Efficient Deep Learning Approach for the Classification of Pneumonia in Infants from Chest X-Ray Images

Ahmed Saeed AlGhamdi[iD]

Department of Computer Engineering, College of Computer and Information Technology, Taif University, PO Box. 11099, Taif 21994, Saudi Arabia

Corresponding Author Email: asjannah@tu.edu.sa

ABSTRACT

Pediatric pneumonia is a major infectious disease which caused more than 5,00,000 deaths of infants and young children below the age of 5 years. This number is equivalent to the death of one child per minute. These statistics are daunting and more focused studies and development of tools are required to tackle this challenge. It has also been seen that pediatric pneumonia is curable through antibiotics and oxygen therapy if it is diagnosed at early stages. Current radiology techniques used are not able to diagnose pneumonia at early stages because specific arrangements are to be made for childcare. They cannot be treated as normal adults. Special care is to be taken for childcare as improper handling of radiology techniques may harm the child or generate inaccurate diagnosis. Recent advancements in computer aided diagnosis with the help of deep learning techniques has improved the quality of medical imaging techniques such as CT-Scan, X-ray images, etc. However, limited attention is given to the application of deep learning techniques for diagnosis and classification of pediatric pneumonia. Moreover, conducting manual tests, image classification and analysis of radiological images are prone to human errors due to lack of expertise of the radiologists. Inaccurate analysis of radiological images like X-rays can suggest inappropriate treatment for children which can prove to be fatal. To fasten the interpretation of radiological images, this paper suggests a deep learning model, focusing specifically on classifying pediatric pneumonia from chest X-ray images. In this paper, a MobileNeT-V3 architecture is implemented and tested against various datasets containing more than 10,000 chest X-ray images. Analysis of suggested implementation shows that this is better than various classical tools and techniques used for classification of chest X-ray. The suggested technique achieved a classification accuracy of 95.8% over dataset-1 and 97.8% over dataset-2, shows the efficiency of the technique. The model in this study not only demonstrated high classification accuracy but also excelled in other key metrics, achieving a precision of 97% and 94%, recall of 97% and 98%, and F1 scores of 97% for both datasets, underscoring its precision and reliability in diagnosing pediatric pneumonia.

1. INTRODUCTION

Pneumonia is a widespread illness that majorly affects young children, with the World Health Organization (WHO) reporting it as the single largest infectious cause of death in children worldwide [1-3]. Statistically, it is more prevalent among children under five, with recent studies indicating that up to 156 million new cases are recorded annually, out of which around 20 million cases are severe enough to require hospital admission [4, 5]. Specifically, the incidence of pneumonia in infants younger than two years old is about 13%, which is a significant portion of the pediatric population [4-7]. The disease is marked by respiratory symptoms such as a persistent cough, labored breathing, and fever, which can often be mistaken for less severe viral infections [7]. The severity of these symptoms can escalate rapidly, necessitating urgent medical attention. The WHO also notes that pneumonia accounted for approximately 15% of the 5.2 million deaths of

children under five in 2019, which translates to almost 800,000 deaths, primarily in low- and middle-income countries [7, 8]. Such high morbidity and mortality rates emphasize the need for improved diagnostic tools that can facilitate early detection and treatment. Rapid and accurate diagnosis is critical because pneumonia in children can be effectively treated with antibiotics, and in severe cases, with oxygen therapy, which are both highly successful if the illness is caught in time [9]. The challenge lies in the fact that pneumonia often shares symptoms with other illnesses, making it difficult to diagnose without the proper tools and expertise [10]. The Global Action Plan for Pneumonia and Diarrhea (GAPPD) by WHO and UNICEF aims to protect against, prevent, and treat pneumonia among children, with targets to reduce mortality rates [11]. However, to achieve these targets, there is a pressing need to bridge the gap in pneumonia diagnosis, particularly in regions where healthcare systems are burdened or under-resourced [12]. In many regions, especially in poorer parts of the world,

there is a significant lack of skilled diagnosis professionals. This shortage can lead to incorrect readings of X-ray images, which is a serious issue because it might mean that a child with pneumonia does not get diagnosed and treated properly and quickly [12]. Even in places with enough radiologists, the process of examining chest X-rays is still slow and complex. This is due to various factors such as the quality of the X-ray, the patient's positioning, and the subtlety of the signs of pneumonia on the images [13]. The delay in diagnosis due to these factors can be critical because quick and accurate identification of pneumonia is crucial for starting the right treatment in time. Moreover, Pneumonia is a significant health problem not just in one location but across various regions, each with its own set of challenges and impact levels. In the Asia-Pacific region, for example, community-acquired pneumonia (CAP) is a leading cause of death due to factors like an aging population, dense cities, and limited healthcare access [14]. The presence of pathogens like Klebsiella pneumoniae and Burkholderia pseudomallei, along with a high resistance to antibiotics, makes this region particularly vulnerable [14].

In a more global context, pneumonia poses a large burden on the elderly population, with hospital admission rates and case-fatality ratios (CFR) increasing with age [14, 15]. The systematic review in the paper [15] estimated a staggering 6.8 million hospital admissions for clinical pneumonia in older adults worldwide in 2015, with a higher rate among men and an increase in admissions with advancing age. This study underscores the heavy toll of pneumonia on healthcare systems, especially when considering the underestimated disease burden when using radiologically confirmed pneumonia as the diagnostic criterion. Focusing on Saudi Arabia, severe pneumonia in children under five results in millions of episodes and a significant number of deaths, with states like Uttar Pradesh and Bihar bearing the brunt of the disease [15]. Despite the vast number of cases, the actual morbidity and mortality rates at the sub-national level remain unclear. The disparity in the distribution of pneumonia cases and deaths across different states within Saudi Arabia highlights the need for targeted healthcare interventions and the importance of vaccines in reducing the disease burden. The varying prevalence and impact of pneumonia in these regions illustrate the necessity for region-specific strategies in managing and preventing pneumonia. They also highlight the importance of considering local epidemiological factors when designing public health initiatives and the potential benefits of vaccines in mitigating the disease's impact. Given these challenges, there's a pressing need for more reliable and faster diagnostic methods. This has led to research on deep learning systems that can analyze X-ray images and assist or even automate the diagnosis process. However, some challenges persist. For example, the authors [16] reports on a system utilizing VGG19, which achieved a classification accuracy of 86.97%, and upon enhancement with an Ensemble Feature Scheme, reached an accuracy of 95.70%. However, this system's reliance on a complex blend of handcrafted and deep-learned features can become cumbersome and less practical for everyday clinical use. Jaiswal et al. [17] introduces a Mask-RCNN based model with an innovative post-processing step, aiming for robust identification and localization of pneumonia. Despite its strengths, the model's need for significant training process adjustments and intricate post-processing can be a barrier in time-sensitive clinical environments. In the paper of Hashmi et al. [18], a model combining various advanced deep

learning architectures boasts a high accuracy of 98.43% and an impressive AUC score of 99.76%. Yet, this amalgamation of different models entails high computational costs and increased operational complexity. The approach explained in the paper of Chakraborty et al. [19] achieved a promising accuracy of 95.62% using a convolutional neural network. Nonetheless, it did not emphasize the solution's efficiency, which is critical for rapid diagnosis in real-world settings. In contrast, the MobileNet-V3 architecture proposed in this study is engineered to address these challenges effectively. It streamlines the diagnostic process by obviating the need for complex feature engineering and multiple-model ensembles. Its design leverages efficient computing blocks and lightweight layers, enhancing speed without sacrificing accuracy. This results in a fast, reliable, and resource-efficient model, suitable for real-time clinical application, capable of high accuracy that is in line with or surpasses the reported results, thus overcoming the drawbacks of the previous techniques. The proposed model is designed to classify chest X-rays into pneumonia and normal categories more efficiently than current methods. The MobileNeT-V3 architecture was selected for its efficient use of mobile computing resources, and its implementation of the H-Swish activation function is expected to enhance performance, especially in processing speed and accuracy. By automating the classification process, my aim to reduce the reliance on scarce radiological expertise and speed up the diagnosis, allowing for faster and potentially more accurate treatment interventions.

The contributions of this paper are as:

• This paper presents an X-ray classification technique based on MobileNeT-V3 architecture for classifying X-rays into pneumonia and normal type. This architecture is chosen because of the advantages like, efficient mobile building blocks, layer removal and H-Swish for improving the network, which will be discussed in upcoming sections in detail.

• Manual interpretation of X-rays is a tedious task and are prone to human errors. The classification technique suggested in this paper will resolve this challenge by automating the classification process and saving treatment time.

Rest of the paper is structured as: Literature review is covered in Section 2. Significance of using the suggested technique is explained in Section 3. Section 4 explains the techniques used and architecture in detail followed by results obtained and discussion under Section 5. The paper is concluded in Section 6.

2. RELATED WORK

X-ray is the most common medical imaging technique that is employed for analysis of internal organs such as lungs, bones, etc. to detect any abnormalities in the human body. However, X-rays are conducted manually by radiologists and are prone to errors during interpretation of X-rays. This could be fatal for the patients. To resolve this challenge, various studies have been done on chest X-rays. An innovative method was introduced utilizing the well-known VGG16 model [20], enhanced by fine-tuning its deep layers. This approach was not confined to merely identifying the presence of pneumonia but also extended to gauge its severity—a crucial aspect for clinical decision-making. The method achieved a test accuracy of 86.67%, with exceptionally high recall rates indicating its effectiveness in identifying true positive cases. However, the precision rate of 83% suggests that there was still a

considerable rate of false positives, which could lead to over-diagnosis. Additionally, the requirement for data augmentation due to imbalanced datasets introduces a limitation as it may not fully represent the natural variety in real-world clinical data. The authors [21] employed a CNN architecture, leveraging the strengths of VGG16 and Inception models. The use of transfer learning and ensembling aimed to capitalize on the combined strengths of these models to enhance prediction accuracy. Although the exact accuracy metrics are not detailed, the inherent complexity of ensembling models poses a challenge, as it requires significant computational resources and may complicate the model's interpretability in clinical settings. The authors [22] have developed a VGG-based model with a reduced number of layers, aiming to simplify the architecture. To address the issue of low contrast in chest X-ray images, which can obscure diagnostic details, the Dynamic Histogram Enhancement technique was employed as a preprocessing step. The model showed superior performance across several metrics, including an accuracy of 96.068%. Despite the reduction in complexity, the model still had a 4% increase in parameters compared to MobileNet, indicating that while it has made strides in efficiency, there is still room for improvement in creating a model that is both lightweight and highly accurate.

The authors [23] have examined the potential of machine learning over deep learning to automate the early detection of pediatric pneumonia, a crucial step to reduce its high morbidity and mortality. A Quadratic SVM model was used, resulting in an accuracy of 97.58%. While this study demonstrated machine learning's potential in medical imaging, its limitations lie in the necessity for extensive data augmentation to balance the dataset and optimize feature extraction, which can be resource-intensive and may not scale well in different clinical environments. The authors [24] delved into the utility of several advanced image recognition models to improve pneumonia detection from chest X-rays. This investigation deployed models like VGG16, ResNet, InceptionNet, and DenseNet, as well as a tailored CNN model. The VGG16 model stood out, achieving the lowest Mean Absolute Error (MAE), a statistical measure indicating the average magnitude of errors in predictions, clocking in at 66.19. To expedite the learning process of their model, the researchers employed TensorFlow's TPU strategy, a sophisticated technology designed to accelerate deep learning tasks. TPUs, or Tensor Processing Units, are Google's custom-developed application-specific integrated circuits (ASICs) used to boost performance in training large-scale neural network models. The use of TPUs enabled the researchers to slash training times by over 68% compared to traditional CPUs and nearly 55% in comparison to GPUs. However, the study's reliance on this advanced technology could pose accessibility issues in settings lacking such specialized hardware. In the paper of Malik et al. [25], the focus was on distinguishing COVID-19 from other respiratory diseases, an important task given their symptomatic similarities. A deep learning model named CDC Net was developed, which utilized concepts from residual networks and dilated convolution to enhance image analysis. When evaluated against public benchmark data, the CDC Net model achieved impressive metrics, including an AUC of 0.9953, demonstrating high accuracy in multi-disease classification from chest X-rays. Despite these strong performance indicators, the real-world application may be hampered by the variability of clinical imaging conditions not represented in public datasets.

The authors [26] introduce a technique combining transfer learning with adversarial training to refine pneumonia detection from X-ray images. By applying adversarial training, the authors aimed to improve the model's performance by including synthetic X-rays along with real ones during the training phase, thus enhancing the model's adaptability to new, unseen images. While this approach yielded a higher accuracy rate, its effectiveness depends on the continuous generation of high-quality synthetic images, which can be computationally demanding. The researchers [27] proposed a novel network called QCSA (Quaternion Channel-Spatial Attention Network) for pneumonia detection. This network blends spatial and channel attention mechanisms with Quaternion algebra to process chest X-ray images. The study reported an accuracy of 94.53% and an AUC of 0.89. Despite this, the integration of attention mechanisms, while beneficial for performance, introduces additional complexity to the model, which may challenge real-world clinical implementation without sufficient computational resources. The authors [28] sought to address the issue of dataset imbalance in chest X-ray pathology classification by constructing a weakly-labeled database from publicly available medical articles. Using advanced text extraction and image verification techniques, the researchers enhanced the detection of various thoracic diseases. Although their approach showed promising results, the dependency on text-based labels for image identification raises concerns about the accuracy of such weak labels and their impact on the model's diagnostic reliability.

The authors [29] tackled the challenge of developing a lightweight model for pneumonia detection that could be deployed in under-resourced regions. The proposed model, which used a combination of CNN architectures with varying kernel sizes, achieved a high recall value of 99.23% and an F1-score of 88.56%. The absence of deep neural networks and transfer learning in this model makes it less computationally intensive. However, the model's reliance on a novel weighted ensemble approach could introduce variability in diagnostic outcomes based on the adjustable threshold, possibly requiring fine-tuning to align with clinical expectations. The authors [30] addressed the need for accurate classification of COVID-19 from chest X-rays by utilizing an ensemble of pre-trained CNNs and textural features. The model was trained on a substantial dataset, resulting in a binary classification accuracy of 98.34% for COVID-19. While the study demonstrates the utility of large datasets in improving model performance, it also hints at potential overfitting when models are trained on smaller datasets, questioning the generalizability of the findings.

The authors [31] have implemented a deep learning model employing the VGG16 architecture in combination with Neural Networks (NN) to analyze chest X-ray (CXR) images. This approach is particularly crucial in developing countries where rapid and accurate diagnosis is hindered by poor living conditions and inadequate healthcare infrastructure. The technique used here harnesses VGG16's robust feature extraction capabilities, supplemented by the pattern recognition power of neural networks. The study reports substantial success with this method, as evidenced by an accuracy of 92.15%, a recall of 0.9308, a precision of 0.9428, and an F1-score of 0.937 for the first dataset. These metrics are indicative of the model's ability to correctly identify cases of pneumonia with a high degree of reliability. When applied to a more diverse second dataset, which included images indicative of pneumonia, normal conditions, and COVID-19,

the model maintained a high level of accuracy at 95.4%, with equally high recall, precision, and F1-score values. A comparative analysis within the study indicates that the VGG16-NN combination outstrips the performance of VGG16 when paired with other machine learning classifiers such as SVM, KNN, RF, and NB. This suggests that neural networks provide a more harmonious complement to VGG16 for this application. Despite the promising results, challenges remain implicit in the study. While not explicitly stated, such challenges likely include ensuring the model's robustness across varying quality of CXR images, adapting to the diversity of pneumonia manifestations across different demographics, and managing the computational load typical of deep learning models. These potential hurdles highlight the need for models that not only perform well under test conditions but also maintain their efficacy in the less-controlled environments of real-world clinical settings. Collectively, these studies underscore the significant progress made in using deep learning for pneumonia detection in chest X-rays, while also highlighting the need for solutions that balance accuracy, computational efficiency, and generalizability to diverse clinical environments.

A summary of the comparative analysis between the related works is done in Table 1.

Table 1. Overview of deep learning techniques for pneumonia detection in chest X-rays

Ref.	Method / Techniques	Datasets	Performance Metrics
[20]	VGG16 model	Imbalanced datasets	Accuracy: 86.67%, Recall: High, Precision: 83%
[21]	CNN, VGG16 and Inception models	Not detailed	Not specifically detailed, only stated high accuracy
[22]	Reduced VGG-based model	Low contrast chest X-rays	Accuracy: 96.068%, Parameters increased by 4%
[23]	Quadratic SVM	Extensive data augmentation	Accuracy: 97.58%
[24]	VGG16, ResNet, InceptionNet, DenseNet, tailored CNN model	Not detailed	MAE: 66.19, Training time reduced (TPUs used)
[25]	CDC Net (residual networks, dilated convolution)	Public benchmark data	AUC: 0.9953
[26]	Transfer learning with adversarial training	Synthetic + real X-rays	Higher accuracy, computationally demanding
[27]	QCSA (Quaternion Channel-Spatial Attention Network)	Not detailed	Accuracy: 94.53%, AUC: 0.89
[28]	Weakly-labeled database construction	Public medical articles	Not detailed
[29]	Combination of CNN architectures	Not detailed	Recall: 99.23%, F1-score: 88.56%
[30]	Ensemble of pre-trained CNNs and textural features	Large dataset	Accuracy: 98.34%
[31]	VGG16 architecture with Neural Networks	Two datasets	Accuracy: 92.15% - 95.4%, Recall: 0.9308, Precision: 0.9428, F1-score: 0.937

The existing studies in the literature review offer a breadth of approaches to pneumonia classification, each with distinct advantages and shortcomings. Many of these methods, such as VGG16 and its variations [20-22, 31] although accurate, suffer from high computational costs and complexity, which can be prohibitive in clinical settings, especially in developing countries. Others, like the QCSA network [27], introduce additional complexity through attention mechanisms, while methods involving adversarial training or large ensembles of CNNs may not be feasible due to the need for substantial computational resources and potential overfitting issues. MobileNetV3, the model selected for this study, is engineered to address these gaps by providing a balance between accuracy and efficiency. It is specifically designed for mobile and edge devices, which means it requires significantly less computational power without a substantial trade-off in performance [32]. By leveraging lightweight depth wise convolutions and an architecture optimized for mobile devices, MobileNetV3 offers a more accessible and practical solution for pneumonia classification in chest X-ray images of infants. This makes it particularly suitable for real-world clinical applications where resources are limited and quick, reliable diagnostics are needed. Thus, MobileNetV3 has the potential to make the detection of pneumonia in infants more feasible in diverse clinical environments.

3. SIGNIFICANCE AND MOTIVATION

Detecting pneumonia in children using chest X-rays presents unique challenges. The developing anatomy of a child's lungs often conceals the radiographic indicators of pneumonia that are more apparent in adults. Moreover, the discomfort and symptoms associated with pneumonia, such as persistent coughing, can cause children to move during the X-ray procedure, resulting in blurred images that make accurate diagnosis difficult [33]. Furthermore, the process of obtaining a high-quality chest X-ray from a child is fraught with difficulties [34]. The fidgeting and lack of cooperation often seen due to their discomfort, or the intimidating nature of the procedure can lead to compromised image quality. This is problematic because the subtleties of pneumonia in pediatric X-rays require clarity for proper identification [35, 36]. This model, MobileNet-V3 stands out as a solution due to its design which prioritizes speed and efficiency without sacrificing accuracy. Its architecture is fine-tuned to perform well even with images that aren't perfectly clear, which is common in pediatric radiographs. This means it can still accurately detect the presence of pneumonia from images that are less than ideal, reducing the need for repeat X-rays and minimizing stress for young patients. Furthermore, MobileNet-V3's streamlined design allows for rapid processing of images. This quick turnaround is crucial in clinical settings where time is of the essence, ensuring that children receive a swift diagnosis and treatment. It also has the potential to be implemented in a wider range of clinical settings, including those with limited resources, due to its lower computational demands compared to more complex models. In essence, MobileNet-V3 helps bridge the gap between the inherent challenges of pediatric pneumonia diagnosis and the need for accurate, timely, and

resource-efficient detection. This makes it a valuable tool in the clinical management of pediatric pneumonia. Moreover, studies like [37, 38] have concluded that despite following the best practice in manual radiography in children, up to 40% of radiography is done with little or no benefit to the patients. This situation arises because current healthcare systems often fail to fully appreciate the radiologist's role. Instead of recognizing the value of accurate diagnostics, the focus tends to be on the costs associated with radiological tests, which sidelines the true importance of radiology in patient care outcomes. In clinical practice, the heavy workload means that radiologists and referring physicians rarely discuss the best imaging approach for each patient, leading to a mismatch between the patient's needs and the imaging performed. This disconnect can result in unnecessary or suboptimal imaging, contributing to the high percentage of low-benefit procedures. The effectiveness of an imaging procedure is also dependent on the quality of the radiology report, which must be clear, accurate, and useful for guiding treatment. However, how these reports are interpreted and acted upon involves various healthcare providers, not just radiologists. Misinterpretations of the information provided can further reduce the effectiveness of the imaging, adding to the instances where radiology does not serve its full potential in aiding patient care [37-39].

To tackle these challenges and to fasten the classification process, automated tools and techniques should be used. In this paper, I have implemented the MobileNeT-V3 architecture over various publicly available datasets of child chest-x ray images. This study has chosen MobileNeT-V3 for several reasons. Firstly, its lightweight structure allows it to process images rapidly, which is crucial in clinical settings where quick decision-making can significantly impact patient outcomes. MobileNet-V3 achieves this by using a smaller model size and fewer computations than traditional neural networks, without a substantial drop in accuracy. Secondly, MobileNet-V3's efficiency does not demand extensive computational resources, making it accessible for use in diverse clinical environments, including regions with limited healthcare infrastructure. This can be particularly beneficial for pediatric care, where timely and accurate diagnosis of conditions such as pneumonia is critical for effective treatment. An automated classification technique like as suggested in the paper will also provide cost-effective treatment to the patients followed by enhancement in the clinical governance. This is because MobileNetV3 is designed to be incredibly efficient, allowing for the rapid processing of images. While specific speed gains can vary depending on the hardware used, it's generally understood that MobileNetV3 can classify images in real-time or near-real-time on mobile devices. This speed is a substantial improvement over older convolutional neural network models that require more computational power and time to analyze images. Furthermore, the speed and accuracy of MobileNetV3 directly translates to improved patient care. Quick, reliable diagnoses mean treatment can start sooner, which is especially crucial for conditions like pneumonia in children, where every minute counts. Additionally, the model's ability to handle variable image quality means fewer repeat scans are necessary, reducing the child's exposure to radiation and stress. Finally, the model's efficiency means less computational resources and time are needed for image analysis, which can reduce the operational costs for healthcare providers. With fewer repeat scans needed due to the model's robustness to image quality, there are savings on consumables

and less wear on imaging equipment. Moreover, faster, and accurate diagnoses lead to optimized use of hospital resources, potentially shortening hospital stays and reducing the overall cost of care. These claims are further supported by the authors of the studies [38, 39], where they confirmed that automated techniques fasten the classification process by discussing the example of bone X-ray rather than using time consuming manual classification techniques like Greulich-Pyle or Tanner-Whitehouse technique. Thus, these benefits of automated classification techniques will act as the base of this study.

4. METHODOLOGY

This study has leveraged two extensive labelled datasets, collected from the references [40-43] to assess the performance of the MobileNet-V3 architecture on a broad spectrum of X-ray images. The first dataset encompasses over 14,000 X-ray images, categorized into training, testing, and validation subsets, with a total of 3814 normal and 3875 pneumonia-infected X-ray images for training, 350 normal and 390 pneumonia cases for testing, and 8 images for each class in the validation set. An additional 2500+ X-ray images are included under 'Other Data'. The second dataset includes over 5,000 images, with 1349 normal and 3883 pneumonia X-rays for training, along with 234 normal and 390 pneumonia images for testing. These datasets are publicly accessible and have been integral in developing a robust model that addresses the data scarcity issue highlighted in previous studies. While the resolution of the images is not explicitly stated, it is implied that they are of high quality, as this is a prerequisite for the accuracy of deep learning models. The comprehensive classification of the datasets is meticulously detailed in Tables 2 and 3 of this paper, ensuring transparency and reproducibility of research findings.

Table 2. Classification of dataset 1

Training Data	Normal X-ray Images = 3814, Pneumonia Infected X-rays = 3875
Testing Data	Normal X-ray Images = 350, Pneumonia Infected X-rays = 390
Validation Data	Normal X-ray Images = 8, Pneumonia Infected X-rays = 8
Other Data	2500+ X-ray Images

Table 3. Classification of dataset 2

Training Data	Normal X-rays = 1349, Pneumonia Infected X-rays = 3883
Testing Data	Normal X-rays = 234, Pneumonia Infected X-rays = 390

The accuracy of proposed deep learning model is highly dependent on the availability and quality of the X-ray images it is trained with. In previous research highlighted by literature review, the limited size of datasets raised concerns about the precision of the resulting models. To address this, my study incorporates two extensive datasets. By utilizing a larger and more diverse dataset, this model is better equipped to accurately identify pneumonia, which improves its reliability and applicability in a clinical setting. In preparing chest X-ray image dataset for the MobileNet-V3 model, I employed a series of preprocessing steps to ensure the data was in the best possible state for training. Initially, I cleaned the dataset by

removing any corrupt or unreadable images. I then standardized the resolution across all images, which is essential for maintaining consistency as the model learns. To highlight the critical features for pneumonia detection, I adjusted the contrast in each image, making patterns such as lung textures and fluid opacities more distinguishable. Data augmentation was a key transformation step, where I applied various alterations to the images, such as rotating, zooming, and flipping them. This process creates a more comprehensive set of training examples, teaching the model to recognize pneumonia in a variety of appearances and orientations. Following augmentation, I normalized the pixel values in the images to have a consistent scale, allowing for faster and more stable model training. Finally, I divided the dataset into distinct training, validation, and testing sets. This division is crucial to evaluate the model's learning capability, hyperparameter settings, and ultimately its diagnostic accuracy on new, unseen images. Each of these preprocessing and transformation steps was meticulously carried out to enhance the model's learning efficiency and improve its predictive performance in real-world clinical applications.
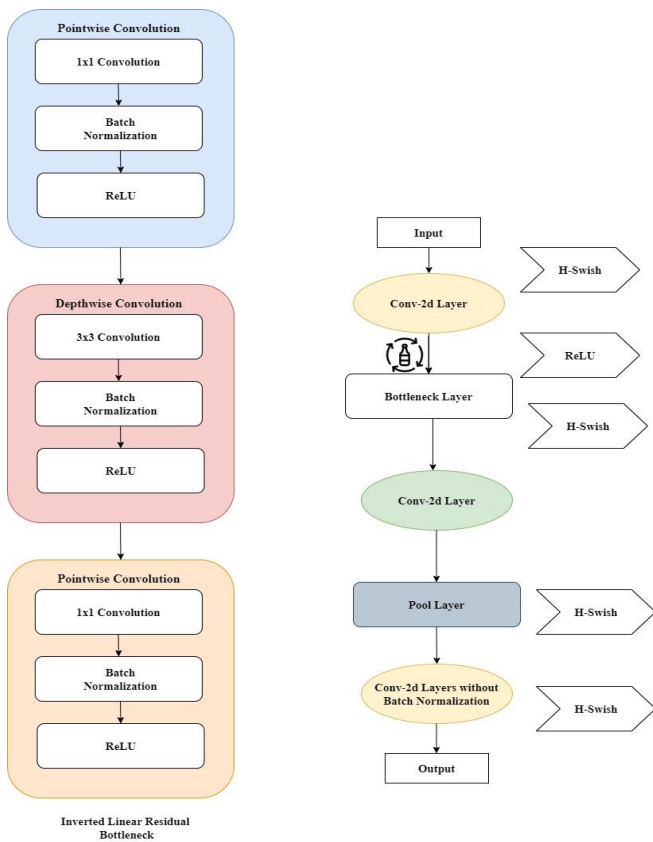

Figure 1. Schematic representation of the general MobileNeT architecture

In this paper, I have used the MobileNeT-V3 architecture for training the pneumonia X-ray images and classify them as pneumonia and normal class respectively. MobileNeT-V3 is a convolutional neural network (CNN), tuned to mobile phone CPU through a combination of hardware-aware network architecture search scheme and complemented by the NetAdapt algorithm [44]. Initially, MobileNeT architecture was derived from the VGG network by addition of depth wise separable convolutions. The latest version at the time of this study is MobileNeT-V3 which is designed after network optimizations to increase its efficiency [45]. According to the

resource usage by the network, MobileNeT is classified in to two types, MobileNeT-V3-Small and MobileNeT-V3-Large. However, I have followed a generalized architecture for this paper as described by Qian et al. [45]. This generalized architecture is depicted by Figure 1.

The various characteristics and layers of the network architecture [45-47] is explained as:

• Depth-wise Separable Convolutions: This is introduced in the network to increase the efficiency of computation. The convolution is divided into two parts. In first part, a single convolutional filter is applied to the depth wise convolution (DWC). In second part, 1×1 convolution is applied to all the channels of the output generated by the DWC. DWC increases the speed of the overall network and is the heart of all the versions of MobileNeT.

• Linear Bottlenecks: This feature is used by the MobileNeT architecture to extract features from a high dimensional space without the loss of information. This layer consists of a 1×1 filter, integrated with a linear activation function. This is because the transformation of ReLU generates a non-linearity in the network, causing the loss of information.

• Inverted Residual Blocks: To extract the necessary information efficiently, ReLU layers are replaced by the bottleneck layers. MobileNeT utilizes shorter paths to prevent loss of gradient and explosion. An inverted residual block (IRB) is found to be valid to act almost similar to the residual blocks. This feature helps in reducing memory costs.

• Neural Architecture Search (NAS): This is performed to determine the optimal architecture for a constrained hardware platform. During NAS, it constructs a search space for the neural network architecture for efficiently searching in the hierarchical search space with reinforcement learning to get the best structure of the model for specific tasks [48]. The general representation of the NAS is given by Figure 2 [49].
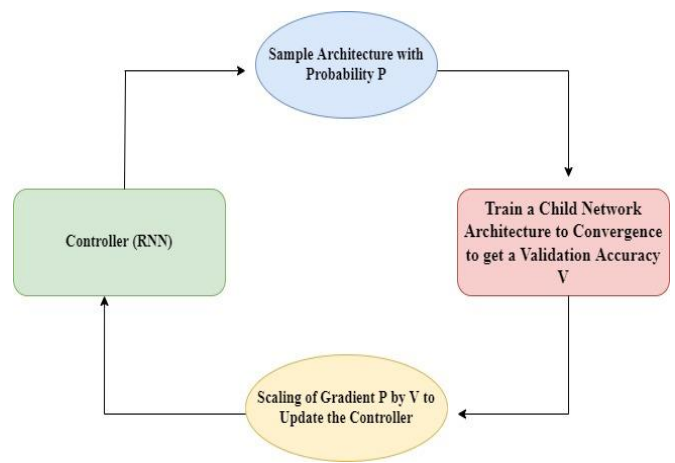

Figure 2. Representation of NAS

The network described by figure 2 works on the similar procedure on which the reinforcement learning works. The controller depicted in the figure gets updated according to the reward function. Normally, the overall model moves from the current state to the state where this reward function is increasing [44, 49]. This reward function is defined as [44]:

$$Accuracy(X) * \left(\frac{Latency(L)}{TargetLatency(T)} \right)^{y} \qquad (1)$$

where, X = resultant model from search and y is a constant.

This resultant network architecture obtained after NAS is refined layer wise using the NetAdapt algorithm. The NetAdapt algorithm works on each filter of every convolution, as described by algorithm 1 [50]. Algorithm 1 optimizes the number of filters for each convolution and chooses the highest accurate model. The best model among the M models with peak highest accuracy is chosen that comes from working on one of the convolutional layers. An improved version of this algorithm is used by MobileNeT network, starting with the NAS output. After this step, a set of proposals are generated with a reduced latency as compared with the previous models. Further, weights are set for new proposals by utilizing the weights generated from previous networks as well as the random initialization of new filters. Finally, the selected proposals are finer tuned until the target latency is achieved [44-50]. To increase the efficiency of the overall architecture, a few improvements to the network is made. These network improvements can be made in two ways, removal of layers and non-linearity/swish. Layer removal is applied to the initial layers and last layers. in the last layer block, 1×1 expansion layer is taken from the inverted residual block and is moved past the pooling layer. Furthermore, it is seen that the expansion layer takes higher computation time but since it is moved behind the pooling layer, there is no need to compress by projection layer from the last layer of the previous block and the projection layer along with the filtering layer can be removed from the previous bottleneck layers [44]. This helps in saving a lot of computational time. The second way to improve network is swish non-linearity. Swish non-linearity is used to improve the accuracy of the model.

Algorithm-1: NetAdapt Algorithm
Input: A Pretrained Network → N_0 (with M conv and FC layers), A Resource Budget → B, A Resource Reduction Schedule → ϕ
Output: λ → An Adaptive Network Meeting B
1. j = 0
2. β_j = empirical measurement of N_j
3. **while** β_j > B **do**
4.   CON = $\beta_j - \phi$;
5.   **for** m from 1 to M **do**
6. $N_{filter\ m}$, $\beta_{simp\ m}$ = select number of filters (N_j, m, CON)
7. $N_{simp\ m}$ = CWF (N_j, m, $N_{filter\ m}$);   // CWF = choose which type of filter
8. $N_{simp\ m}$ = STFT ($N_{simp\ m}$);   // STFT = short-term fine tuning
9. N_{j+1}, β_{j+1} = Choose Highest Accuracy (N_{simp}, β_{simp})
10.   j = j+1
11. λ = LTFT (N_j);   // LTFT = Long-Term Fine Tune
12. **return** λ

However, it is shown by the references [44-48] that sigmoid function is expensive computationally which could affect the accuracy of the architecture. Thus, the authors [44] suggested a h-swish or hard swish non-linearity defined as:

$$hard - swish(Z) = Z\frac{ReLU6(z+3)}{6} \qquad (2)$$

This function helps in improving the computational power of the model.

5. RESULT ANALYSIS

Before training my model, I started by creating helper functions. According to the standard definition of the helper functions explained by Creating Helper Functions [51] is, a helper function is a function that does the computation of other functions to save time. These functions are used to make the models easier to understand by giving descriptive names to the computation. These functions also allow us to reuse the computations like general functions. Various helper functions that are utilized in my model are summarized in Table 4 along with their features.

Table 4. Summarization of helper functions

Helper Functions	Features
Create tensor board callback	This function helps to prevent model overfitting, visualize the model training, saving checkpoints and creating tensor board.
Plot loss curve	This helps in visualizing the loss curves of the model.
Unzip data	This helps in extracting the data.
Compare history	This function is used for comparisons.
Walk through dir	This function helps in checking of contents of various directories in the dataset.
Pred and plot	This function helps in rounding the probability generated and to better visualize the plots.

Once the helper functions are defined, data preprocessing and data transformation is done. Initially the data collected from the real world contains missing values and is in unable format by the deep learning models. Data preprocessing helps in cleaning the dataset and converting it to a usable format for the deep learning models. Data transformation helps to increase the data analysis process and improve the data driven decision making capability. It also helps in determining the structure of the data given to the model for training. Before the model training process, the dataset is split into training, testing, and validation sets. The training set is used for training the model, testing set is used for testing the model, and validation set is used for fine tuning the parameters of the deep learning model. Finally, the model is evaluated based on test data or data which is unseen previously. For training, I have used the same model architecture and parameters for both dataset 1 and dataset 2. The rationale for employing the same MobileNet-V3 architecture and training parameters across both datasets is anchored in the need for consistency and comparability. Using the same architecture allows us to directly compare the performance of the model on different datasets, ensuring that any variation in results can be attributed to the data itself rather than differences in model configuration. MobileNet-V3 is designed to be versatile and efficient, which makes it well-suited to handle the variability inherent in chest X-ray images from diverse sources. Furthermore, consistency in parameters such as batch size, image dimensions, and learning rate ensures that the model trains under uniform conditions, allowing us to isolate and evaluate the effectiveness of the architecture without the confounding effects of differing training regimes. This approach simplifies the evaluation process and provides clearer insights into the model's capabilities and potential areas for optimization specific to pediatric pneumonia diagnosis from X-ray images. During training, the data is taken as size 224 × 224 and with batch size

32. The shuffling of data is enabled as True, and the class mode was chosen as categorical. This means that the model will generate a 2D output as mutually exclusive labels, for example, either NORMAL or PNEUMONIA. A seed value of 42 is assigned for each training, testing, and validation set to enable reproducibility of the results. Finally, before training the model, an error level analysis (ELA) is done. ELA is a technique that is used to enhance the efficiency of differentiating copy-move images produced by deep-fake from real images [52]. I have also done the ELA for initial image analysis for both the datasets. The outcomes generated by the ELA process are depicted by Figure 3.
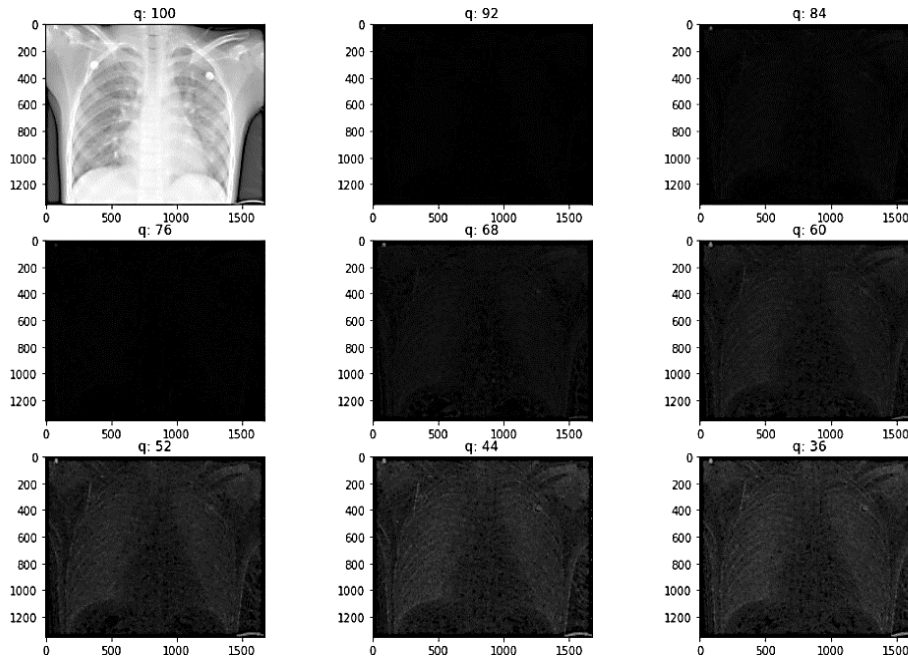

Figure 3. Illustration of error rate analysis (ELA) for X-ray images of dataset 1
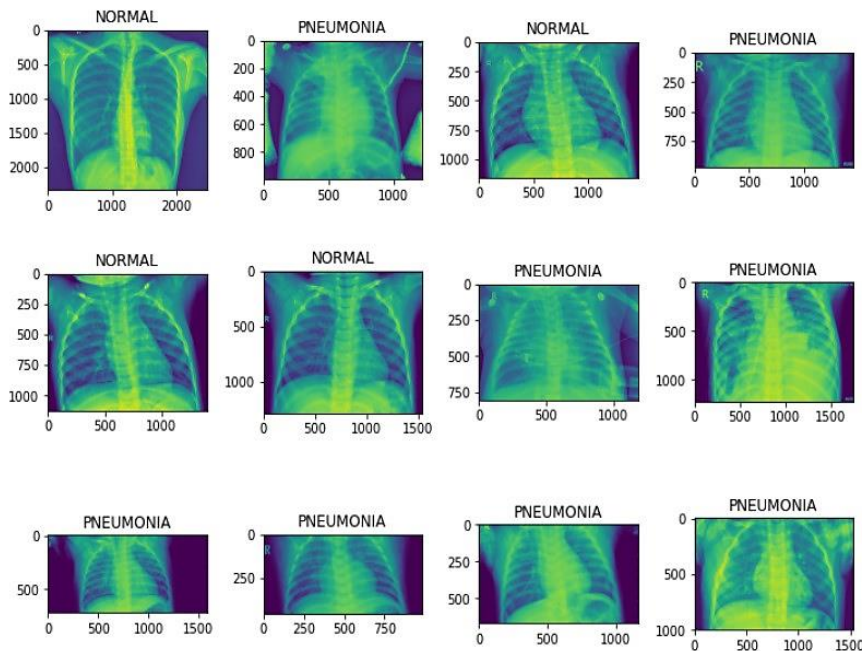

Figure 4. Random images from the processed dataset 1

Now the data is ready for training. The model is trained for 30 epochs with 2 output layers and input size of $224 \times 224 \times 3$ and batch size of 32. The default ImageNet weights were used while training the model with an average pooling layer. During training, a model checkpoint callback is used to save the model/weights in a checkpoint so that the model can be reloaded to continue for training after some time. In this model, I have defined the checkpoint callback as save the best only.

This is because, I want to save the best outcomes to the file checkpoint. Another parameter defined is monitor = val_accuracy. This helps us in monitoring the validation accuracy generated by the models. Another parameter, verbose is set to 0, which means that no message will be prompted to the user when callback actions are taken. I have also set up the early stopping callback which is defined to stop training the model if the validation loss during the training of the model

doesn't shows any improvement for 3 epochs. I have also used optimizers to enhance the accuracy of the model. Generally, a deep learning model is difficult to optimize due to complex architecture. However, certain optimizers are still able to handle the complex architecture and generate efficient results. Optimizers are used to adjust the weights of a deep learning model and maximize the loss function. A loss function is used for measuring the performance of the model. An optimizer should be used during training of the deep learning model. In this paper, I have used the Adam optimizer with a vale set at 0.00001, whereas the loss is defined as categorical cross entropy.

This optimizer was chosen because the results generated by this optimizer are better than another optimizer, have a fast computational time and require a few tuning parameters. In this study, I did not employ explicit regularization techniques such as dropout or L1/L2 regularization. This decision was based on the extensive and varied nature of datasets, which inherently reduces the risk of overfitting. Moreover, the use of data augmentation and a robust architecture like MobileNet-V3, which is designed to generalize well, further mitigates the need for additional regularization. Before starting training of

the model, I need to look at some random data from the representation of the dataset is shown in Figure 4.

Now the model is trained for 30 epochs with input size 224 × 224 × 3 and batch size 32. The outcomes generated from training the model is recorded in Table 5. A summary of the training parameters and methods used is provided in Table 6.

The overall training accuracy achieved over dataset 1 is 96%. In 30 epochs, each epoch was run for 118 rounds. Since it is difficult to analyze the results with Table 5. So, a plot between validation loss and validation accuracy and training loss and training accuracy is plotted for better visualizations of results. these plots are shown by Figure 5 and Figure 6 respectively.

Once the model is developed, it needs to be validated using various datasets. Normally, there is a limited quantity of accurate data available for training purposes and validating it is necessary to be able to develop a reliable model. To evaluate the reliability of the model developed, the plots mentioned by Figure 5 and Figure 6 are used. During model training, both the accuracy and loss for the validation data can be different. With increase in epochs, the loss should decrease and accuracy increases.

Table 5. Results generated from dataset 1

# Epochs	Loss	Accuracy	Val Loss	Val Accuracy
1	0.5130	0.7492	0.3906	0.8226
2	0.3686	0.8412	0.2909	0.8964
3	0.2917	0.8826	0.2405	0.9092
4	0.2534	0.8949	0.2146	0.9124
5	0.2256	0.9120	0.1965	0.9177
6	0.2049	0.9210	0.1851	0.9199
7	0.1966	0.9197	0.1767	0.9252
8	0.1876	0.9301	0.1706	0.9295
9	0.1838	0.9269	0.1630	0.9327
10	0.1783	0.9312	0.1576	0.9359
11	0.1670	0.9386	0.1533	0.9359
12	0.1585	0.9402	0.1487	0.9370
13	0.1537	0.9442	0.1447	0.9444
14	0.1499	0.9442	0.1404	0.9391
15	0.1481	0.9469	0.1374	0.9434
16	0.1426	0.9456	0.1351	0.9444
17	0.1400	0.9520	0.1337	0.9423
18	0.1383	0.9490	0.1302	0.9487
19	0.1354	0.9528	0.1271	0.9476
20	0.1334	0.9512	0.1248	0.9476
21	0.1297	0.9562	0.1236	0.9487
22	0.1286	0.9528	0.1217	0.9541
23	0.1240	0.9562	0.1185	0.9551
24	0.1198	0.9562	0.1177	0.9509
25	0.1214	0.9554	0.1161	0.9519
26	0.1145	0.9605	0.1153	0.9530
27	0.1101	0.9624	0.1129	0.9551
28	0.1108	0.9592	0.1113	0.9562
29	0.1101	0.9629	0.1104	0.9562
30	0.1078	0.9618	0.1096	0.9562

Table 6. Summarized of the training parameters and methods used

| Parameter/Method	Description	Value/Setting
Loss Function	Categorical cross-entropy	-
Optimization Algorithm	Adam optimizer	Learning rate: 0.00001
Number of Epochs	Total epochs for training	30
Regularization	Not applied due to dataset diversity	-
Early Stopping	Stops training if no improvement in validation loss	Patience: 3 epochs
Model Checkpoints	Saves the best model based on validation accuracy	'Save best only'
ELA for Image Analysis	Pre-training analysis for image authenticity	Conducted for both datasets
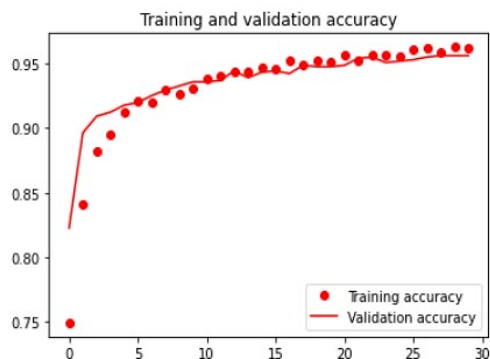
Figure 5. Training vs validation accuracy for dataset 1

However, certain cases are possible for validation accuracy and validation loss. If the validation loss increases, then the validation accuracy decreases. This means that the model developed contains errors and the model is not learning properly. If the validation loss increases, then validation accuracy also increases. This means that the model could be overfitting. However, for my model, as evident from Figure 5 and Figure 6, validation loss decreases and validation accuracy increases. This means that my model is working fine and is learning things accurately. This model achieved an overall testing accuracy of 95.82% with an overall loss of 0.11627 on dataset 1, which shows the effectiveness of this model.
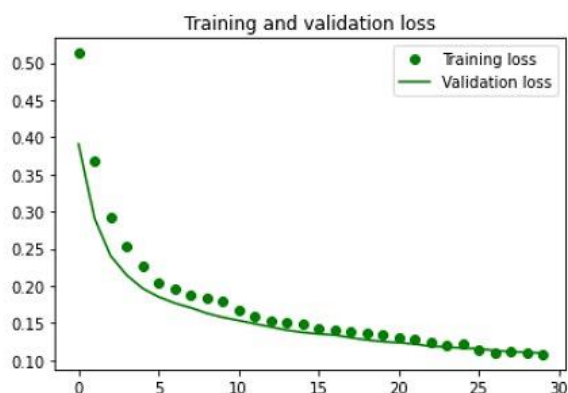

Figure 6. Training vs validation loss for dataset 1

For further evaluation of reliability of the model, I have used certain metrics. These metrics are defined as:

Precision- it is defined as the fraction of correct predictions (True Positive) from the total results. Precision measures the accuracy of the positive predictions made by the model. In other words, it tells us what proportion of positive identifications was actually correct. For this study, this would mean how many of the X-rays identified as having signs of pneumonia truly had the condition. It is calculated as:

$$Precision = \frac{TruePositive}{TruePositive + FalsePositive}$$

Recall- it is defined as the fraction of true positives from the total number of true positives and false negatives. It assesses the model's ability to find all the relevant cases within a dataset. It's the proportion of actual positives the model correctly identified. In the context of pneumonia detection, recall would indicate how many X-ray images of pneumonia the model was

able to correctly identify out of all the cases that were actually pneumonia. It is given as:

$$Recall/Sensitivity = \frac{TruePositive}{TruePositive + FalseNegative}$$

F-1 Score- The F1 score is a way of combining the precision and recall of the model into a single metric that balances both concerns. It's the harmonic mean of precision and recall, giving us an overall indication of the model's accuracy. An F1 score reaches its best value at 1 (perfect precision and recall) and worst at 0. It is evaluated by,

$$F - 1\ Score = \frac{2 * (TruePositive * FalsePositive)}{(TruePositive + FalsePositive)}$$

Sensitivity- It is used for evaluating the ability of the model to predict the true negative in each category. This is given by,

$$Specificity = \frac{TrueNegative}{TrueNegative + FalsePositive}$$

These metrics are especially important in medical diagnostics, where it's crucial not only to correctly identify conditions but also not to miss any cases that could lead to further health complications. These metrics will also act as the basis of the confusion matrix as depicted by Figure 7. The values generated from these evaluation metrics are recorded in Table 7.

Table 7. Evaluation metrics for dataset 1

	Precision	Recall	F-1 Score	Support
NORMAL (0)	0.94	0.91	0.93	333
PNEUMONIA (1)	0.96	0.98	0.97	839
Macro Average	0.95	0.94	0.95	1172
Weighted Average	0.96	0.96	0.96	1172
Specificity = 0.97				

For better understanding of these metrics, a confusion matrix is drawn and depicted by Figure 7. A confusion matrix helps in visualizing the summary of the predictive performance of the developed model. Classification accuracy generated by the model can be misleading if there are more than two classes in the dataset. Analysis of confusion matrix gives us a better understanding of the working and classification generated by my model.

The predictions made by the model over dataset 1 are depicted by Figure 8. Furthermore, to expand the scope of this study to large datasets and to enhance reliability of my model and to resolve the challenge of data scarcity, the model was again trained on new dataset 2. The model training process was carried out by keeping the architecture the same as for the dataset 1. The training parameters and batch size as described for dataset 1 are kept same for dataset 2. The model is again trained for 30 epochs of 286 rounds for dataset 2. The results generated by my model for dataset 2 are recorded in Table 8. The error analysis ELA for dataset 2 is shown by Figure 9.
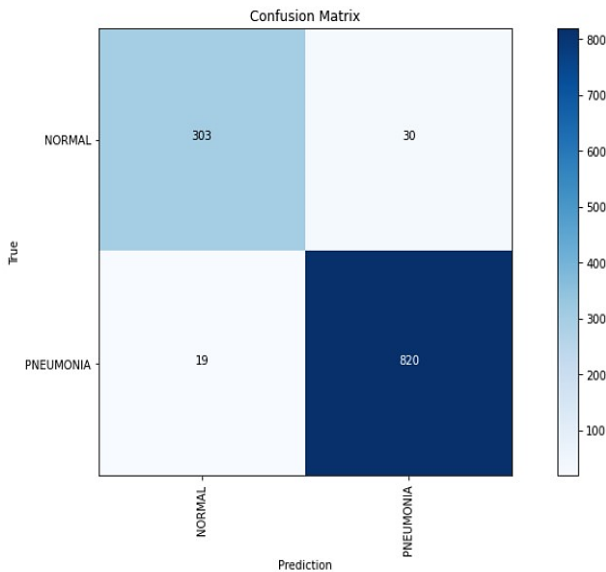
Figure 7. Confusion matrix for dataset 1

Figure 9 showcases the Error Level Analysis (ELA) for dataset 2, which serves as a tool to visualize the consistency of image quality across a dataset by detecting variations in error levels due to image compression. In medical imaging, and particularly in the analysis of X-rays with deep learning models, ELA can offer insights into the model's sensitivity to image quality. From the ELA, I observe a range of error levels indicated by 'q' values, where a high 'q' value denotes a lower level of compression and potentially higher image quality. As the 'q' value decreases, I see an increase in the granularity of the image, which is indicative of higher compression or lower image quality. When considering the performance of proposed deep learning model, these ELA results do not necessarily point to errors in diagnosis but rather provide a measure of the variability in image quality that the model encounters. Analyzing the patterns in the ELA images, there is no evident correlation suggesting that the model's accuracy is compromised at specific error levels. This implies that proposed model maintains a consistent performance across the spectrum of image qualities present in dataset 2. The model's robustness to such variations can be attributed to its training on a diverse set of images, which likely included varying degrees of compression artifacts akin to those visualized in the ELA. This training would have equipped the model to differentiate between pertinent features of pneumonia and artifacts introduced by image compression. Furthermore, the accuracy metrics such as precision, recall, and F1 score for proposed model remain high, affirming that the model's diagnostic capabilities are not hindered by the variations in image quality represented by the ELA. It suggests that the model has effectively learned to identify the key features indicative of pneumonia, despite the noise introduced by different compression levels. In conclusion, the ELA serves as a testament to the model's ability to handle a wide array of image qualities within dataset 2. The absence of a negative impact on model accuracy due to varying error levels suggests that the model is well-tuned and could potentially be deployed in a real-world clinical setting where X-ray image quality can be variable. Furthermore, to enhance the model's performance, several areas for improvement are identified. Image preprocessing can be optimized to ensure the highest quality of images are used for model training, focusing on noise

reduction to mitigate the impact of compression artifacts. Data augmentation should be expanded to include a broader spectrum of image qualities, which would train the model to handle real-world variations in X-ray images more effectively. The feature extraction capabilities of the model could be refined, helping it to better differentiate between important features indicative of pneumonia and noise resulting from image compression. Adjustments to the MobileNet-V3 architecture could be explored to optimize its ability to process images of varying quality levels. Finally, training the model on a more diverse set of datasets would expose it to a wider array of pathologies and image conditions, further improving its diagnostic accuracy and generalizability.

Table 8. Results generated for dataset 2

# Epochs	Loss	Accuracy	Val Loss	Val Accuracy
1	0.5126	0.7390	0.2802	0.9087
2	0.2710	0.8957	0.1979	0.9305
3	0.2133	0.9149	0.1681	0.9379
4	0.1830	0.9299	0.1510	0.9436
5	0.1638	0.9410	0.1386	0.9506
6	0.1526	0.9438	0.1328	0.9576
7	0.1435	0.9466	0.1230	0.9589
8	0.1321	0.9524	0.1166	0.9615
9	0.1299	0.9515	0.1126	0.9615
10	0.1226	0.9572	0.1073	0.9650
11	0.1173	0.9586	0.1039	0.9663
12	0.1100	0.9619	0.1014	0.9677
13	0.1080	0.9603	0.0975	0.9681
14	0.1044	0.9649	0.0949	0.9690
15	0.0993	0.9647	0.0921	0.9698
16	0.0959	0.9662	0.0900	0.9712
17	0.0939	0.9667	0.0881	0.9712
18	0.0938	0.9677	0.0861	0.9725
19	0.0866	0.9705	0.0842	0.9738
20	0.0853	0.9695	0.0826	0.9738
21	0.0832	0.9715	0.0804	0.9764
22	0.0784	0.9724	0.0789	0.9674
23	0.0773	0.9741	0.0771	0.9773
24	0.0725	0.9742	0.0756	0.9781
25	0.0765	0.9729	0.0743	0.9768
26	0.0718	0.9740	0.0738	0.9799
27	0.0675	0.9776	0.0718	0.9790
28	0.0660	0.9777	0.0704	0.9799
29	0.0650	0.9786	0.0717	0.9790
30	0.0644	0.9795	0.0679	0.9786

Random data representation for dataset 1 as represented by Figure 4 is not necessary for dataset 2 as that representation was only for the understanding of how the data looks like. Chest X-rays are similar, so there is no need to show random images from dataset 2. One can proceed with the model training results directly. Various plots involved in interpretation of results are shown by Figure 10 and Figure 11 respectively.

The developed model achieved a training accuracy of 99% and testing accuracy of 97.80% with an overall loss of 0.06169 over dataset 2 which proves the efficiency of the model on large datasets. Furthermore, it is evident from Figure 10 and Figure 11 that validation loss decreases and validation accuracy increases. This means that my model is working fine and is learning accurately. For further evaluation of reliability of the model, I have used the same metrics that were used for evaluating results from dataset 1. The values obtained for dataset 2 is recorded in Table 9.
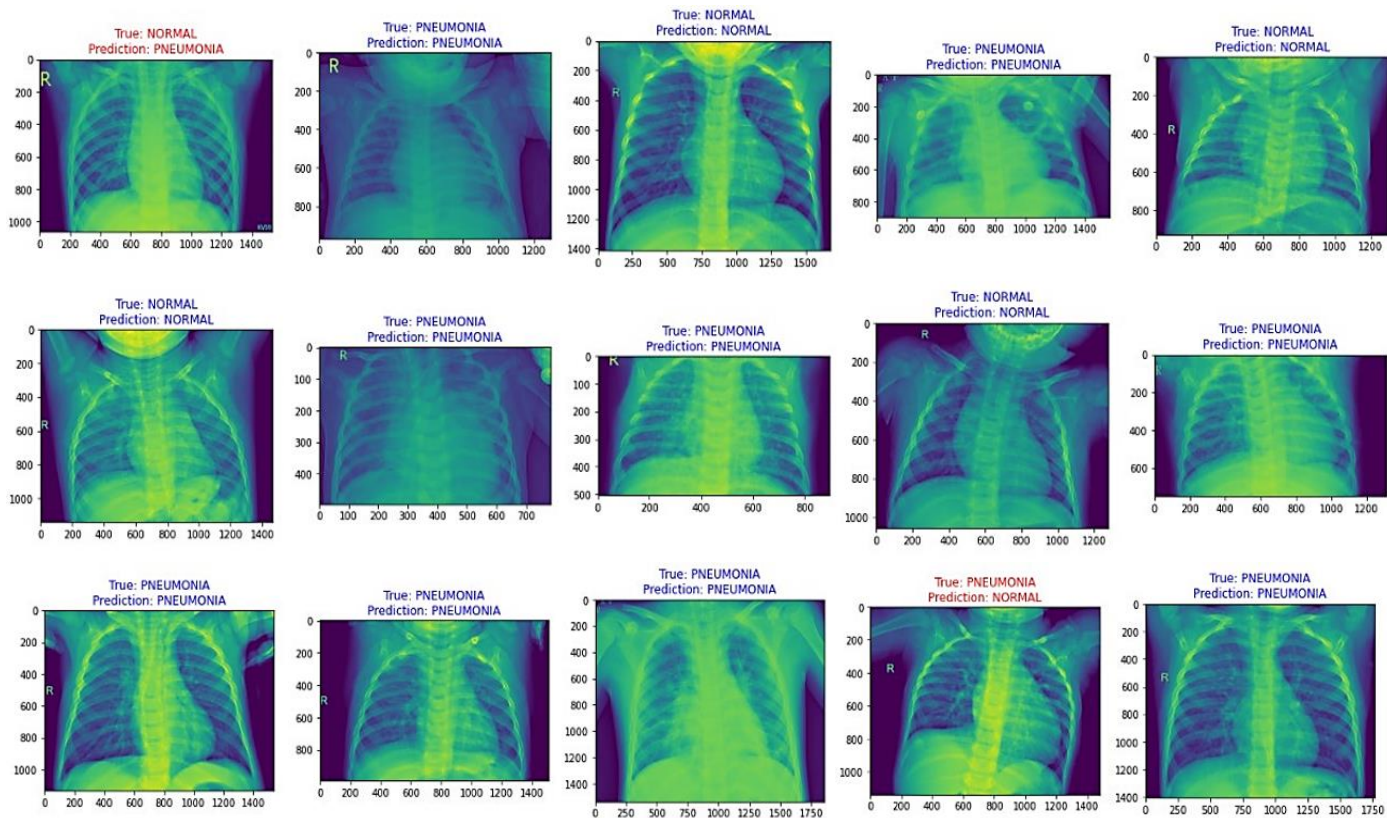
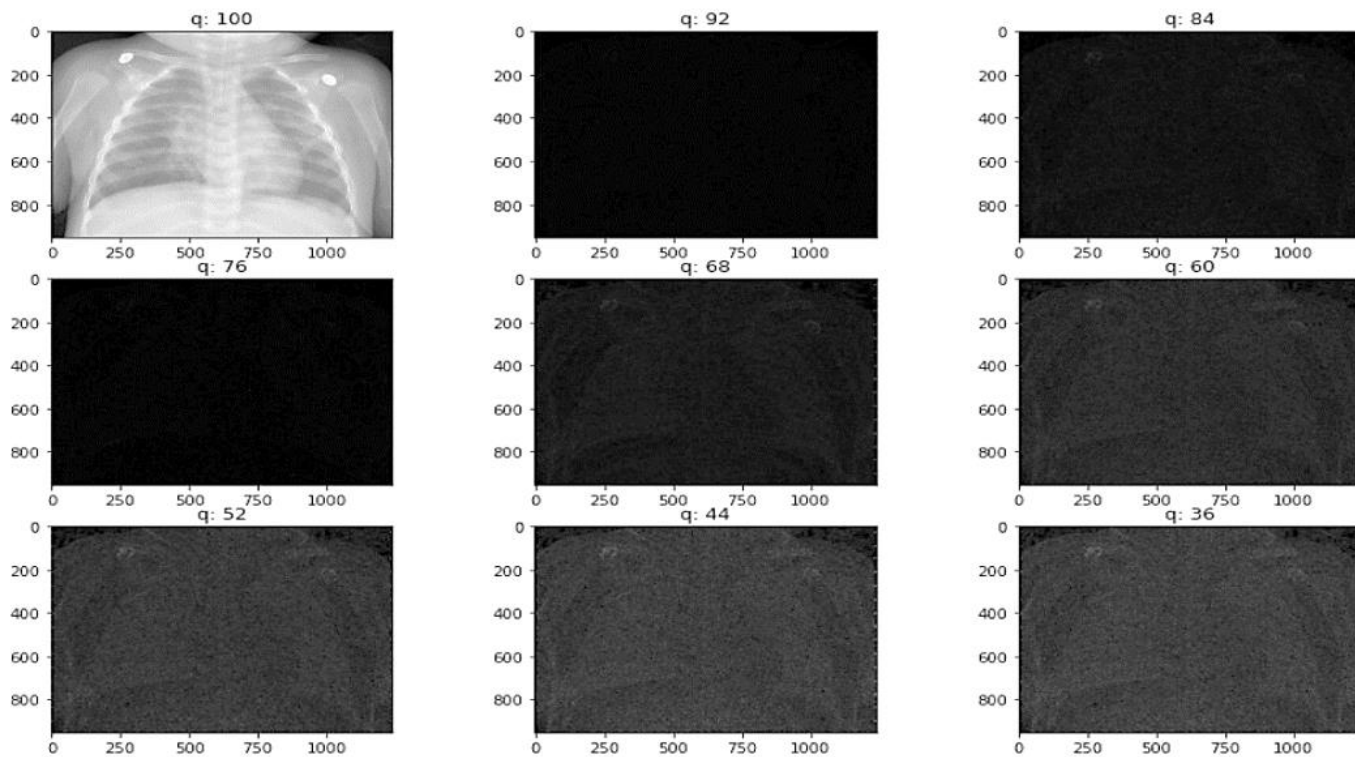**Figure 8.** Representation of prediction outcomes generated from dataset 1

Figure 9. Representation of error rate analysis (ELA) for dataset

Table 9. Evaluation metrics for dataset 2

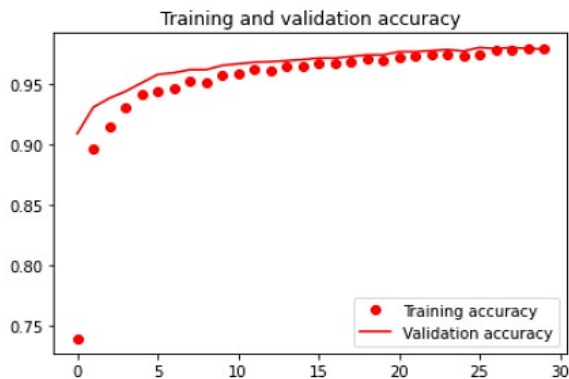|	Precision	Recall	F-1 Score	Support
NORMAL (0)	0.97	0.97	0.97	1080
PNEUMONIA (1)	0.98	0.98	0.98	1791
Macro Average	0.98	0.98	0.98	2861
Weighted Average	0.98	0.98	0.98	2861
Specificity = 0.98				

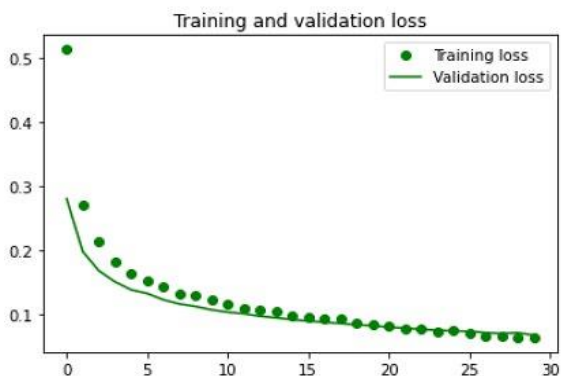**Figure 10.** Training vs validation accuracy for dataset 2

Figure 11. Training vs validation loss for dataset 2

For better understanding of these metrics, a confusion matrix is drawn and depicted by Figure 12. A confusion matrix helps in visualizing the summary of the predictive performance of the developed model. Classification accuracy generated by the model can be misleading if I have more than two classes in the dataset. Analysis of confusion matrix gives us a better understanding of the working and classification generated by my model.
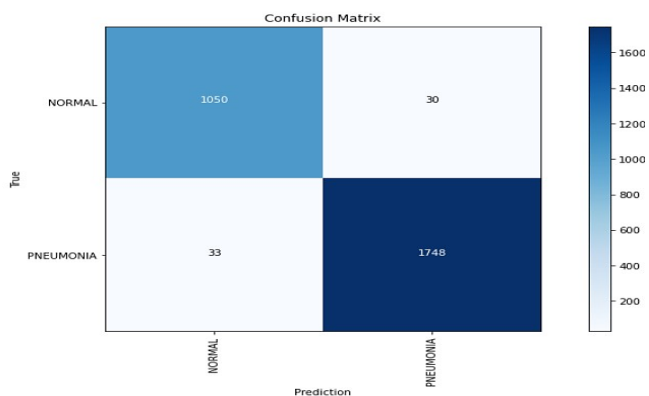

Figure 12. Confusion matrix for dataset 2

The predictions made by the model over dataset 1 is depicted by Figure 13 During my analysis, I found that time consumption is also a challenge for manual X-ray technique. Manual X-ray takes a lot of time from hours to days, from sample collection to final interpretation report generation. After these long procedures, the patient is refereed for treatment. This might not be good for young children. Furthermore, waiting for long time may prove to be fatal for

infants who might already infected with pneumonia. The technique suggested in the paper would reduce the time and speed up the treatment process. To support this, a comparative analysis between the time taken for training dataset 1 vs time taken for training dataset 2 is recorded in Table 10. In Table 10, the observed variation in training times between datasets 1 and 2 is attributed to several definitive factors. The primary factor is the dataset size; dataset 1 has fewer images compared to dataset 2, resulting in shorter training times per epoch. Additionally, the inherent complexity of the images in dataset 2 is higher, which necessitates more time for the model to process and learn from the intricate features present in the X-rays. Another contributing factor is computational resources; any discrepancies in the hardware specifications or resource allocation during the training process directly impact the training duration. Preprocessing and augmentation procedures also play a significant role; dataset 2 undergoes more rigorous preprocessing and augmentation, leading to longer training steps. Lastly, batch processing efficiency varies inherently with the data content, with dataset 2 presenting more challenging batches that extend the model's training time. These factors combined explain the consistent difference in training times, underscoring the need for tailored optimizations in the training process for each dataset.

Table 10. Comparative analysis of training time for both datasets

# Epochs	Training Time Per Step for Dataset 1 (in Sec)	Training Time Per Step for Dataset 2 (in Sec)
1	117.945	303.1
2	94.794	292.1
3	94.795	259.906
4	94.796	266.931
5	96.815	251.877
6	94.796	259.906
7	94.798	258.901
8	95.800	249.872
9	94.798	250.871
10	95.802	248.868
11	96.808	247.864
12	94.800	249.872
13	93.786	250.874
14	95.807	250.875
15	93.784	251.879
16	92.781	249.872
17	93.785	253.885
18	95.806	247.864
19	93.791	249.872
20	93.790	250.875
21	95.806	251.879
22	96.812	249.872
23	95.801	249.872
24	94.793	249.872
25	94.799	248.868
26	93.787	247.864
27	102.864	248.866
28	98.829	248.865
29	94.798	249.870
30	94.796	254.878

The average time taken for training dataset 1 is 97 seconds, that is, 1.62 minutes. The average time taken for training dataset 2 is 255 seconds, that is, 4.25 minutes. Thus, it can be concluded from the Figure 14 that the suggested technique reduces time drastically. Automated systems can generate high accurate results by processing larger datasets in less possible

time. Thus, the technique suggested in the paper is better than manual techniques. For further analysis of performance, the suggested technique is compared with some of the closely related studies in this field based on accuracy percentage. This comparison is done in Table 11.

From Table 11, it is evident that proposed model performs well and outperforms other models suggested by recent studies. The comparative analysis across various studies, including proposed own work, provides a diverse look at the methodologies applied, the datasets used, and the results achieved in the context of pneumonia detection using deep learning models.

Table 11. Comparative analysis of proposed work with recent works

| Study Number	Methodologies	Datasets	Results
[53]	15-layer CNN, MLD, Correntropy feature selection, One-class kernel ELM classifier	Radiopeadia images	95.1% accuracy, 95.1% sensitivity, 95% specificity, 94% precision
[54]	Pretrained AlexNet model, two-way/three-way/four-way classification	Public databases not explicitly mentioned	Up to 99.62% accuracy, 98.19% sensitivity, 100% specificity
[55]	Neuromorphic spiking neural network on AIRBiS	Collected test data for AIRBiS project	92.1% accuracy for pneumonia detection
[56]	Several pre-trained CNN models	CXR images	Over 80% accuracy, SqueezeNet achieved 81.62%
[57]	CNN architectures with/without data augmentation	Plain chest radiography	98.5% accuracy with DenseNet, 98.6% on validation with Inception
[58]	VGG16-based CADx system, conventional and mixup data augmentation	Two public datasets not explicitly mentioned	83.6% three-category accuracy, >90% sensitivity for COVID-19
[59]	TensorFlow-based CNN model	Chest X-ray images	Over 95% accuracy for pneumonia detection
[60]	DCNN with an explainable method	Chest X-ray dataset for COVID-19 classification	Average accuracy above 96% for COVID-19 pneumonia classification
Proposed Work	MobileNet-V3 architecture	Dataset 1: Over 14,000 images, Dataset 2: Over 5,000 images	95.82% accuracy for Dataset 1, 97.80% accuracy for Dataset 2
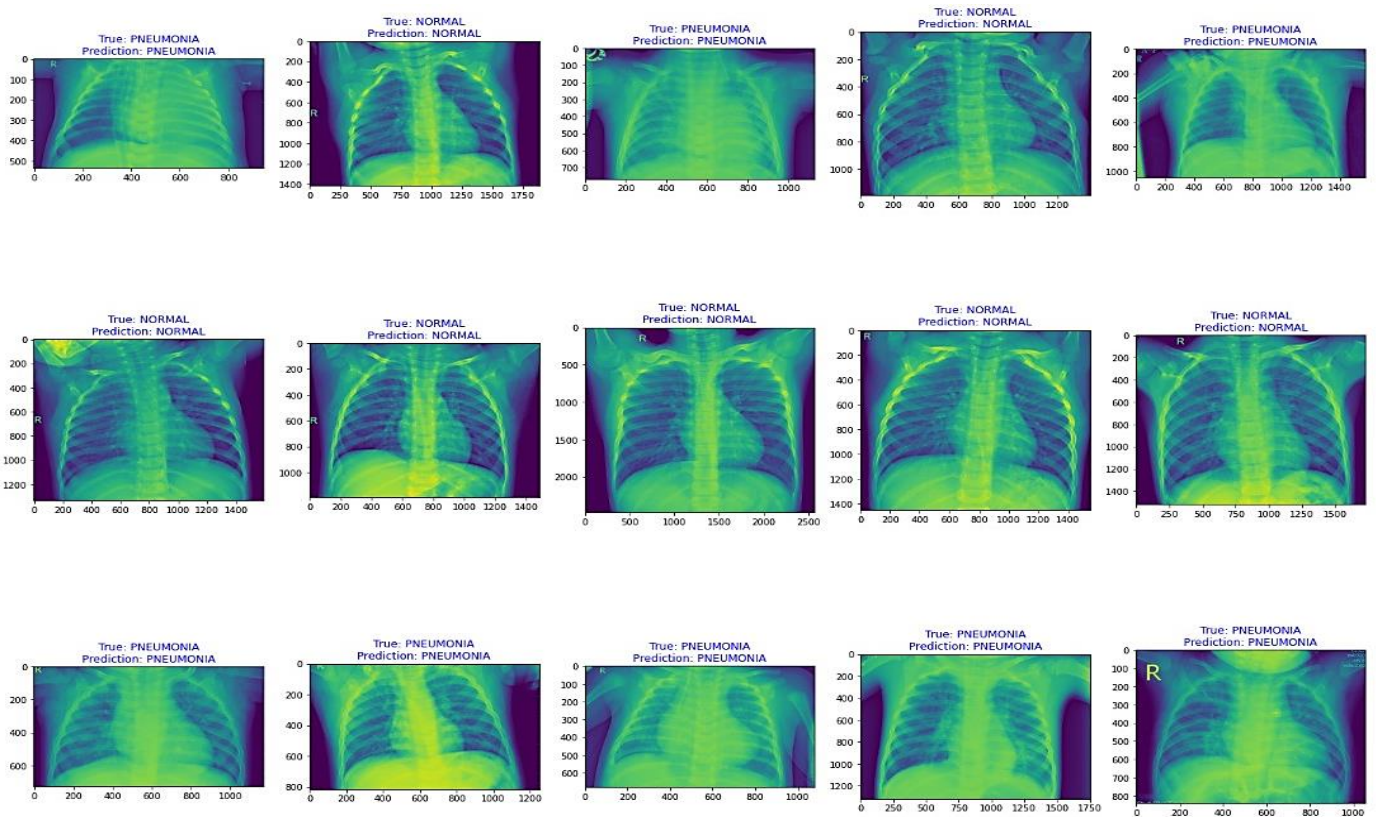
Figure 13. Representation of prediction outcomes generated from dataset 2

Starting with the paper of Khan et al. [53], the research utilizes a 15-layer convolutional neural network and employs a unique combination of deep feature extraction and the Max-Layer Detail (MLD) approach, enhanced by a Correntropy feature selection technique. This method, along with a one-class kernel extreme learning machine classifier, was tested on images from Radiopeadia and resulted in a balanced performance with accuracy, sensitivity, specificity, and precision all hovering around the 95% mark. Study [54] adopts a different strategy by leveraging the pretrained AlexNet model to classify various conditions of chest X-rays into multiple categories, ranging from two-way to four-way classifications. Their results were impressive, particularly in distinguishing COVID-19 pneumonia from other types,

achieving nearly perfect specificity in several instances. This indicates the AlexNet model's strong discriminative power when trained on comprehensive public databases. Study [55] introduces an innovative approach by implementing a neuromorphic spiking neural network within the AIRBiS framework, which is noteworthy for its suitability in edge computing environments due to its low-power requirements. While their accuracy is slightly lower than others, at 92.1%, the method stands out for its application in resource-constrained settings.
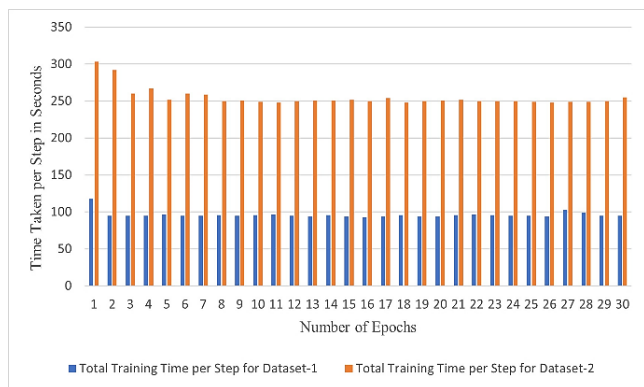

Figure 14. Comparative analysis between dataset 1 and dataset 2 based on training time

Study [56] explores the efficacy of various well-known pretrained CNN models in detecting pneumonia from chest X-rays. SqueezeNet, in particular, achieved the highest accuracy among the models tested, although the accuracy rate was the lowest among the studies at just over 80%. However, the speed of detection was a highlight for SqueezeNet, suggesting its potential for rapid diagnosis. Study [57] delves into the performance of different CNN architectures, both with and without data augmentation, on chest radiography images. They reported very high accuracy rates, especially with DenseNet and Inception models, indicating the strength of these architectures in feature learning and classification tasks. Study [58] used a VGG16-based CADx system, which combined conventional methods with mixup data augmentation, showing that this blend of augmentation techniques can be more effective than using just one type. The system achieved a respectable three-category accuracy of 83.6% and was particularly sensitive to COVID-19 pneumonia. Study [59] utilized a TensorFlow-based CNN model for pneumonia detection, achieving high accuracy. This study emphasizes the capability of deep learning algorithms, supported by CNNs, to analyze chest X-ray images with high precision. Study [60] reports on the use of a deep convolutional neural network with an explainable AI component to differentiate COVID-19 pneumonia in chest X-rays, achieving an average accuracy above 96%. The inclusion of explainability is significant, as it provides insights into the model's decision-making process. In contrast, proposed work utilized the MobileNet-V3 architecture across two large datasets, achieving 95.82% and 97.80% accuracy, respectively. The high accuracy across both datasets demonstrates the model's effectiveness and the benefits of using a large and diverse dataset to improve the model's generalization capabilities. In summary, these studies illustrate a range of deep learning approaches applied to pneumonia detection, each with varying degrees of complexity and success. The work in this study aligns with the higher-performing models, signifying the MobileNet-V3

architecture's suitability for this application due to its high accuracy and robustness across diverse image qualities. The comparative analysis underscores the importance of not only model selection but also the breadth and depth of the datasets used for training to achieve high performance in medical image analysis tasks.

6. CONCLUSION

Pneumonia remains one of the leading causes of mortality among infants and young children worldwide. Although it is a treatable condition, its rapid and accurate diagnosis is crucial for preventing deaths. Current automated diagnostic systems are predominantly tailored for adults and do not address the specific nuances of pediatric cases. My initial analysis highlighted a reliance on smaller datasets in existing techniques, which casts doubt on their applicability in real-world scenarios. Recognizing this, this study proposes a rapid classification method utilizing the MobileNet-V3 architecture. The results from the extensive dataset evaluation confirm that this technique not only boasts high accuracy but also overcomes the time constraints associated with manual diagnostic methods. The demonstrated reliability and efficiency of my approach indicates that it is ready for real-world deployment, potentially transforming the landscape of pediatric pneumonia diagnosis and significantly reducing mortality rates in young children. Moreover, proposed MobileNet-V3 model achieved accuracies of 95.82% on dataset 1 and 97.80% on dataset 2, which, when compared to the results of similar studies, showcases this model's enhanced performance. For instance, study [53]'s model achieved an accuracy of 95.1%, making this model's performance better by 0.72% on dataset 1 and by 2.7% on dataset 2. Study [55]'s reported accuracy is 92.1%, indicating this model outperforms it by 3.72% and 5.7% for datasets 1 and 2, respectively. Even when compared to study [56]'s best-performing model, SqueezeNet, which reached an accuracy of 81.62%, this model shows a substantial improvement of 14.2% and 16.18% for each dataset. This specific improvement in accuracy not only demonstrates the effectiveness of this model but also underscores the advantages of utilizing the MobileNet-V3 architecture and the extensive datasets were employed. Furthermore, The MobileNet-V3 architecture demonstrated in this study holds significant potential for real-world applications, particularly in enhancing diagnostic capabilities in medical facilities. Its high accuracy and efficiency make it an ideal candidate for deployment in telemedicine platforms, where rapid and reliable diagnostics are crucial. It could serve as a support tool for radiologists, helping to reduce the workload and providing a second opinion in busy or understaffed clinical environments. Moreover, due to its computational efficiency, MobileNet-V3 can be integrated into mobile applications, enabling point-of-care diagnostics and thereby expanding access to healthcare services in remote or resource-limited regions. This accessibility could be pivotal in outbreak situations, where swift diagnosis is essential to control the spread of diseases like COVID-19. Additionally, the architecture's adaptability suggests that it could be repurposed for other imaging-based diagnostic tasks, making it a versatile tool in the broader context of healthcare AI solutions. This study, while achieving significant accuracy in detecting pneumonia from chest X-ray images using the MobileNet-V3 architecture, does present certain limitations

that pave the way for future research. One limitation is the reliance on existing, publicly available datasets. While extensive, these datasets may not cover the full spectrum of pneumonia cases, such as varying degrees of disease severity or presentations in diverse populations. Future research could focus on collecting and including more heterogeneous data that encapsulate a wider range of pathological features, particularly from underrepresented regions. Another area is the interpretability of the model's decision-making process. While MobileNet-V3 provides efficiency and accuracy, understanding the 'why' behind its predictions is crucial for clinical acceptance. Future work could explore methods to increase the transparency of the model, possibly by integrating explainable artificial intelligence (XAI) techniques that provide insights into the model's reasoning.

ACKNOWLEDGMENT

REFERENCES

[1] Nascimento-Carvalho, C.M. (2020). Community-acquired pneumonia among children: The latest evidence for an updated management. Jornal de Pediatria, 96: 29-38. https://doi.org/10.1016/j.jped.2019.08.003

[2] Marangu, D., Zar, H.J. (2019). Childhood pneumonia in low-and-middle-income countries: An update. Paediatric Respiratory Reviews, 32: 3-9. https://doi.org/10.1016/j.prrv.2019.06.001

[3] O'Brien, K.L., Baggett, H.C., Brooks, W.A., et al. (2019). Causes of severe pneumonia requiring hospital admission in children without HIV infection from Africa and Asia: the PERCH multi-country case-control study. The Lancet, 394(10200): 757-779. https://doi.org/10.1016/S0140-6736(19)30721-4

[4] Tegenu, K., Geleto, G., Tilahun, D., Bayana, E., Bereke, B. (2022). Severe pneumonia: Treatment outcome and its determinant factors among under-five patients, Jimma, Ethiopia. SAGE Open Medicine, 10. https://doi.org/10.1177/20503121221078445

[5] Mirkarimi, M., Alisamir, M., Saraf, S., Heidari, S., Barouti, S., Mohammadi, S. (2020). Clinical and epidemiological determinants of lower respiratory tract infections in hospitalized pediatric patients. International Journal of Pediatrics, 2020: 1-7. https://doi.org/10.1155/2020/8844420

[6] Kudagammana, S.T., Karunaratne, R.R., Munasinghe, T.S., Kudagammana, H.D.W.S. (2020). Community acquired paediatric pneumonia; experience from a pneumococcal vaccine-naive population. Pneumonia, 12(1): 1-7. https://doi.org/10.1186/s41479-020-00071-6

[7] Ekirapa-Kiracho, E., De Broucker, G., Ssebagereka, A., Mutebi, A., Apolot, R.R., Patenaude, B., Constenla, D. (2021). The economic burden of pneumonia in children under five in Uganda. Vaccine: X, 8: 100095. https://doi.org/10.1016/j.jvacx.2021.100095

[8] Aschenaki, A. (2023). Determinants of pneumonia among under-five aged children attending in Gebretsadik Shawo general hospital, Kaffa Zone, South West Ethiopia (Doctoral dissertation, HU).

[9] McAllister, D.A., Liu, L., Shi, T., et al. (2019). Global, regional, and national estimates of pneumonia morbidity and mortality in children younger than 5 years between 2000 and 2015: A systematic analysis. The Lancet Global Health, 7(1): e47-e57. https://doi.org/10.1016/S2214-109X(18)30408-X

[10] Tasaka, S., Ohshimo, S., Takeuchi, M., et al. (2022). ARDS clinical practice guideline 2021. Journal of intensive care, 10(1): 32. https://doi.org/10.1186/s40560-022-00615-6

[11] World Health Organization. (2022). Stakeholder Consultative Meeting on Prevention and Management of Childhood Pneumonia and Diarrhoea: Report, World Health Organization, 12-14 October 2021.

[12] Jahan, Y. (2022). Management of Severe Childhood Pneumonia in Low-Resource Setting Countries. In Regionalized Management of Medicine, Singapore: Springer Nature Singapore, 177-185. https://doi.org/10.1007/978-981-16-7893-6

[13] Mahomed, N., Manduku, V., Hlabangana, L.T., et al. (2023). Evaluation of paediatric chest X-ray quality in the Pneumonia Aetiology Research for Child Health (PERCH Study), a prospective international multicentre study in 7 developing countries. Wits Journal of Clinical Medicine, 5(2): 85-96. https://doi.org/10.18772/26180197.2023.v5n2a2

[14] Song, J. H., Huh, K., Chung, D.R. (2016). Community-acquired pneumonia in the Asia-Pacific region. In Seminars in Respiratory and Critical care Medicine, Thieme Medical Publishers, 37(6): 839-854. https://doi.org/10.1055/s-0036-1592075

[15] Farooqui, H., Jit, M., Heymann, D.L., Zodpey, S. (2015). Burden of severe pneumonia, pneumococcal pneumonia and pneumonia deaths in Indian states: Modelling based estimates. PloS One, 10(6): e0129191. https://doi.org/10.1371/journal.pone.0129191

[16] Dey, N., Zhang, Y.D., Rajinikanth, V., Pugalenthi, R., Raja, N.S.M. (2021). Customized VGG19 architecture for pneumonia detection in chest X-rays. Pattern Recognition Letters, 143: 67-74. https://doi.org/10.1016/j.patrec.2020.12.010

[17] Jaiswal, A.K., Tiwari, P., Kumar, S., Gupta, D., Khanna, A., Rodrigues, J.J. (2019). Identifying pneumonia in chest X-rays: A deep learning approach. Measurement, 145: 511-518. https://doi.org/10.1016/j.measurement.2019.05.076

[18] Hashmi, M.F., Katiyar, S., Keskar, A.G., Bokde, N.D., Geem, Z.W. (2020). Efficient pneumonia detection in chest Xray images using deep transfer learning. Diagnostics, 10(6): 417. https://doi.org/10.3390/diagnostics10060417

[19] Chakraborty, S., Aich, S., Sim, J.S., Kim, H.C. (2019). Detection of pneumonia from chest X-rays using a convolutional neural network architecture. In International Conference on Future Information & Communication Engineering, 11(1): 98-102.

[20] Nath, M., Choudhury, C. (2020). Automatic detection of pneumonia from chest X-rays using deep learning. In Machine Learning, Image Processing, Network Security and Data Sciences: Second International Conference, MIND 2020, Silchar, India, pp. 175-182. https://doi.org/10.1007/978-981-15-6315-7_14

[21] Rathi, R., Balayan, N., Kumar, C. V. (2020). Pneumonia Detection using Chest X-ray. International Journal of

Pharmaceutical Research (09752366), 12(3): 1150. https://doi.org/10.31838/ijpr/2020.12.03.181

[22] Zhang, D., Ren, F., Li, Y., Na, L., Ma, Y. (2021). Pneumonia detection from chest X-ray images based on convolutional neural network. Electronics, 10(13): 1512. https://doi.org/10.3390/electronics10131512

[23] Barakat, N., Awad, M., Abu-Nabah, B.A. (2023). A machine learning approach on chest X-rays for pediatric pneumonia detection. Digital Health, 9: 205520762311800008. https://doi.org/10.1177/205520762311800

[24] Kundur, N.C., Anil, B.C., Dhulavvagol, P.M., Ganiger, R., Ramadoss, B. (2023). Pneumonia detection in chest X-rays using transfer learning and TPUs. Engineering, Technology & Applied Science Research, 13(5): 11878-11883. https://doi.org/10.48084/etasr.6335

[25] Malik, H., Anees, T., Din, M., Naeem, A. (2023). CDC_Net: Multi-classification convolutional neural network model for detection of COVID-19, pneumothorax, pneumonia, lung Cancer, and tuberculosis using chest X-rays. Multimedia Tools and Applications, 82(9): 13855-13880. https://doi.org/10.1007/s11042-022-13843-7

[26] Athar, A., Asif, R.N., Saleem, M., Munir, S., Al Nasar, M.R., Momani, A.M. (2023). Improving Pneumonia Detection in chest X-rays using Transfer Learning Approach (AlexNet) and Adversarial Training. In 2023 International Conference on Business Analytics for Technology and Security (ICBATS), Dubai, United Arab Emirates, pp. 1-7. https://doi.org/10.1109/ICBATS57792.2023.10111193

[27] Singh, S., Kumar, M., Kumar, A., Verma, B.K., Shitharth, S. (2023). Pneumonia detection with QCSA network on chest X-ray. Scientific Reports, 13(1): 9025. https://doi.org/10.1038/s41598-023-35922-x

[28] Lin, M., Hou, B., Mishra, S., et al. (2023). Enhancing thoracic disease detection using chest X-rays from PubMed Central Open Access. Computers in Biology and Medicine, 159: 106962. https://doi.org/10.1016/j.compbiomed.2023.106962

[29] Bhatt, H., Shah, M. (2023). A convolutional neural network ensemble model for pneumonia detection using chest X-ray images. Healthcare Analytics, 3: 100176. https://doi.org/10.1016/j.health.2023.100176

[30] Kuzinkovas, D., Clement, S. (2023). The detection of COVID-19 in chest X-rays using ensemble CNN techniques. Information, 14(7): 370. https://doi.org/10.3390/info14070370

[31] Sharma, S., Guleria, K. (2023). A deep learning based model for the detection of pneumonia from chest X-ray Images using VGG-16 and neural networks. Procedia Computer Science, 218: 357-366. https://doi.org/10.1016/j.procs.2023.01.018

[32] Anh, P.T., Duc, H.T.M. (2021). A benchmark of deep learning models for multi-leaf diseases for edge devices. In 2021 International Conference on Advanced Technologies for Communications (ATC), Ho Chi Minh City, Vietnam, pp. 318-323. https://doi.org/10.1109/ATC52653.2021.9598196

[33] Pediatric Pneumonia Chest X-ray. (2022). Retrieved 1 October 2022, from https://www.kaggle.com/datasets/andrewmvd/pediatric-pneumonia-chest-xray.

[34] Thukral, B.B. (2015). Problems and preferences in

pediatric imaging. Indian Journal of Radiology and Imaging, 25(4): 359-364. https://doi.org/10.4103/0971-3026.169466

[35] Pine, D.S. (2019). Seeking Clinically useful pediatric imaging measures: Insights from research on autism spectrum disorder. American Journal of Psychiatry, 176(12): 976-978. https://doi.org/10.1176/appi.ajp.2019.19090983

[36] Donnelly, L.F. (2021). Fundamentals of Pediatric Imaging. Academic Press.

[37] Gray, M. (2017). Value based healthcare. BMJ, 356.

[38] European Society of Radiology (ESR) eu-affairs@myesr.org. (2017). ESR concept paper on value-based radiology. Insights into Imaging, 8: 447-454.

[39] Davendralingam, N., Sebire, N.J., Arthurs, O.J., Shelmerdine, S.C. (2021). Artificial intelligence in paediatric radiology: Future opportunities. The British Journal of Radiology, 94(1117): 20200975. https://doi.org/10.1259/bjr.20200975

[40] Pediatric_pnemonia. (2022). https://www.kaggle.com/datasets/mlrocksters/pediatric-pnemonia, accessed on 4 October 2022.

[41] Kermany, D.S., Goldbaum, M., Cai, W., et al. (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. Cell, 172(5): 1122-1131.

[42] Kermany, D., Zhang, K., Goldbaum, M. (2018). Labeled optical coherence tomography (oct) and chest X-ray images for classification. Mendeley Data, 2(2): 651. https://doi.org/10.1016/j.cell.2018.02.010

[43] Pediatric Pneumonia Chest X-ray. (2022). from https://www.kaggle.com/datasets/andrewmvd/pediatric-pneumonia-chest-xray, accessed on 4 October 2022.

[44] Howard, A., Sandler, M., Chu, G., et al. (2019). Searching for mobilenetv3. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Korea (South), pp. 1314-1324. https://doi.org/10.1109/ICCV.2019.00140

[45] Qian, S., Ning, C., Hu, Y. (2021). MobileNetV3 for image classification. In 2021 IEEE 2nd International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE), Nanchang, China, pp. 490-497. https://doi.org/10.1109/ICBAIE52039.2021.9389905

[46] Al Majid, A.H., Wibowo, A., Ramadhanti, S., Zaneta, D., Faoziya, R., Maris, V. (2021). Comparison of models architecture on chest X-ray image classification with transfer learning algorithms. In 2021 5th International Conference on Informatics and Computational Sciences (ICICoS), Semarang, Indonesia, pp. 171-175. https://doi.org/10.1109/ICICoS53627.2021.9651771

[47] Haase, D., Amthor, M. (2020). Rethinking depthwise separable convolutions: How intra-kernel correlations lead to improved mobilenets. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, pp. 14600-14609. https://doi.org/10.1109/CVPR42600.2020.01461

[48] Everything you need to know about MobileNetV3 and its comparison with previous versions. (2019). https://towardsdatascience.com/everything-you-need-to-know-about-mobilenetv3-and-its-comparison-with-previous-versions-a5d5e5a6eeaa, accessed on 2 October 2022.

[49] Zoph, B., Le, Q.V. (2016). Neural architecture search with reinforcement learning. arXiv preprint

arXiv:1611.01578.
https://doi.org/10.48550/arXiv.1611.01578

[50] Yang, T.J., Howard, A., Chen, B., Zhang, X., Go, A., Sandler, M., Sze, V., Adam, H. (2018). Netadapt: Platform-aware neural network adaptation for mobile applications. arXiv:1804.03230. https://doi.org/10.48550/arXiv.1804.03230

[51] Creating Helper Functions. (2022). https://web.cs.wpi.edu/~cs1101/a05/Docs/creating-helpers.html, accessed on 4 October 2022.

[52] Sri, C.G., Bano, S., Deepika, T., Kola, N., Pranathi, Y.L. (2021). Deep neural networks based error level analysis for lossless image compression based forgery detection. In 2021 International Conference on Intelligent Technologies (CONIT), Hubli, India, pp. 1-8. https://doi.org/10.1109/CONIT51480.2021.9498357

[53] Khan, M.A., Kadry, S., Zhang, Y.D., Akram, T., Sharif, M., Rehman, A., Saba, T. (2021). Prediction of COVID-19-pneumonia based on selected deep features and one class kernel extreme learning machine. Computers & Electrical Engineering, 90: 106960. https://doi.org/10.1016/j.compeleceng.2020.106960

[54] Ibrahim, A.U., Ozsoz, M., Serte, S., Al-Turjman, F., Yakoi, P.S. (2021). Pneumonia classification using deep learning from chest X-ray images during COVID-19. Cognitive Computation, 1-13. https://doi.org/10.1007/s12559-020-09787-5

[55] Fukuchi, T., Ogbodo, M.I., Wang, J., Dang, K.N., Ben Abdallah, A. (2022). Efficient pneumonia detection method and implementation in chest X-ray images based on a neuromorphic spiking neural network. In Conference on Computational Collective Intelligence Technologies and Applications, Springer, Cham, pp. 311-321. https://doi.org/10.1007/978-3-031-16014-1_25

[56] Rangasamy, K., Fuzi, N.S.M., As'ari, M.A., Rahmad, N. A., Sufri, N.A.J. (2022). Deep learning-based fine-grained automated pneumonia detection model. Journal of Engineering Science and Technology, 17(4): 2373-2389.

[57] Ranjith, M.S., Aruna Devi, K., Lakshminarayana, M. (2021). Performance analysis of pneumonia using convolutional neural networks. Nveo-Natural Volatiles & Essential Oils Journal, 8(5): 2161-2169.

[58] Nishio, M., Noguchi, S., Matsuo, H., Murakami, T. (2020). Automatic classification between COVID-19 pneumonia, non-COVID-19 pneumonia, and the healthy on chest X-ray image: Combination of data augmentation methods. Scientific Reports, 10(1): 1-6. https://doi.org/10.1038/s41598-020-74539-2

[59] Babic, D., Jovovic, I., Popovic, T., Cakic, S., Filipovic, L. (2022). Detecting pneumonia with tensorflow and convolutional neural networks. In 2022 IEEE International Conference on Omni-layer Intelligent Systems (COINS), Barcelona, Spain, pp. 1-4. https://doi.org/10.1109/COINS54846.2022.9854948

[60] Hou, J., Gao, T. (2021). Explainable DCNN based chest X-ray image analysis and classification for COVID-19 pneumonia detection. Scientific Reports, 11(1): 16071. https://doi.org/10.1038/s41598-021-95680-6