figure 2. Summary of the model obtained by adapting the EfficientNetB0 model to the problem.

3. Results and Discussion

In this section, we present the outcomes derived from the training regimen applied to the EfficientNetB0 model. The parameters established prior to commencement of training were meticulously defined, encompassing an input size of (224,224), a batch size of 20, 100 epochs, a learning rate set at 0.001, Adamax optimizer implementation, and categorical crossentropy serving as the loss function. During the training phase, images were initially resized to dimensions of 224x224 before being inputted into the network, facilitating the subsequent training processes.

Subsequent to the training phase, the reliability of the model was assessed through rigorous testing with a set of distinct test images. The precision, recall, and F1 scores for each classification category were meticulously computed and are delineated in Table 1. Notably, the attained scores consistently surpass 0.9, indicative of the model's commendable performance across diverse metrics.

Table 1. Test results for EfficientNetB0

Class	Precision	Recall	F-1 Score
COVID	1	1	1
Normal	1	0.89	0.94
Viral Pneumonia	0.91	1	0.95

Furthermore, a comprehensive confusion matrix, depicted in Figure 3, elucidates the numerical breakdown of correctly and misclassified test data. This matrix not only provides insights into the overall accuracy of the model but also delineates misclassifications, specifying the predicted class for instances where misclassification occurred. The findings collectively underscore the robustness and efficacy of the EfficientNetB0 model in achieving high predictive accuracy and reliability.

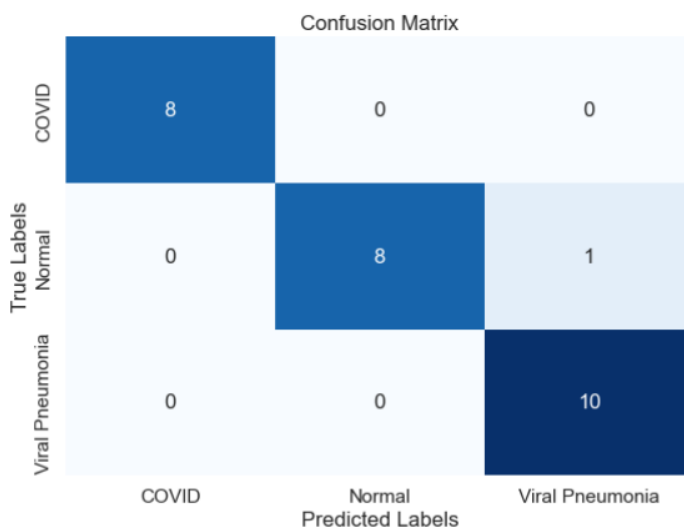


Figure 3. Confusion Matrix for EfficientNetB0

The EfficientNetB0 model's superior performance is evident from Table 1 and Figure 3. It consistently outperformed itself across all classes. The confusion matrix analysis further accentuates this superiority, revealing only one misclassification in the test dataset. Specifically, an image intended to be classified as normal was erroneously identified as viral pneumonia. This misclassification highlights the effectiveness of the model used. Therefore, the results of this study confirm the superior predictive capabilities of the model.

4. Conclusion

In conclusion, this study examines the use of the EfficientNetB0 model for classifying respiratory diseases based on X-ray images, with a focus on distinguishing COVID-19, viral pneumonia, and normal cases. The study's analysis, supported by both quantitative metrics and visual representations, demonstrates the remarkable effectiveness of the EfficientNetB0 model in achieving superior performance across all targeted classes.

Table 1 and Figure 3 demonstrate that EfficientNetB0 consistently outperforms other models, indicating its robustness and generalization capabilities. The confusion matrix analysis revealed only one misclassification in the test dataset, confirming the model's reliability.

Although no model is completely immune to misclassifications, EfficientNetB0 achieved high accuracy and precision in identifying cases of viral pneumonia, with only one instance of misclassification where a normal case was identified as viral pneumonia. This suggests that the model has great potential as a diagnostic tool for respiratory diseases, including the crucial task of distinguishing between COVID-19 and other pulmonary conditions.

The study's findings add to the increasing research on deep learning applications in medical imaging. EfficientNetB0's success in our specific context suggests its potential usefulness in real-world clinical settings, especially considering its efficiency gains in terms of model size and computational resources.

Looking ahead, future investigations could explore the model's performance on larger and more diverse datasets. Additionally, its interpretability in clinical decision-making and potential enhancements to address specific challenges in respiratory disease diagnosis could be examined. Based on the presented outcomes, it is clear that the EfficientNetB0 model is a valuable asset in improving the accuracy and efficiency of respiratory disease classification. This paves the way for enhanced diagnostic capabilities in the field of medical imaging and healthcare.

Leveraging Deep Learning for Critical X-ray Classification in the Era of Respiratory Diseases

ORIGINALITY REPORT

5%

SIMILARITY INDEX

PRIMARY SOURCES

1	www.ncbi.nlm.nih.gov Internet	36 words — 2%
2	Sohaib Asif, Wenhui Yi, Jin Hou, Tao Yi, Jinhai Si. "14 Deep convolutional neural network for the classification of COVID-19 from chest X-ray images", Walter de Gruyter GmbH, 2021 Crossref	27 words — 1%
3	pajes.pau.edu.tr Internet	18 words — 1%
4	timesofindia.indiatimes.com Internet	15 words — 1%
5	Mehmet Ergün Hatir, Mücahit Barstuğan, İsmail İnce. "Deep learning-based weathering type recognition in historical stone monuments", Journal of Cultural Heritage, 2020 Crossref	8 words — < 1%

EXCLUDE QUOTES OFF
EXCLUDE BIBLIOGRAPHY OFF

EXCLUDE SOURCES OFF
EXCLUDE MATCHES OFF