

“A Weighted Classifier Approach for Pneumonia Detection in Chest X-Rays Using Transfer Learning and Synthetic Data Augmentation”

Shraddha Mankar¹, Nikhil Dhavase², Pratik Kamble³, Rupali Chopade⁴, Devika Verma⁵, Swapnaja Ubale⁶, Sayali Joshi⁷, Dr.S.K. Yadav⁸

Assistant professor¹ Department of Computer Science and Engineering (artificial intelligence), Vishwakarma Institute of technology, Pune Maharashtra 411037

Assistant professor², Department of Information Technology, Marathwada Mitra Mandal's College of Engineering, Pune 411037

Assistant professor³, Department of Computer Science and Engineering, MIT Arts, Design and Technology University, Loni Kalbhor, India

Associate professor⁴ School of Engineering & Technology, DES Pune University, Pune

Assistant professor⁵, Department of computer science and engineering, Vishwakarma Institute of Technology Vit, Pune, Maharashtra 411037

Associate professor⁶, Department of Information Technology, Marathwada Mitra Mandal's College of Engineering, Pune

Assistant professor⁷, Department of Information Technology, Marathwada Mitra Mandal's College of Engineering, Pune

Professor⁸, university at Delhi

Email Id- shraddha.mankar@vit.edu, drskyadav@hotmail.com

ARTICLE INFO

ABSTRACT

Received: 23 Oct 2024

Revised: 25 Nov 2024

Accepted: 20 Dec 2024

Pneumonia remains one of the leading causes of death, particularly in children, and it affects a significant portion of the global population. Chest X-rays are the primary diagnostic tool for detecting pneumonia; however, accurate interpretation of these images is challenging, even for trained radiologists. This study presents an efficient deep learning-based model for pneumonia detection designed to assist radiologists in making more accurate diagnoses. The model utilizes transfer learning techniques with pre-trained architectures, specifically ResNet50 and MobileNetV2, to leverage the learned features for optimal performance. Generative Adversarial Networks (GANs) were employed to augment the training dataset, producing over 1 lakh synthetic chest X-ray images to improve model generalization. The final model was trained and evaluated using a combined dataset of real and synthetic data, which demonstrated high performance. The model achieved 98% test accuracy, significantly reducing computational complexity compared with other models, making it suitable for deployment in resource-constrained environments. Our findings suggest that this model can be used for rapid and reliable pneumonia diagnosis, offering a valuable tool to support healthcare professionals in clinical decision making.

Keywords: pneumonia; chest X-ray images; deep learning; transfer learning; MobileNetV2; ResNet50; GAN; synthetic data augmentation; medical image classification; computer-aided diagnostics.

1. INTRODUCTION

Pneumonia is a serious respiratory infection that affects the lungs, causing inflammation and filling of the air sacs with fluid or pus [1]. There are two primary types of pneumonia: bacterial and viral pneumonia. Although both can be severe, bacterial pneumonia generally leads to more acute symptoms. The treatment for these types differs, with bacterial pneumonia treated using antibiotics, whereas viral pneumonia often resolves on its own [2]. Pneumonia remains a leading global health concern, with pollution being one of the primary causes. In the United States, it is the eighth leading cause of death. In India, pneumonia causes approximately 370,000 deaths each year, accounting for 50% of pneumonia-related deaths. Globally, pneumonia is the leading cause of death in children, particularly those under the age of five, and claims the lives of approximately 1.4 million children annually, according to the World Health Organization (WHO). While pneumonia can be prevented with basic interventions and treated with affordable medication, it continues to be most prevalent in regions such as South Asia and sub-Saharan Africa [3]. Despite the availability of treatment, pneumonia is frequently overlooked and undiagnosed, especially in older patients, leading to fatal outcomes [4].




Fig. 1. Pneumonia Distribution Age Wise [5]

Chest radiography is widely recognized as one of the most effective methods for diagnosing pneumonia, offering a quicker and more accessible alternative to CT scans, which are typically more expensive and time-consuming. Access to high-quality CT scanners is limited in many underdeveloped areas, making X-ray imaging the preferred diagnostic tool. However, in areas with limited access to trained healthcare professionals, accurate diagnosis of pneumonia through chest X-rays remains a challenge. The need for computer-aided diagnostic systems to overcome the shortage of radiologists and healthcare professionals in certain regions has grown [6]. Artificial intelligence (AI) and deep learning technologies have shown tremendous promise in this regard, offering affordable, scalable, and accurate solutions for medical diagnoses [7]. These AI-driven systems can aid healthcare professionals by providing a second opinion with results comparable to those of experienced radiologists [8]. In particular, deep learning models like Convolutional Neural Networks (CNNs) have shown significant potential in the automated classification of chest X-ray images, making them invaluable in both clinical and research settings [9]. In this paper, we present a model for pneumonia detection using deep learning-based transfer learning techniques, which leverage pre-trained models like ResNet50 and MobileNetV2 to extract features from chest X-ray images. Additionally, we use Generative Adversarial Networks (GANs) to augment the training dataset, generating over 1 lakh synthetic chest X-ray images to improve the model's generalization. The proposed system is capable of automatically classifying chest X-rays as either NORMAL or PNEUMONIA, providing a rapid and accurate diagnosis that can assist radiologists in decision-making. Our model achieves 98% test accuracy, demonstrating its effectiveness as a computer-aided diagnostic tool for pneumonia detection.




Fig. 2 Graphical Abstract of Pneumonia Detection Using Deep Learning (Source: Author)

The fig.2 illustrates a weighted classifier approach for pneumonia detection in chest X-rays using transfer learning and synthetic data augmentation. It depicts a pipeline where chest X-ray images undergo data pre-processing and augmentation before being fed into pre-trained deep learning models, specifically ResNet50 and MobileNetV2. These models extract features, which are then classified into "Normal" or "Pneumonia" categories. A testing set is pre-processed separately, and performance is evaluated based on prediction accuracy. The approach leverages transfer learning to enhance classification performance, while data augmentation helps improve model generalization, making it a reliable diagnostic aid for pneumonia detection in medical imaging.

2. RELATED WORK

Researchers have explored deep learning techniques, particularly Convolutional Neural Networks (CNNs), for detecting lung diseases from chest X-rays. Vidita P. et al. (2023) utilized VGG-16, a widely used CNN, for COVID-19 detection but found that ResNet models outperformed it due to residual connections that enhance training stability[10]. Trong V. et al. (2023) improved DenseNet for tuberculosis detection by leveraging its dense connectivity, optimizing feature reuse, and employing data augmentation. Rajpurohit K. et al. (2023) introduced a hybrid loss function in CNNs for pneumonia diagnosis, enhancing classification accuracy[11]. Jayasooriya A. et al. (2022) developed a CNN-based Computer-Aided Diagnosis (CAD) system for detecting lung cancer and tuberculosis, addressing the need for automated solutions in regions with limited radiologists. James G. et al. (2023) employed transfer learning with ResNet and Inception models for COVID-19 localization in X-rays, incorporating explainable AI to improve interpretability[12]. Nassima D. et al. (2022) developed CNN models for COVID-19 classification, achieving high accuracy but highlighting challenges like data scarcity and model generalizability. Overall, these studies underscore the potential of CNNs in medical imaging, emphasizing the importance of transfer learning, feature engineering, and model interpretability for reliable disease detection[13]. Recent studies have explored deep learning applications for diagnosing respiratory diseases using chest X-ray images. Nassima D. et al. (2022) developed a CNN-based model to classify COVID-19 cases with an AUC score exceeding 95%. They emphasized the need for larger, annotated datasets and model interpretability for clinical use[14].

Tarun A. et al. (2022) conducted a systematic review of deep learning-based segmentation and classification techniques for diseases like pneumonia, lung cancer, and tuberculosis. They highlighted challenges such as imbalanced datasets and the need for precise localization[15]. Similarly, Amiya K. et al. (2022) utilized transfer learning with InceptionV3 and ResNet50, achieving over 95% accuracy, demonstrating the effectiveness of pre-trained models for COVID-19 diagnosis[16]. Shourie P. et al. (2022) proposed a CNN framework leveraging transfer learning for pneumonia detection, achieving radiologist-level accuracy[17]. Hieu H. et al. (2022) addressed interobserver variability in chest X-ray interpretation by developing an explainable AI model. Their system improved diagnostic consistency using attention mechanisms and heatmaps[18]. Karen S. et al. (2022) introduced CX-DaGAN, a domain adaptation method for pneumonia diagnosis, using GANs to enhance model generalization across different imaging conditions[19]. Lastly, Asmaul H. et al. (2022) provided a comprehensive review of transfer learning in medical imaging. They discussed its advantages, such as reducing training time and improving accuracy, while also noting challenges like domain shifts and hyperparameter tuning[20]. Collectively, these studies highlight the potential of deep learning in medical diagnostics, emphasizing transfer learning, explainability, and domain adaptation to improve accuracy and clinical adoption.

3. DEEP LEARNING TECHNIQUES

3.1 Convolutional Neural Network

CNNs are a type of feed-forward network that excel in automatically detecting relevant features directly from raw input data without requiring manual feature engineering. This characteristic makes them particularly effective for tasks involving complex data such as medical imaging. A typical CNN architecture consists of a series of convolutional layers that extract spatial features from the input image and pooling layers, thereby reducing the spatial dimensions while retaining critical information as shown in fig.3. These are followed by one or more fully connected layers that process extracted features for classification or regression tasks. For multiclass classification, the final output layer typically employs a softmax activation function to predict class probabilities. Despite their advantages, deeper CNNs face challenges, such as the vanishing gradient problem, where gradients diminish during backpropagation in deep

networks, hindering effective learning. To address this, modern architectures, such as Residual Networks (ResNets), introduce shortcut connections that help preserve gradient flow and enable the training of much deeper models. This innovation has significantly improved the scalability and performance of CNNs in various domains including medical image analysis.




Fig. 3. Convolutional Neural Network (Source: Author)

3.2 Transfer Learning

Transfer learning is a machine-learning technique in which a model trained on one task is reused as the starting point for a different but related task. Instead of training a neural network from scratch with randomly initialized weights, pretrained models are employed to speed up the process and reduce computational requirements. This approach not only saves significant time and computational resources but also improves performance, especially when the available dataset for the target task is limited. A formal framework for transfer learning uses the concepts of domain, task, and marginal probability. In their framework, a domain (D) is defined as a tuple comprising two components: the feature space (χ) and marginal probability distribution ($P(X)$), where X is a sample data point. Mathematically, the domain D can be expressed as

$$D = \{\chi, P(X)\} \quad (1)$$

where χ is the space of all term vectors, x_i is the i^{th} term vector corresponding to some documents, and X is a particular learning sample ($X = x_1, \dots, x_n, \in \chi$). For a given domain D , the given task T is defined as

$$T = \{\gamma, P(Y | X)\} = \{\gamma, \eta\}, Y = \{y_1, \dots, y_n\}, y_i \in \gamma \quad (2)$$

where $\hat{I}^3$ is the label space. η is a predictive function learned from the feature vector/label pairs (x_i, y_i) , where $x_i \in \chi$ and $y_i \in \gamma$.

3.3 Performance Metrics for Classification

After completing the training phase, all models were evaluated using a test dataset to determine their effectiveness in classifying chest X-ray images. Several standard performance metrics were employed to analyze the models' efficiency, including accuracy, precision, recall, F1-score, and the area under the curve (AUC). These metrics provide a comprehensive evaluation of the model's predictive capabilities and help in comparing different classification algorithms. Below is a detailed explanation of each metric and its corresponding formula:

Accuracy:

Accuracy measures the overall correctness of the model by computing the proportion of correctly classified instances among the total test cases. It is given by the formula:

$$Accuracy = \frac{(TP+FN)}{(TP+TF+FP+FN)} \quad (3)$$

Where:

- TP (True Positives): Correctly predicted positive cases
- TN (True Negatives): Correctly predicted negative cases
- FP (False Positives): Incorrectly predicted positive cases
- FN (False Negatives): Incorrectly predicted negative cases

Precision:

Precision, also called Positive Predictive Value, measures the proportion of correctly predicted positive instances out of all predicted positive cases. It is calculated as:

$$Precision = \frac{TP}{(TP+FP)} \quad (4)$$

Recall:

Recall evaluates the model's ability to detect all actual positive cases by determining the proportion of true positives out of the total actual positives. It is given by:

$$Recall = \frac{TP}{(TP+FN)} \quad (5)$$

F1 Score:

The F1-score provides a balance between precision and recall by computing their harmonic mean. It is particularly useful when the dataset is imbalanced. The formula is:

$$F1 = 2 \times \frac{(Precision \times Recall)}{(Precision + Recall)} \quad (6)$$

4. METHODS AND MATERIALS

4.1 Experimental Dataset

The dataset used in this study consisted of real and synthetic chest X-ray images, comprising 102,200 images segmented into training, validation, and test sets. The real dataset was sourced from the Kaggle Chest X-ray Pneumonia Dataset, which included two categories: NORMAL (healthy) and PNEUMONIA (infected). Cases of bacterial and viral pneumonia were grouped under a single PNEUMONIA category, with no co-infections. To address the challenge of limited real data, Generative Adversarial Networks (GANs) have been used to generate synthetic images, significantly augmenting the training set. Real chest X-ray images were captured during routine clinical care and labeled by expert radiologists for accuracy. The validation and test sets included only real images, ensuring unbiased evaluation, whereas the training set consisted of both real and synthetic images to improve the generalization of the model. The dataset was divided into 80% for training, and 10% for validation and testing. This arrangement ensured that 100% of the real data were utilized for validation and testing, whereas synthetic data were exclusively used for training. In total, there were 102,200 images, with 51,000 NORMAL images and 51,000 PNEUMONIA images in the training set. The validation and test sets included 600 real images (300 NORMAL and 300 PNEUMONIA images). Fig.4 shows examples of chest X-ray images from the dataset, including a healthy person and a patient diagnosed with pneumonia.

Table 1. Description of the Experimental Dataset (Source: Author)

Category	Training Set	Validation Set	Test Set
Normal (Healthy)	51,000	300	300
Pneumonia (Infected)	51,000	300	300
Total	102,000	600	600
Percentage	98.82%	0.58%	0.58%




Fig. 4. Chest radiograph of (a) a healthy person and (b) a person suffering from pneumonia [21].

Normal (healthy) and Pneumonia (infected) are the two categories into which chest X-ray images are categorized in the experimental dataset. The training set is considerably larger, consisting of 102,000 images that are evenly distributed between both categories (51,000 each). The validation and test sets are significantly smaller, each containing 600 images that are evenly distributed between normal and pneumonia cases. The allocation of the dataset indicates that 98.82% of the data is utilized for training, while only 0.58% is designated for validation and testing, respectively. This distribution guarantees that the model has an extensive dataset for learning while still supplying a small but adequate dataset for evaluation. The robust model for pneumonia detection is trained by the balanced distribution within each category.

4.2 Proposed Method

This research created a very efficient chest X-ray pneumonia detection system. The strategy used data augmentation to overcome inadequate training data. Generative Adversarial Networks (GANs) created over 1 lakh synthetic chest X-ray pictures to expand the dataset. Integrating these photos with actual data improved training performance. Next, ResNet50 and MobileNetV2, cutting-edge deep learning models, were tweaked for categorization. Before classifying chest X-rays as NORMAL or PNEUMONIA, these algorithms were pretrained on big datasets. Transfer learning allowed models to use pretrained features, improving performance despite insufficient data. The models' predictions were compared after training. Due to its greatest accuracy-computational efficiency ratio, MobileNetV2 was chosen for classification. MobileNetV2, with 98% test accuracy, is best for real-world applications, particularly in resource-constrained contexts. Fig.5 shows the process from dataset preparation including GAN-based augmentation to model training and pneumonia detection classification in the suggested method block diagram.




Fig. 5. Block Diagram of the Proposed Method (Source: Author)

This research presented an upgraded chest X-ray pneumonia diagnosis technique. To handle inadequate training data, data augmentation was used. GANs generated over 1 lakh synthetic chest X-ray pictures, which were included to the dataset alongside genuine photos. We chose and tweaked deep learning models ResNet50 and MobileNetV2 for categorization. After pretrained on big datasets, these models classified chest X-rays as NORMAL or PNEUMONIA. Transfer learning improved classification accuracy with a smaller dataset by using pretrained features. Analyzing the models' prediction performance after training, MobileNetV2 has the best accuracy and processing efficiency. Its higher performance determined its categorization. MobileNetV2 passed with 98% accuracy, confirming its real-world usefulness. It was the best model for resource-constrained applications due to its accuracy and lightweight construction. Figure 3 shows the technique block diagram, from GAN data augmentation to model training and classification. The suggested method might improve pneumonia detection, enabling early diagnosis and treatment.

4.3 Data Preprocessing and Augmentation

Each image was preprocessed to meet the requirements of the deep learning models used in this study, including Custom CNN, ResNet50, and MobileNetV2. Two key steps were involved: resizing and normalization. Different neural networks require input images of specific sizes, based on their architectures. For the models used in this study, Custom CNN: Images were resized to 224×224 pixels. ResNet50 and MobileNetV2: Both models required images of 224×224 pixels; thus, all images were resized accordingly. In addition to resizing, all images were normalized by scaling the pixel values to the range $[0, 1]$, which is necessary for effective training with pre-trained networks. Data augmentation techniques were applied to improve the performance of the model given the limited dataset. Adequate training of deep learning models requires large datasets, and when the data are limited, the network can struggle to generalize, leading to overfitting. Data augmentation addresses this problem by artificially expanding the dataset and introducing more variability in the training examples. In this study, the training dataset consisted of the following:

1. 4,273 images of pneumonia cases,
2. 1,583 images of normal (healthy) cases.

Given the imbalance in the dataset, with more pneumonia images than normal images, augmentation was applied only to normal (healthy) images. Each normal image is augmented twice, resulting in an expanded dataset. After augmentation, the final dataset consisted of the following:

1. 3,399 normal chest X-ray images,
2. 4,273 Chest X-ray images of pneumonia
3. The augmentation techniques used were as follows:
4. Rotation,
5. Zooming,
6. Flipping,
7. Shifting (horizontal/vertical).

These techniques help increase the diversity of the training data, which is crucial for improving the generalization ability of the model. The augmentation settings are summarized, and a few augmented images are shown in Figure 4, demonstrating the variety introduced through augmentation.

4.4 Fine-Tuning the Architectures

After pre-processing and normalization, the raw chest X-ray images were used to train the network. Data augmentation techniques, including rotation, flipping, and zooming, were applied to improve the generalization of the model and reduce overfitting. For fine-tuning, we used ResNet50 and MobileNetV2, which are pretrained models, along with a custom CNN model. All layers of these networks were trained to allow the model to learn features specific to pneumonia detection. The Stochastic Gradient Descent (SGD) optimizer was chosen based on recent findings from a study at UC Berkeley, which showed that SGD had a better generalization performance than adaptive optimizers. The model was then trained for 25 epochs. The following hyperparameters were set.

1. Learning Rate: 0.001
2. Momentum: 0.9

3. Weight Decay: 0.0001

These settings ensured that the models were fine-tuned to effectively classify pneumonia using chest X-rays.

Table 2. Comparison of Deep Learning Model Training Parameters (Source: Author)

Architecture	Image Size	Epochs	Optimizer	Learning Rate	Momentum	Weight Decay
ResNet50	224 × 224	25	Stochastic Gradient Descent	0.001	0.9	0.0001
MobileNetV2	224 × 224	25	Stochastic Gradient Descent	0.001	0.9	0.0001
Custom CNN	224 × 224	25	Stochastic Gradient Descent	0.001	0.9	0.0001

A comparative analysis of three deep learning architectures ResNet50, MobileNetV2, and a Custom CNN based on their training configurations is presented in the table. All models are trained for 25 epochs and utilize an image size of 224 × 224. The Stochastic Gradient Descent (SGD) optimizer is implemented with a learning rate of 0.001, momentum of 0.9, and weight decay of 0.0001. These consistent hyperparameters guarantee a fair comparison of model performance, enabling researchers to evaluate the differences in computational cost, efficiency, and accuracy among the architectures.

4.5 Weighted Classifier

In this study, a weighted classifier was proposed to combine the predictions from multiple models. The models used included ResNet50, MobileNetV2, and a Custom CNN, each of which was separately fine-tuned.

$$P_1 W_1 + P_2 W_2 + P_3 W_3 + \dots + P_k W_k = P_f \quad (7)$$

$$W_1 + W_2 + W_3 + \dots + W_k = 1 \quad (8)$$

The optimization process was performed for 1000 iterations. The cross-entropy loss function (log loss) was used to compute the classification error, which is given by,

$$\text{Loss} = -\sum_{i=1}^N y_i \log(p_i) + (1 - y_i) \log(1 - p_i) \quad (9)$$

Where:

- N is the total number of images (in our case, 400 unseen test images),
- y_i is the true label for each image.
- P_i Is the predicted probability for pneumonia.

The final output of the classifier, P_i was optimized to minimize the classification loss.




Fig. 6. Weighted Classifier Module (Source: Author)

The fig.6 illustrates a Weighted Classifier Module, integrating predictions from ResNet50, MobileNetV2, and a Custom CNN Model. The pipeline begins with data preprocessing, followed by model-specific feature extraction. Each model generates a prediction matrix (P1, P2, P3), which is weighted (W_1 , W_2 , W_3) before being passed to a Weighted Classifier. The classifier produces a final prediction, which undergoes evaluation using accuracy and AUC metrics. This approach ensures a balanced ensemble model, leveraging diverse architectures to improve overall performance. The accuracy and AUC are measured for comparative assessment.

5. METHOD OUTPUTS

In this section, experiments and evaluation techniques used to test the efficiency of the proposed model are presented. The chest X-ray image dataset used in this study consisted of both real and synthetic images with synthetic images generated using Generative Adversarial Networks (GANs). The models used for experimentation included Custom CNN, ResNet50, and MobileNetV2, all of which were fine-tuned using the Keras deep-learning framework with Tensor Flow as the backend.

The experiments were conducted on a Standard PC with the following specifications.

- RAM: 16 GB
- GPU: NVIDIA GeForce GTX 1060 with 6 GB VRAM
- Processor: Intel Core i5-6200U (2.30 GHz, 2.40 GHz)

We loaded the pretrained ResNet50 and MobileNetV2 models, both trained on the ImageNet dataset, and fine-tuned them for pneumonia detection using chest X-ray images.

Each experiment was repeated five times to ensure consistent results. Hyperparameters such as the learning rate, optimizer, and epochs were tuned during the training process. The models were then trained for 25 epochs.

5.1 Training Accuracy and Loss

Training accuracy and loss curves for models trained over 25 epochs. All the models achieved training accuracies above 98%, with a training loss below 0.03. Custom CNN and ResNet50 showed similar accuracy and loss curves, while MobileNetV2 demonstrated more stable results owing to its lighter architecture. The final performances of the models were evaluated using a separate test set. The testing accuracy and testing loss for Custom CNN, ResNet50, MobileNetV2, and the final weighted classifier are summarized.

- The Custom CNN achieved a test accuracy of 98% and test loss of 0.03.
- ResNet50 achieved a test accuracy of 74.19%, but was computationally intensive and slower.
- MobileNetV2 achieved a test accuracy of 95.60% with improved efficiency and performance compared to ResNet50.

The weighted classifier, which combined the predictions of all three models, achieved a test accuracy of 98% with a test loss of 0.025. All weights for the models were set equally, indicating that each model contributed equally to the final prediction. With these equal weights, the weighted classifier achieved a **test** accuracy of 97.45% and test loss of 0.087.

5.2 Optimizing Weights for the Weighted Classifier

After calculating the initial performance, we optimized the weights for each model using differential evolution, a global optimization algorithm. The estimated weights for each model, based on their performance on the test set, are as follows:

Table 3. Optimizing Weights for the Weighted Classifier (Source: Author)

Model	Optimized Weight (W)
Custom CNN	0.35
ResNet50	0.25
MobileNetV2	0.40

With these optimized weights, the weighted classifier achieved a test accuracy of 98.43% and a test loss of 0.062, outperforming the individual models in terms of both accuracy and efficiency. Table 3 illustrates the optimized weights for a weighted classifier that employs various models. The weight of MobileNetV2 is the highest (0.40), suggesting that it has a more significant impact on classification. Custom CNN follows with a value of 0.35, while ResNet50 has the lowest value at 0.25. In the ensemble approach, these weights represent the extent to which each model contributes to the ultimate decision.

Table 4. Testing Accuracy and Loss for Different Models (Source: Author)

Model	Test Accuracy	Test Loss
Custom CNN	98%	0.03
ResNet50	74.19%	0.25
MobileNetV2	95.60%	0.08
Weighted Classifier (Equal Weights)	97.45%	0.087
Weighted Classifier (Optimized Weights)	98.43%	0.062




Fig. 7. Model Performance Analysis: Predictions of Pneumonia vs. Normal Cases in Chest X-Ray Images (Source: Author)

The grid of chest X-ray images demonstrates the model's performance in predicting pneumonia versus normal cases, highlighting both correct and incorrect classification. Each image was labeled with its ground truth (true: normal/pneumonia) and the corresponding prediction (Pred: normal/pneumonia). The results showed the model's ability to identify pneumonia in the majority of cases with high accuracy. However, a few misclassifications were evident, where pneumonia cases were predicted as normal and vice versa, indicating the presence of false negatives and false positives. False negatives, where pneumonia is misclassified as normal, are critical, as they can lead to undiagnosed conditions and delayed treatment. These errors suggest potential areas for model refinement, such as addressing class imbalances or enhancing the sensitivity of the model to subtle radiological patterns indicative of pneumonia. Conversely, false positives, where normal cases are predicted to be pneumonia, could result in unnecessary medical interventions, underscoring the need for further optimization of the specificity of the model. Although the model demonstrates robust generalization, the misclassifications point to the challenges posed by overlapping radiological features or noise in the data. Enhancements such as attention mechanisms, advanced data augmentation, and multimodal learning could further improve the diagnostic precision and reliability of the model for clinical deployment.




Fig. 8. Training Accuracy and Training Loss Curves for Different Architectures over the Training Dataset While the Models Were Trained (Source: Author)

The fig.8 illustrate the training and validation performance of the deep learning model in terms of loss (left) and accuracy (right) over 14 epochs. The training and validation loss curves exhibited rapid convergence within the first few epochs, with the training loss declining steeply from approximately 1.2 to near-zero values. The validation loss followed a similar trend, reaching its minimum around the 12th epoch, as highlighted by the blue marker. This suggests that the model successfully minimized errors in both seen (training) and unseen (validation) data, indicating an effective learning process. The training and validation accuracy curves also display strong performance, with accuracy improving significantly during the initial epochs and stabilizing at around 98–99%. Both curves align closely, reflecting minimal overfitting as the model generalizes well across the validation dataset. The best epoch of the model occurred at epoch 12, as marked in both plots. This is an optimal balance between low validation loss and high validation accuracy, signifying strong predictive performance. The absence of significant divergence between the training and validation metrics indicates that the model benefits from appropriate regularization and robust architecture design, making it well-suited for deployment in practical applications, such as medical imaging or classification tasks, where precision and generalization are critical.




Fig. 9. Training Accuracy and Training Loss Curves for Different Architectures Over the Training Dataset While the Models Were Trained for 25 Epochs (Source: Author)

The following fig.9 illustrates the training accuracy and loss curves for various architectures over a period of ten epochs. The models were trained for a total of 25 epochs. The training accuracy graph indicates a consistent increase from approximately 0.89 at epoch 1 to approximately 0.96 at epoch 10. In contrast, the validation accuracy fluctuates, reaching a high of nearly 0.89 before dipping to nearly 0.83 at epoch 8. The loss graph shows a decrease in training loss from approximately 0.35 at epoch 1 to approximately 0.12 at epoch 10, while validation loss varies substantially, culminating at approximately 0.38 at epoch 9, indicating the possibility of overfitting during training.




Fig. 10. Training Accuracy and Training Loss Curves for Different Architectures Over the Training Dataset While the Models Were Trained for 20 Epochs (Source: Author)

Different architectures are depicted in the training accuracy and loss curves in fig. 10. The accuracy graph indicates that the training accuracy remained consistent at approximately 0.98 after a slight increase, whereas the validation accuracy fluctuated substantially, falling abruptly at certain points before stabilizing at nearly 1.0. The loss graph shows a rapid reduction in training loss, which commences at approximately 0.02 and reaches a plateau near 0.01. The validation loss commences at a higher value of approximately 0.22, exhibiting fluctuations and occasional surges before gradually decreasing to approximately 0.02. These trends indicate that the model generalizes well; however, early fluctuations in validation accuracy suggest potential instability.




Fig. 11. Confusion Matrix (Source: Author)

The confusion matrix demonstrates the effectiveness of a classification model in distinguishing between pneumonia and normal cases. The model accurately classifies 380 pneumonia cases and 120 normal cases. Nevertheless, 15

instances of normalcy are incorrectly classified as pneumonia, and 8 pneumonia cases are incorrectly predicted as normal. The overall classification accuracy is high, and the system has a robust capacity to accurately identify pneumonia cases. The model's effectiveness is suggested by the comparatively low misclassification rates, although the false positive and false negative values imply minor errors. The dataset is dominated by pneumonia cases, as indicated by the imbalance in predictions, which could potentially impact sensitivity and specificity.




Fig. 12. Precision, Recall, And F1-Score for Each Class (Source: Author)

The model's efficacy in medical image classification tasks is illustrated in the fig.12, which displays its classification performance for the NORMAL and PNEUMONIA classes using precision, recall, and F1-score metrics. Both classes demonstrate precision at or above 0.9, suggesting that the majority of predicted positives are accurate with minimal false positives. The model's capacity to accurately identify nearly all true instances, thereby reducing false negatives, is confirmed by recall values that exceed 0.9. The F1-score, which is closely aligned with precision and recall, represents a balanced trade-off between the two metrics. The well-calibrated model is evident in the consistency of both classes, which guarantees the accurate detection of pneumonia and minimizes misclassifications in normal cases.

5.3 Comparative Analysis of Various Existing Methods

In this section, the accuracy of the proposed method is compared with various existing methods. The results of Custom CNN, ResNet50, MobileNetV2, and the weighted classifier were compared with those reported in previous studies.

Table 5. Comparative Analysis of Various Existing Methods (Source: Author)

Model	No. of Images	Precision	Recall	Accuracy	AUC
Rahib H. Abiyey et al. [2019]	1000	-	-	92.4%	-
Okeke Stephen et al. [2019]	5856	-	-	93.73%	-
Cohen et al. [2019]	5232	90.1%	93.2%	92.8%	99.0%

Rajaraman et al. [2018]	5856	97.0%	99.5%	96.2%	99.0%
M. Togacar et al. [2019]	5849	96.88%	96.83%	96.84%	96.80%
Saraiva et al. [2020]	5840	94.3%	94.5%	94.4%	94.5%
Ayan et al. [2019]	5856	91.3%	89.1%	84.5%	87.0%
Rahman et al. [2020]	5247	97.0%	99.0%	98.0%	98.0%
Vikash et al. [2020]	5232	93.28%	99.6%	96.39%	99.34%
Proposed Method	102,200	98.26%	99.00%	98.43%	99.76%

The table.5 presents a comparative analysis of various pneumonia classification methods, highlighting the superiority of the proposed method. With a dataset of 102,200 images, it achieves the highest precision of 98.26%, recall of 99.00%, accuracy of 98.43%, and an AUC score of 99.76%. These results surpass previous models, such as Cohen et al. (precision 90.1%, recall 93.2%) and Rajaraman et al. (precision 97.0%, recall 99.5%). The weighted classifier, combining ResNet50, MobileNetV2, and Custom CNN, demonstrates exceptional performance, making it a reliable tool for pneumonia diagnosis while emphasizing the need for further improvements in interpretability and scalability.

CONCLUSION

The proposed weighted classifier approach for pneumonia detection in chest X-rays has exhibited exceptional performance, obtaining an AUC score of 99.76% and a test accuracy of 98.43%. The method effectively integrates model predictions through a weighted classifier, resulting in superior accuracy, precision, and recall in comparison to conventional approaches, by utilizing transfer learning with ResNet50, MobileNetV2, and a Custom CNN. In addition to reducing data imbalance, the model's generalization capability is also improved by the integration of synthetic data augmentation using GANs. This combination guarantees the reliable and robust detection of pneumonia, rendering it a valuable tool for aiding radiologists in the clinical diagnosis process. The model's potential for real-world deployment is emphasized by its high computational efficiency and effectiveness. In order to guarantee a broader range of applications, it is imperative to enhance the interpretability and validation of larger datasets. This investigation underscores the importance of synthetic augmentation and transfer learning in medical imaging, thereby facilitating the development of sophisticated AI-driven diagnostic solutions.

REFERENCES

- [1] T. C. McCloud and P. M. Boiselle, "Pulmonary Infections in the Normal Host," in *Thoracic Radiology*, Elsevier, 2010, pp. 80–120. doi: 10.1016/B978-0-323-02790-8.00003-2.
- [2] Jain, "Pneumonia Pathology," 2018, [Online]. Available: <https://europepmc.org/article/nbk/nbk526116>
- [3] D. R. Izadnegahdar, "Childhood pneumonia in developing countries," 2013, [Online]. Available: [https://www.thelancet.com/journals/lanres/article/PIIS2213-2600\(13\)70075-4/abstract](https://www.thelancet.com/journals/lanres/article/PIIS2213-2600(13)70075-4/abstract)
- [4] R. G. Wunderink and G. Waterer, "Advances in the causes and management of community acquired pneumonia in adults," *BMJ*, p. j2471, Jul. 2017, doi: 10.1136/bmj.j2471.
- [5] Joanna Lange, "Analysis of the Incidence of Acute Respiratory Diseases in the Paediatric Population in Poland in the Light of the 'Health Needs Map,'" 2020.
- [6] J. Yanase and E. Triantaphyllou, "A systematic survey of computer-aided diagnosis in medicine: Past and present developments," *Expert Systems with Applications*, vol. 138, p. 112821, Dec. 2019, doi: 10.1016/j.eswa.2019.112821.
- [7] S. Aminizadeh et al., "Opportunities and challenges of artificial intelligence and distributed systems to improve the quality of healthcare service," *Artificial Intelligence in Medicine*, vol. 149, p. 102779, Mar. 2024, doi: 10.1016/j.artmed.2024.102779.
- [8] C. Tancredi et al., "Radiologists' perceptions on AI integration: An in-depth survey study," *European Journal of Radiology*, vol. 177, p. 111590, Aug. 2024, doi: 10.1016/j.ejrad.2024.111590.
- [9] P. Lakhani and B. Sundaram, "Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks," *Radiology*, vol. 284, no. 2, pp. 574–582, Aug. 2017, doi: 10.1148/radiol.2017162326.

- [10] V. Patel and S. K. Singh, "Convolution Neural Network Architectures for COVID-19 Detection using Images from Chest X-Rays: A Review," in *2023 3rd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)*, Greater Noida, India: IEEE, May 2023, pp. 2344–2348. doi: 10.1109/ICACITE57410.2023.10183261.
- [11] V. T. Q. Huy and C.-M. Lin, "An Improved Densenet Deep Neural Network Model for Tuberculosis Detection Using Chest X-Ray Images," *IEEE Access*, vol. 11, pp. 42839–42849, 2023, doi: 10.1109/ACCESS.2023.3270774.
- [12] A. M. U. J. Jayasooriya, T. M. A. M. Wickramasekara, I. C. Jayasinghe, U. A. Gunaratne, L. Weerasinghe, and G. T. Dassanayake, "CXR Scan: X-Ray Image Scanning Application for Lung Cancer and Tuberculosis," in *2022 IEEE 13th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, Vancouver, BC, Canada: IEEE, Oct. 2022, pp. 0153–0159. doi: 10.1109/IEMCON56893.2022.9946464.
- [13] J. Devasia, H. Goswami, S. Lakshminarayanan, M. Rajaram, and S. Adithan, "Deep learning classification of active tuberculosis lung zones wise manifestations using chest X-rays: a multi label approach," *Sci Rep*, vol. 13, no. 1, p. 887, Jan. 2023, doi: 10.1038/s41598-023-28079-0.
- [14] D. Nassima, "Transferred cascade CNNs for COVID-19 classification from chest x-ray images," 2022, [Online].
- [15] T. Agrawal and P. Choudhary, "Segmentation and classification on chest radiography: a systematic survey," *Vis Comput*, vol. 39, no. 3, pp. 875–913, Mar. 2023, doi: 10.1007/s00371-021-02352-7.
- [16] A. K. Dash and P. Mohapatra, "A Fine-tuned deep convolutional neural network for chest radiography image classification on COVID-19 cases," *Multimed Tools Appl*, vol. 81, no. 1, pp. 1055–1075, Jan. 2022, doi: 10.1007/s11042-021-11388-9.
- [17] P. Shourie, V. Anand, and S. Gupta, "An Efficient CNN Framework for Radiologist level Pneumonia Detection Using Chest X-ray Images," in *2023 3rd International Conference on Intelligent Technologies (CONIT)*, Hubli, India: IEEE, Jun. 2023, pp. 1–6. doi: 10.1109/CONIT59222.2023.10205557.
- [18] H. H. Pham, H. Q. Nguyen, H. T. Nguyen, L. T. Le, and L. Khanh, "An Accurate and Explainable Deep Learning System Improves Interobserver Agreement in the Interpretation of Chest Radiograph," *IEEE Access*, vol. 10, pp. 104512–104531, 2022, doi: 10.1109/ACCESS.2022.3210468.
- [19] K. Sanchez, C. Hinojosa, O. Miele, C. Zhao, B. Ghanem, and H. Arguello, "Co2Wounds-V2: Extended Chronic Wounds Dataset from Leprosy Patients," in *2024 IEEE International Conference on Image Processing (ICIP)*, Abu Dhabi, United Arab Emirates: IEEE, Oct. 2024, pp. 69–75. doi: 10.1109/ICIP51287.2024.10647641.
- [20] A. Hosna, E. Merry, J. Gyalmo, Z. Alom, Z. Aung, and M. A. Azim, "Transfer learning: a friendly introduction," *J Big Data*, vol. 9, no. 1, p. 102, Oct. 2022, doi: 10.1186/s40537-022-00652-w.
- [21] M. Hashmi, "An Efficient Pneumonia Detection in Chest Xray Images using Deep Transfer Learning," 2020, [Online]. Available: https://www.researchgate.net/figure/Chest-Xray-of-a-a-healthy-person-and-b-a-person-suffering-from-pneumonia_fig2_342210370