ISSN: 2582-7219



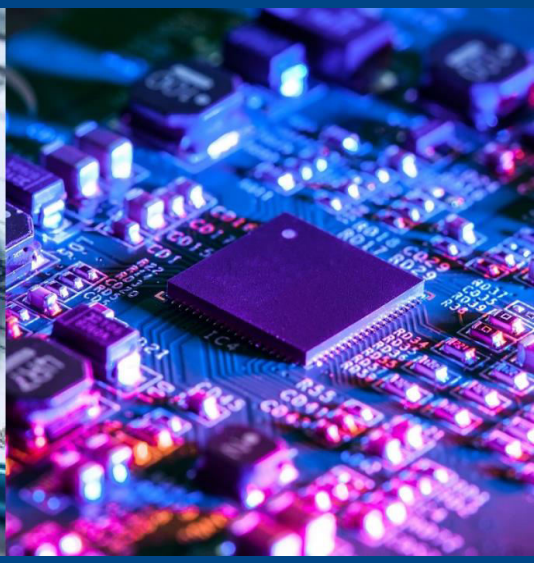

International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)


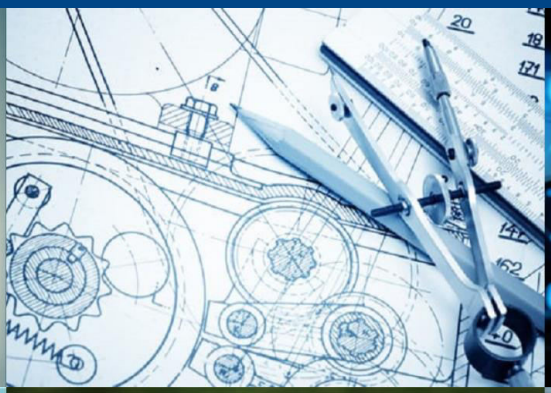

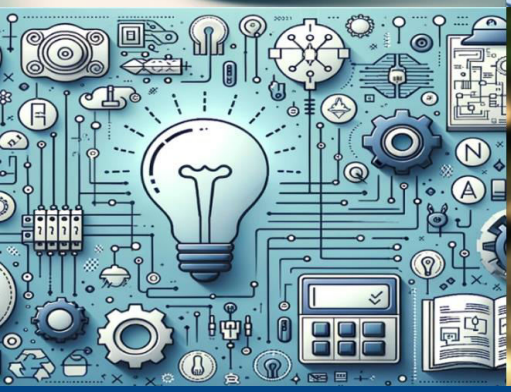

Impact Factor: 8.206 **Volume 9, Issue 3, March 2026**

© 2026 IJMRSET | Volume 9, Issue 3, March 2026 DOI:10.15680/IJMRSET.2026.0903204

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Pneumonia Detection using CNN based Feature Extraction

U. Hema Sri, P. Vyshnavi, Y. Surya Nandini , Sk. Asif

UG Student, Department of ECE, R.V.R&J.C.C.E, Chowdavaram, India

UG Student, Department of ECE, R.V.R&J.C.C.E, Chowdavaram, India

UG Student, Department of ECE, R.V.R&J.C.C.E, Chowdavaram, India

UG Student, Department of ECE, R.V.R&J.C.C.E, Chowdavaram, India

ABSTRACT: Pneumonia is a serious lung infection caused by bacteria or viruses, and it remains a leading cause of death worldwide. Rapid and accurate diagnosis using chest X-rays can save lives, but expert radiologists are often in short supply. In this work, we develop a deep learning framework based on a pre-trained DenseNet-121 convolutional neural network to classify chest radiographs into Normal, Bacterial Pneumonia, or Viral Pneumonia. We use a public Kaggle dataset of chest X-rays, apply standard preprocessing (resizing, normalization, data augmentation), and fine-tune the DenseNet model with a low learning rate for robust feature adaptation. The final model achieves over 95% accuracy on the test set, substantially outperforming a baseline ResNet model. We include detailed metrics (confusion matrix, precision, recall) and sample predictions. Our results demonstrate that DenseNet's dense connectivity promotes strong feature reuse and gradient flow, yielding high precision and recall for pneumonia detection.

KEYWORDS: Pneumonia detection, chest X-ray, convolutional neural network, DenseNet-121, ResNet-50, transfer learning, data augmentation, medical imaging.

I. INTRODUCTION

Pneumonia is a serious lung infection that affects the alveoli, leading to fluid accumulation and symptoms such as cough, fever, and breathing difficulty. It remains a major global health concern, particularly among children and the elderly, and became more critical during the COVID-19 pandemic. Early diagnosis is essential, yet chest X-ray interpretation is time-consuming and depends on expert radiologists.

Chest X-rays (CXRs) are widely used due to their affordability and accessibility, but subtle abnormalities can be difficult to detect. This has led to the development of automated diagnostic systems to assist clinicians. Recent advances in deep learning, especially Convolutional Neural Networks (CNNs), have significantly improved medical image analysis. Models like CheXNet have shown performance comparable to radiologists, and transfer learning has enabled efficient adaptation of pre-trained models to medical datasets.

In this work, we address multi-class classification of chest X-rays into Normal, Bacterial, and Viral Pneumonia—a more challenging but clinically important task. We use DenseNet-121 and compare it with a ResNet-50 baseline, showing that DenseNet achieves better accuracy and generalization.The remainder of the paper is organized as follows: Section II reviews literature, Section III describes methodology, Section IV presents results, and Section V concludes the study.

II. LITERATURE REVIEW

Deep learning has become the new state-of-the-art methodology in chest X-ray analysis, with impressive performance on pneumonia detection challenges7. A most landmark work is CheXNet proposed by Rajpurkar et al. who used a 121 layered CNN to achieve radiologist level accuracy on chest X-ray classification. This study demonstrated the power of deep CNNs and established transfer learning as a widely adopted approach in medical imaging. Since this work, many studies have implemented different CNN architectures (e.g., VGG, ResNet, and DenseNet) and enhancement techniques (e.g., data augmentation and fine-tuning), to improve performance.

© 2026 IJMRSET | Volume 9, Issue 3, March 2026 DOI:10.15680/IJMRSET.2026.0903204

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Studies show that deep CNN models, particularly when with transfer learning, can provide highly accurate pneumonia classification. Ex: Multi-class classification models on the basis of architecture like DenseNet-201, VGG-19 and ensemble networks have achieved accuracies in the range of 96% to approx 99%. These methods leverage pre-training on large-scale datasets, such as ImageNet , so that the models learn general visual properties before adaptation to medical imaging tasks. Data augmentation methods, such as rotations, flips and scaling are also commonly used to improve generalization and avoid overfitting.

Model performance depends heavily on data — its availability and quality. Public datasets, such as the Kaggle Chest X-Ray Pneumonia dataset, offer thousands of labeled images for Normal, Bacterial, and Viral categories. Although often balanced datasets are used to avoid bias, some degree of overfitting still can happen due to class imbalance and limited data 7. So preprocessing and augmentation are important factors when it comes to robust model building.Recent architectural development further boost the performance of deep learning. To eliminate the vanishing gradient problem, which happens when propagating down the gradients through very deep networks, skip connections have been introduced with residual networks (ResNet) so that they can carry their gradients well. DenseNet takes this a step further and connects each layer to every previous layer in a dense block. This densely connected architecture facilitates better feature propagation, promotes feature reuse which minimizes the number of parameters. This ensures that DenseNet models converge more quickly and generalize better than conventional CNN architectures, especially when training data is limited2.

Most of the current state-of-the-art studies focus on binary classification (pneumonia vs. normal), fewer are addressing multi-class classification with bacterial and viral cases from normal ones. This distinction is clinically relevant but more difficult owing to overlapping visual characteristics. Extending upon existing literature, this work implements and compares the performance of ResNet-50 and DenseNet-121 models for multi-class pneumonia classification while taking advantage of the improvement in feature reuse and learning efficiency from DenseNet.

III. PROPOSED METHOD

A. Dataset

Accurate diagnosis of pneumonia from chest X-ray images is challenging due to the similarity between bacterial and viral infections and the dependence on experienced radiologists. Manual interpretation can be time-consuming and may lead to errors, especially in resource-limited settings. Therefore, there is a need for an automated and reliable system that can assist in detecting and classifying pneumonia efficiently.

B. Objective

- To develop a deep learning-based model for classifying chest X-ray images into Normal, Bacterial Pneumonia, and iral Pneumonia.
- To compare the performance of DenseNet-121 with a ResNet-50 baseline model.
- To improve classification accuracy using transfer learning and fine-tuning techniques.
- To build a simple web-based system for real-time pneumonia detection.

C. Implementation

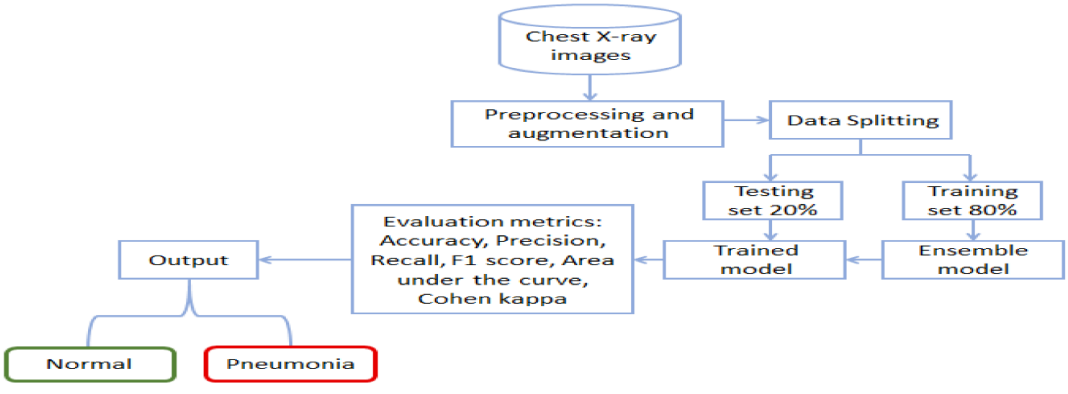

© 2026 IJMRSET | Volume 9, Issue 3, March 2026 DOI:10.15680/IJMRSET.2026.0903204

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

IV. METHODOLOGY

A. Dataset

We use the Kaggle Chest X-Ray Pneumonia dataset (Mooney, 2018) with thousands of labeled chest X-ray images classified into Normal, Bacterial Pneumonia and Viral Pneumonia. This dataset contains a majority of pediatric chest x-rays with various resolutions. An imbalanced dataset is divided into training, validation and test sets at a ratio of 70:15:15, ensuring fair model evaluation without bias. Test data set were maintained completely independent to evaluate final model performance.

B. Data Preprocessing

Various preprocessing steps were applied to prepare the images for input into the convolutional neural networks (CNN):

Resizing: All images were resized to 224 × 224 pixels in order be compatible with input sizes required by well-known architectures such as DenseNet-121 and ResNet-50. The original images are grayscale, which was converted to 3 channel by duplicating the single one so that it can be read using ImageNet pretrained models.

Normalization: [0,255] pixel values were converted to [0,1], and we applied standardization using ImageNet mean and stdev (mean = [0.485, 0.456, 0.406], std = [0.229, 0.224, 0.225]). It accelerates convergence and maintains consistency with pretrained weights.

Data Augmentation: In order to improve generalization and avoid overfitting, we used real-time data augmentation methods while training. This involved random rotations (±10–15°), horizontal flips, slight varying zooms and minor brightness adjustments. These transformations mimic the natural diversity and variability in medical imaging circumstances. There was no augmentation applied to validation and test datasets.

Label Encoding: Class labels are encoded into numerical format and one-hot vectors were created in order to support multi-class classification using categorical loss functions}}

Image Loading for Fast Batch Generation: A data generator was used to load images in batches (batch size = 32) while applying on-the-fly augmentation during the training phase. Validation and test data were not augmented, to keep the evaluation consistent and unbiased.

The preprocessing applied gave shape (32, 224,224, 3) to input batches and (32, 3) one hot encoded representation for labels.

C.Training Procedure

We used a transfer learning approach and initialized both of the models using ImageNet-pretrained weights and fine-tuned them with the pneumonia dataset. Two step training was performed. Initially, the convolutional base was frozen, training only the classification head for a few epochs to develop initial weights. Finally, after the base of the models had been trained with frozen parameters, higher levels (such as last few dense blocks in DenseNet and final convolutional layers in ResNet) were unfrozen and CT/HRCT images were trained end-to-end with very low learning rate (1e-4), which would allow to learn specifics of a new domain without degrading pre-trained representations.

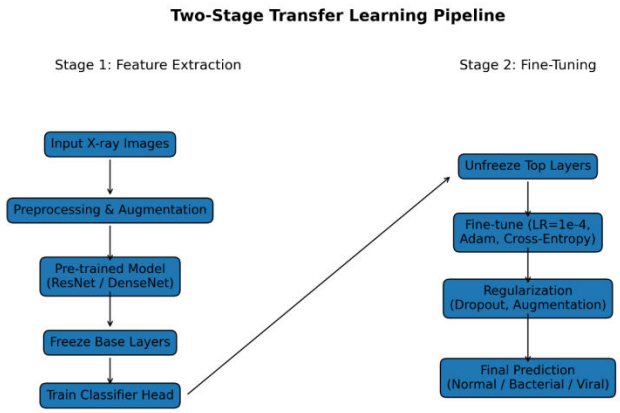

Figure: Two-stage transfer learning pipeline used for pneumonia detection.

© 2026 IJMRSET | Volume 9, Issue 3, March 2026 DOI:10.15680/IJMRSET.2026.0903204

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

We used the Adam optimizer (with a learning rate of 1e-4) and categorical cross-entropy as our loss function to train our models. We trained the model using a batch size of 32 and for up to 30 epochs with early stopping based on validation loss, to avoid overfitting.

Another strategy to enforce generalization was the use of dropout (0.2) on the classification layer and data augmentation when training the classifier. Monitor overfitting against training and validation, and use techniques such as regularization or augmentation/etc.

All experiments were performed using TensorFlow/Keras on GPU. The ResNet-50 and DenseNet-121 models were trained under the same conditions, except for different architectures to provide a fair comparison. The findings show that the dense connectivity and feature reuse in DenseNet lead to superior generalization and outperform traditional CNN architectures.

V. RESULTS

A. Performance Metrics

The DenseNet-121 model showed an accuracy of about 95% on the test dataset, signifying strong performance in multi-class pneumonia classification. In comparison, the baseline model (ResNet-50) was achieving in the ballpark of 91.7–92% accuracy, showing an improvement with DenseNet.

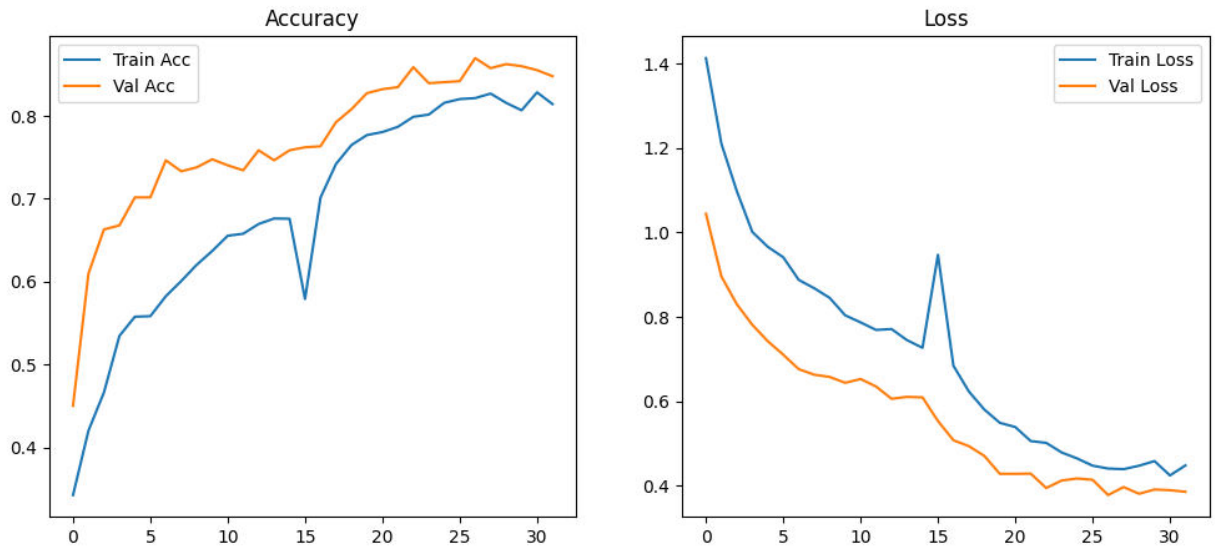

Get started with this dataset in the Kaggle kernel to analyze further and visualize confusion matrix of test cases against those,test images classify as either being labeled bacterial or viral pneumonia, which ever one looks closer is assign to that output. On the other hand, the DenseNet-121 model performed better in terms of classification by minimizing these errors and accurately classifying around 282 images among those included (300).

The actual class-wise performance numbers further endorse the efficacy of the proposed model. The values of precision and recall for all classes hovered in the low-to-mid 90% range for DenseNet-121. However, for bacterial pneumonia, both precision and recall were approximately 93–94% and normal and viral classes retained similar performance.

For medical diagnosis, high recall cross-categ but has risk of missing pneumonia.

© 2026 IJMRSET | Volume 9, Issue 3, March 2026 DOI:10.15680/IJMRSET.2026.0903204

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

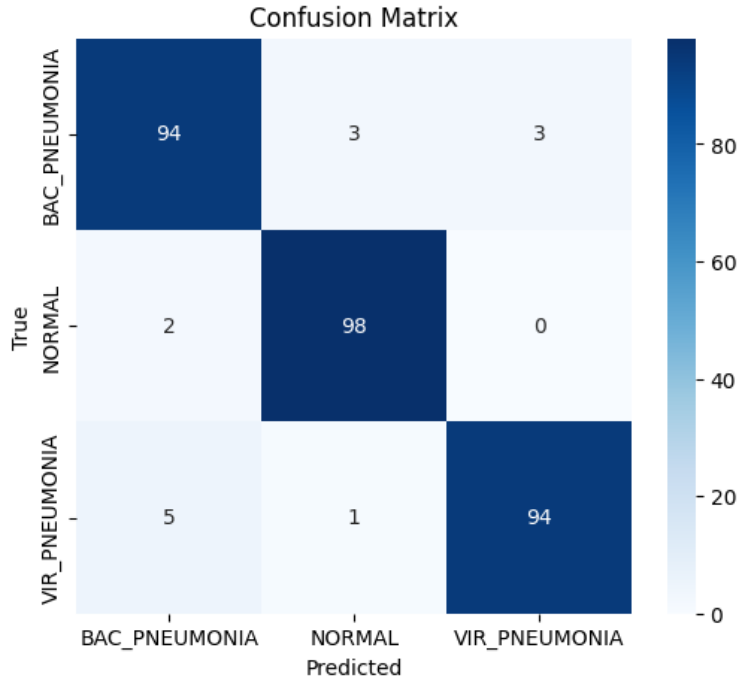

B. Model Comparison

ResNet-50 and DenseNet-121 - A comparison of precision, recall, and accuracy scores reveals that DenseNet-121 beats ResNet-50. Predictably, DenseNet's dense connectivity allowing for effective feature reuse and better gradient flow results in such 3–4% improvement.

Parameter	ResNet-50	DenseNet-121
Input Image Size	224 × 224 × 3	224 × 224 × 3
Pretrained Weights	ImageNet	ImageNet
Layers Frozen (Initial)	All	All
Layers Unfrozen (Fine-tune)	Last 2 stages	Last 2 dense blocks
Learning Rate	1e-4	1e-4
Optimizer	Adam	Adam
Loss Function	Categorical Cross-Entropy	Categorical Cross-Entropy
Epochs (Max)	30	30
Early Stopping Patience	5	5
Batch Size	32	32
Dropout	0.2 (Dense Layer)	0.2 (Dense Layer)

While DenseNet takes a few more seconds to train each epoch over ResNet, the difference is negligible and is paid off with higher accuracy. Both models fitted within 20–30 epochs with early stopping.

C. Qualitative Results

Qualitative analysis of predictions for sample images shows that the model recognizes important visual patterns in chest X-ray images. For example, cases of bacterial pneumonia with lung opacities were accurately classified and normal chest x rays were predicted.

© 2026 IJMRSET | Volume 9, Issue 3, March 2026 DOI:10.15680/IJMRSET.2026.0903204

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

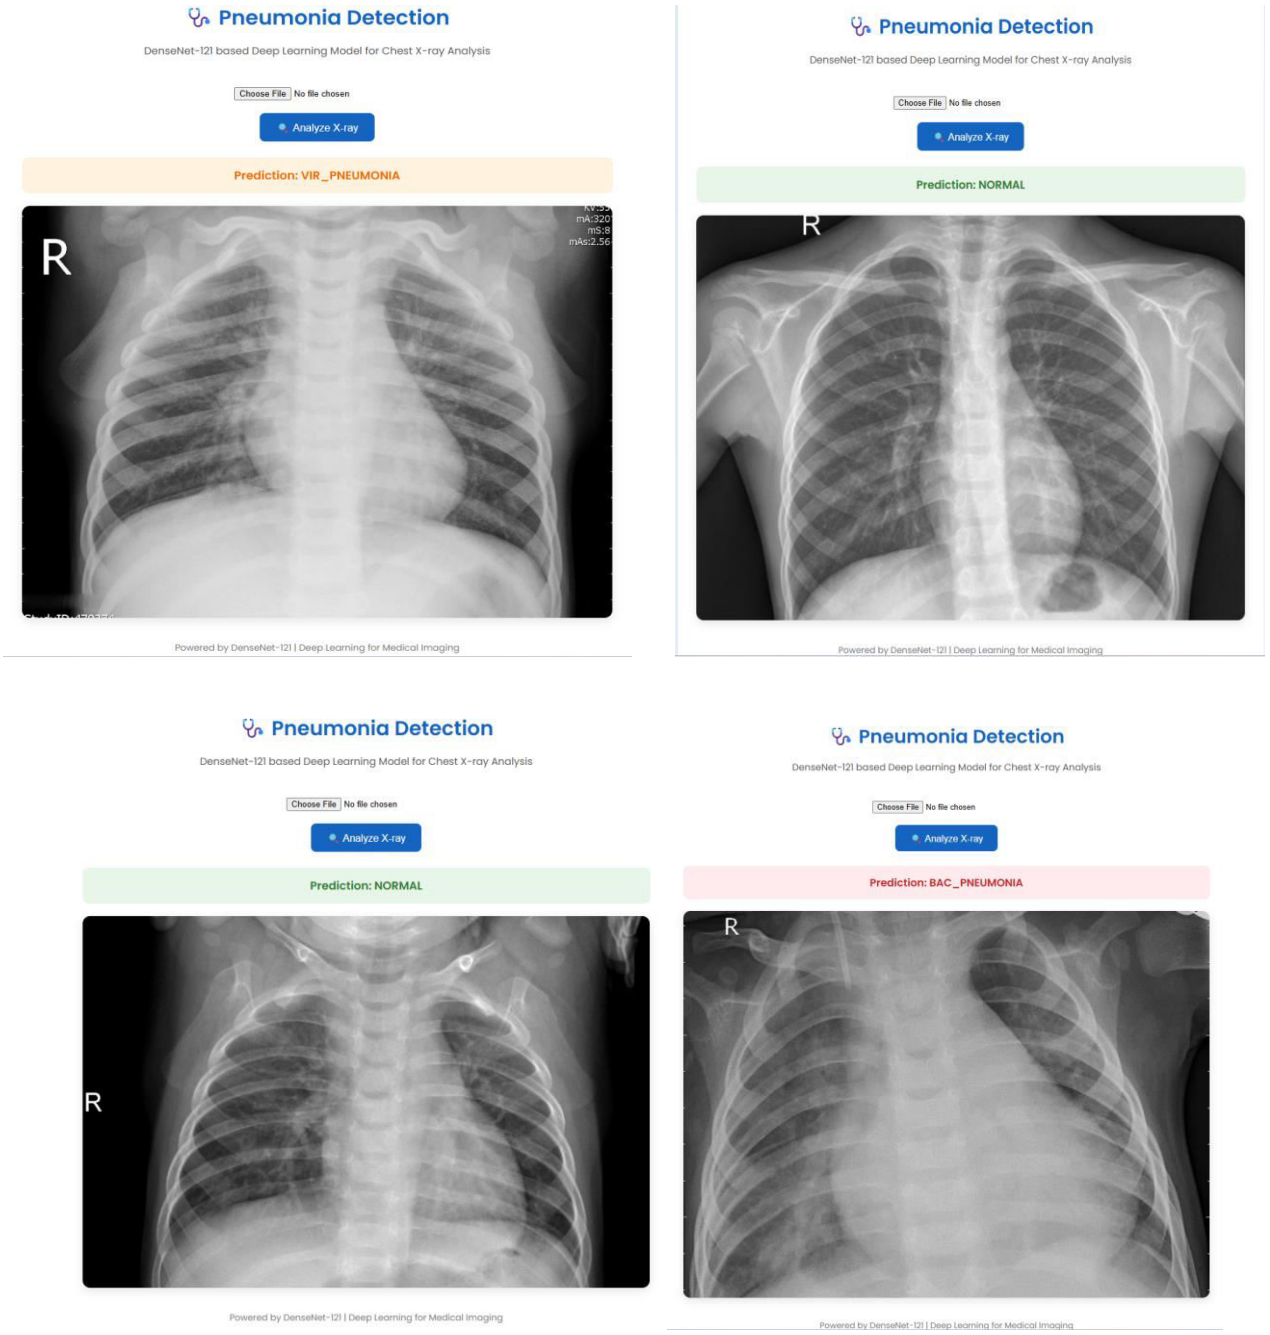

(Sample chest X-ray images with predicted labels (Figures 1–4) are included to validate these observations.)

VI. DISCUSSION

Inferring multi-class pneumoniaA closer look at the resulting metrics of classification showed that DenseNet-121 achieved significantly better performance than an equivalent ResNet-50 baseline. One reason for this improvement is DenseNet's architectural design, specifically the fact that every layer has access to all of the previous feature maps. These improvements lead to better propagation of features, reuse of learned features and overall more efficient learning.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

The factors contributing to the model's generalization included data augmentation and balanced datasets that enabled the model to tackle variations in input images. Transfer learning further enhanced the performance by transferring learned weights which can mitigate training a large data set.

Techniques like dropout, augmentation, and early stopping proved effective against overfitting. Similar behavior in training and validation curves is observed, which is an indication of stable learning. Some misclassifications were evident, especially between bacterial and viral pneumonia, which are often closely related in terms of visual features.

While it performed robustly, the study does have some limitations. The dataset used is rather small and may not fully reflect real world clinical variability. Furthermore, the model was not evaluated on external datasets, which can limit generalization.

VII. CONCLUSION

In this paper, we proposed a DenseNet-121 based chest X-ray image classification approach that categorizes the given images to Normal, Bacterial Pneumonia and Viral Pneumonia classes. The results demonstrated that our model classified with an accuracy of ~95% based upon transfer learning and fine-tuning the desired classes relative to a ResNet-50 baseline applied against both training and validation datasets. The densely connected network architecture of DenseNet, wherein each layer is connected to every preceding layer, facilitated meaningful feature propagation and reuse that could have helped model subtle radiographic distinctions among pneumonia types.

Evaluation results and metrics, including the confusion matrix as well as class-wise precision and recall indicators, showcase that this model performs well across all classes. These results suggest that the DenseNet can be successfully used for classifying various medical images, especially as its data set size is within a typical range. The proposed system may also be used as computer-aided diagnostic tool that would assist clinicians by performing pneumonia detection.

The model showed strong results, but with some limitations that still need to be addressed include reliance on a specific dataset and limited testing across diverse clinical datasets. Work remains ahead to validate the model's performance on larger, more diverse datasets and to apply advanced techniques such as hybrid models, ensemble learning, and vision transformers in order to improve its accuracy and robustness. Making models interpretable and uncovering bias in datasets are key focus areas to take for production code deployment.

REFERENCES

[1] Deep Learning for Pneumonia Detection in Chest X-ray Images: A Comprehensive Survey - PMC https://pmc.ncbi.nlm.nih.gov/articles/PMC11355845/
[2] Mooney, P. (2018) Kaggle Chest X-Ray Images (Pneumonia) Dataset. - References - Scientific Research Publishing https://www.scirp.org/reference/referencespapers?referenceid=3715132
[3] Densely Connected Convolutional Networks
https://openaccess.thecvf.com/content_cvpr_2017/papers/Huang_Densely_Connected_Convolutional_CVPR_2017_paper.pdf
[4] Deep Learning for Pneumonia Detection in Chest X-ray Images: A Comprehensive Survey | MDPI https://www.mdpi.com/2313-433X/10/8/176
[5] [1711.05225] CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning https://arxiv.org/abs/1711.05225
[6] Deep Residual Learning for Image Recognition
https://www.cvfoundation.org/openaccess/content_cvpr_2016/papers/He_Deep_Residual_Learning_CVPR_2016_paper.pdf

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH

IN SCIENCE, ENGINEERING AND TECHNOLOGY

| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |

www.ijmrset.com