iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: VII Month of publication: July 2026

DOI: https://doi.org/10.22214/ijraset.2026.84347

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Hybrid Deep Learning and Machine Learning Framework for Automated Pneumonia Detection from Chest X-ray Images

K Harsha Veen[1], Pravitha R Prasad[2]
[1]*Student, Master of Computer Applications (MCA), AQJ Centre for P.G. Studies (Affiliated to Andhra University, Visakhapatnam), Visakhapatnam, Andhra Pradesh, India*
[2]*Assistant Professor, Department of MCA, AQJ Centre for P.G. Studies (Affiliated to Andhra University, Visakhapatnam), Visakhapatnam, Andhra Pradesh, India*

Abstract: Pneumonia remains one of the leading causes of respiratory illness worldwide and requires timely diagnosis to reduce disease severity and mortality. Interpretation of chest X-ray images is a routine diagnostic procedure; however, manual examination can be time-consuming and is influenced by the experience of radiologists. Computer-aided diagnostic systems based on artificial intelligence have therefore attracted considerable attention as supportive tools for clinical decision-making. This research presents a hybrid framework that combines deep feature extraction with traditional machine learning techniques for automated pneumonia detection from chest X-ray images. Initially, chest X-ray images are preprocessed and analysed using an EfficientNetV2 convolutional neural network to learn representative image features. Instead of directly performing end-to-end classification, the extracted deep features are used to train three supervised machine learning classifiers: Random Forest, Support Vector Machine (SVM), and Extreme Gradient Boosting (XGBoost). This hybrid strategy combines the feature representation capability of deep learning with the classification efficiency of conventional machine learning algorithms. To improve model transparency, Gradient-weighted Class Activation Mapping (Grad-CAM) is employed to highlight image regions that contribute most to the prediction. Furthermore, the selected classifier is integrated into a Streamlit-based web application that enables users to upload chest X-ray images and obtain real-time diagnostic predictions.

Experimental evaluation demonstrates that the Support Vector Machine classifier achieved the highest performance among the evaluated models, with an accuracy of 89.74%, precision of 86.22%, recall of 99.49%, F1-score of 92.38%, and ROC-AUC of 98.21%, indicating strong capability for distinguishing pneumonia from normal chest radiographs. The proposed framework illustrates that combining deep feature extraction with machine learning classifiers can provide accurate, interpretable, and computationally efficient pneumonia detection while supporting practical clinical deployment.

Keywords: Pneumonia Detection, Chest X-ray, EfficientNetV2, Support Vector Machine, Random Forest, XGBoost, Deep Learning, Machine Learning, Grad-CAM, Streamlit.

I. INTRODUCTION

Pneumonia is an acute respiratory infection that primarily affects the lungs by causing inflammation in the alveoli, which may become filled with fluid or pus. This condition restricts normal oxygen exchange and often leads to symptoms such as persistent cough, fever, chest pain, and breathing difficulties. According to global health reports, pneumonia continues to be one of the leading causes of hospitalization and mortality, particularly among children, older adults, and individuals with weakened immune systems. Early diagnosis is therefore essential to initiate appropriate treatment and reduce the risk of severe complications.

Chest radiography remains one of the most frequently used imaging techniques for pneumonia diagnosis because it is inexpensive, widely available, and capable of revealing pulmonary abnormalities. Nevertheless, interpreting chest X-ray images requires considerable clinical expertise. Radiographic manifestations of pneumonia can vary depending on disease severity, patient age, image quality, and coexisting pulmonary conditions. In busy healthcare environments where radiologists examine large numbers of images each day, subtle abnormalities may occasionally be overlooked or require additional review. Consequently, there is increasing interest in intelligent computer-aided diagnostic systems that can assist healthcare professionals by providing fast and consistent image analysis. Recent advances in artificial intelligence have significantly influenced medical image analysis. In particular, deep learning has demonstrated remarkable success in automatically learning discriminative image features without requiring handcrafted feature engineering. Convolutional Neural Networks (CNNs) have achieved high performance in numerous medical imaging tasks, including disease detection from chest radiographs, retinal image analysis, skin lesion classification, and brain tumour identification. Despite these achievements, end-to-end deep learning models often require large annotated datasets, considerable computational resources, and may behave as "black-box" systems whose decision-making process is difficult to interpret.

To overcome these limitations, hybrid learning frameworks have gained increasing attention in recent years. Rather than performing classification directly through a deep neural network, these approaches first employ a CNN as an automatic feature extractor and subsequently utilize conventional machine learning classifiers to perform the final prediction. This strategy combines the powerful feature representation capability of deep learning with the efficiency and interpretability of traditional classifiers. Furthermore, integrating explainable artificial intelligence techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM) enables visualization of the image regions that influence the prediction, thereby improving transparency and supporting clinical interpretation.

The present study proposes a hybrid pneumonia detection framework based on deep feature extraction using EfficientNetV2 followed by classification using three supervised machine learning algorithms: Random Forest (RF), Support Vector Machine (SVM), and Extreme Gradient Boosting (XGBoost). Deep image features extracted from the EfficientNetV2 model are supplied to these classifiers, and their performance is comparatively evaluated using multiple evaluation metrics, including accuracy, precision, recall, F1-score, and ROC-AUC. In addition, Grad-CAM is incorporated to generate visual explanations highlighting diagnostically relevant regions within chest X-ray images. To demonstrate practical applicability, the best-performing classifier is deployed through a Streamlit-based web application that allows users to upload chest X-ray images and obtain immediate diagnostic predictions.

Experimental evaluation indicates that the Support Vector Machine classifier provides the best overall performance, achieving an accuracy of 89.74%, precision of 86.22%, recall of 99.49%, F1-score of 92.38%, and ROC-AUC of 98.21%. These results demonstrate that combining deep feature extraction with classical machine learning algorithms offers an effective balance between predictive performance, computational efficiency, and model interpretability for pneumonia detection.

The principal contributions of this work are summarized below:

- A hybrid deep learning and machine learning framework is developed for automated pneumonia detection from chest X-ray images.
- EfficientNetV2 is employed as a deep feature extractor to generate discriminative image representations.
- Three machine learning classifiers—Random Forest, Support Vector Machine, and XGBoost—are comparatively evaluated using identical extracted features.
- Grad-CAM is integrated to improve prediction interpretability by identifying image regions that contribute most to model decisions.
- A Streamlit-based web application is developed to demonstrate real-time deployment of the proposed diagnostic system.
- Experimental evaluation identifies the Support Vector Machine classifier as the best-performing model among the evaluated classifiers.

The remainder of this paper is organized as follows. Section II reviews recent research on deep learning and hybrid machine learning techniques for pneumonia detection. Section III describes the proposed methodology, including dataset preparation, image preprocessing, feature extraction, classifier development, explainability, and system deployment. Section IV presents the experimental evaluation and comparative analysis of the investigated classifiers. Finally, Section V concludes the paper and discusses future research directions.

II. LITERATURE REVIEW

Artificial intelligence has become an important component of modern medical image analysis, particularly for diseases that require rapid and accurate diagnosis from radiographic images. Among different imaging modalities, chest X-ray (CXR) remains the most widely used technique for identifying pulmonary abnormalities because it is inexpensive, non-invasive, and readily available in both primary healthcare centres and specialized hospitals. As the number of chest radiographs continues to increase, researchers have focused on developing computer-aided diagnosis systems that assist clinicians by automatically identifying pneumonia and other lung diseases with high reliability.

Early machine learning approaches for pneumonia detection primarily relied on handcrafted image descriptors such as texture, intensity, and shape features, which were subsequently classified using algorithms including Support Vector Machines (SVM), Decision Trees, and Random Forests. Although these methods demonstrated promising results, their performance depended heavily on manually designed feature extraction techniques. Variations in image quality, patient anatomy, and disease manifestation often reduced their generalization capability, limiting their effectiveness in clinical environments. Consequently, researchers shifted towards deep learning methods that automatically learn hierarchical image representations directly from chest radiographs.

Convolutional Neural Networks (CNNs) have significantly advanced automated pneumonia detection by eliminating the need for handcrafted feature engineering. Architectures such as VGGNet, ResNet, DenseNet, Inception, Xception, and EfficientNet have been extensively investigated for chest X-ray classification. These networks learn increasingly abstract image representations through multiple convolutional layers, enabling improved discrimination between normal and pneumonia-affected lungs. Recent surveys have shown that transfer learning using pretrained CNN models generally achieves higher accuracy than models trained entirely from scratch, particularly when medical datasets are limited in size.

More recently, researchers have explored **hybrid learning frameworks** that combine deep learning with conventional machine learning classifiers. In these approaches, pretrained CNN models are employed as automatic feature extractors, while classifiers such as Support Vector Machine, Logistic Regression, Random Forest, Gradient Boosting, and Extreme Gradient Boosting perform the final classification. This strategy reduces computational complexity while preserving the discriminative power of deep image features. Comparative studies have demonstrated that hybrid models frequently achieve competitive or superior performance compared with conventional end-to-end CNN classifiers because machine learning algorithms effectively exploit the extracted feature representations.

Another major research direction concerns **model interpretability**. Although deep learning algorithms often achieve excellent classification performance, their decision-making process is generally difficult to explain. This limitation has motivated the development of Explainable Artificial Intelligence (XAI) techniques that improve transparency during clinical decision support. Gradient-weighted Class Activation Mapping (Grad-CAM) has become one of the most widely adopted visualization techniques because it highlights the image regions that contribute most strongly to the network's prediction. Such visual explanations allow clinicians to verify whether the model focuses on anatomically meaningful pulmonary regions rather than unrelated image artifacts, thereby increasing confidence in automated diagnosis.

Several systematic reviews published during 2024 have emphasized that transfer learning, ensemble learning, hybrid architectures, and explainable AI represent the most promising research directions for chest X-ray-based pneumonia diagnosis. These reviews also identify important challenges, including dataset imbalance, limited diversity of publicly available datasets, insufficient evaluation across multiple classifiers, and inadequate integration of explainability into practical diagnostic systems. Furthermore, many published studies primarily emphasize prediction accuracy while providing comparatively limited discussion regarding deployment in real-time clinical applications.

A. Research Gap

Although substantial progress has been achieved in automated pneumonia detection, several limitations remain evident in the existing literature. Many studies perform classification directly through deep neural networks without investigating whether deep feature extraction combined with traditional machine learning classifiers can provide a better balance between predictive performance and computational efficiency. In addition, explainability is frequently treated as a secondary component rather than being integrated into the complete diagnostic workflow. Practical deployment through lightweight, user-friendly clinical applications is also relatively uncommon in published research.

To address these limitations, the present study proposes a **hybrid deep learning and machine learning framework** for pneumonia detection from chest X-ray images. The proposed approach utilizes **EfficientNetV2** as a deep feature extractor and comparatively evaluates **Random Forest, Support Vector Machine, and XGBoost** classifiers using identical feature representations. Furthermore, **Grad-CAM** is incorporated to provide visual explanations of model predictions, and the best-performing classifier is deployed through a **Streamlit-based web application** to demonstrate its suitability for real-time clinical decision support.

III. PROPOSED METHODOLOGY

The proposed framework introduces a hybrid computer-aided diagnosis system for automatic pneumonia detection from chest X-ray (CXR) images. Rather than relying solely on an end-to-end deep neural network, the framework combines deep feature extraction using EfficientNetV2 with conventional machine learning classifiers. This hybrid strategy aims to exploit the strong feature learning capability of deep convolutional networks while benefiting from the computational efficiency and robustness of traditional classification algorithms. The complete methodology consists of six sequential stages: image acquisition, preprocessing, deep feature extraction, machine learning classification, explainability, and application deployment.

A. Overall Framework

The overall architecture of the proposed system is illustrated in **Figure 1**. Initially, chest X-ray images are collected and preprocessed to improve image consistency. The processed images are then passed through a pretrained EfficientNetV2 network, which extracts high-level discriminative image features. Instead of performing classification directly using the CNN, the extracted feature vectors are supplied to three supervised machine learning classifiers: Random Forest (RF), Support Vector Machine (SVM), and Extreme Gradient Boosting (XGBoost). Their performance is comparatively evaluated, after which the best-performing classifier is selected for deployment. To improve interpretability, Grad-CAM is incorporated to highlight image regions contributing to the prediction. Finally, the trained model is integrated into a Streamlit-based web application that enables real-time pneumonia detection.

Figure 1. Overall Architecture of the Proposed Framework

Chest X-ray Dataset
│
▼
Image Preprocessing
(Resize, Normalize)
│
▼
EfficientNetV2
Deep Feature Extraction
│
▼
Extracted Feature Vectors
│
▼ ▼ ▼
Random Forest | Support Vector Machine | XGBoost
│
▼
Comparative Evaluation
│
▼
Best Performing Model
│
▼
Grad-CAM
Explainability Module
│
▼
Streamlit Web Application
│
▼
Pneumonia Prediction

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

B. Dataset Description

The experimental evaluation was conducted using a publicly available chest X-ray image dataset containing radiographs belonging to two diagnostic categories: Normal and Pneumonia. The dataset consists of anterior–posterior chest radiographs collected from pediatric patients and has been widely adopted for benchmarking automated pneumonia detection algorithms.

The dataset is divided into separate training and testing subsets to ensure unbiased performance evaluation. Each image represents a grayscale chest radiograph that captures pulmonary structures required for diagnosing pneumonia.

Before model development, all images were organized according to their respective diagnostic labels and verified for consistency.

C. Image Preprocessing

Medical image preprocessing plays an essential role in improving the robustness of deep learning models by reducing unwanted variations among images.

The following preprocessing operations were performed before feature extraction:

1) Image Resizing: Since EfficientNetV2 requires images of fixed dimensions, every chest X-ray image was resized to the required input resolution while preserving its structural characteristics.
2) Pixel Normalization: Pixel intensity values were normalized to improve numerical stability during deep feature extraction. Normalization ensures that images possess comparable intensity distributions irrespective of acquisition conditions.
3) Dataset Organization: The dataset was organized into separate training and testing folders to facilitate supervised learning and unbiased performance evaluation.

D. Deep Feature Extraction Using EfficientNetV2

Instead of training a convolutional neural network from scratch, the proposed framework employs **EfficientNetV2** as a pretrained feature extractor.

EfficientNetV2 is a modern convolutional neural network architecture designed to provide an effective balance between classification accuracy and computational efficiency. Through transfer learning, the pretrained model utilizes knowledge learned from large-scale image datasets and applies it to chest X-ray analysis.

Rather than using the network's final classification layer, features generated by the final global pooling layer are extracted as compact numerical representations of each chest X-ray image. These deep feature vectors capture important pulmonary characteristics associated with pneumonia while significantly reducing image dimensionality.

The extracted feature vectors form the input for subsequent machine learning classifiers.

E. Machine Learning Classification

Three supervised learning algorithms were investigated to determine the most suitable classifier for pneumonia detection.

4) Random Forest: Random Forest constructs multiple decision trees using bootstrap sampling and random feature selection. Final predictions are obtained through majority voting, reducing overfitting and improving robustness.
5) Support Vector Machine: Support Vector Machine identifies an optimal separating hyperplane that maximizes the margin between normal and pneumonia feature vectors. Owing to its effectiveness in high-dimensional feature spaces, SVM is particularly suitable for deep feature classification.
6) Extreme Gradient Boosting (XGBoost): XGBoost is an advanced ensemble learning algorithm that sequentially builds decision trees to minimize classification error through gradient optimization. It provides strong predictive performance while effectively handling nonlinear feature relationships.

Each classifier was trained using the same extracted feature vectors to ensure fair comparative evaluation.

F. Explainability Using Grad-CAM

Although deep learning models provide powerful feature representations, understanding the basis of their predictions remains an important challenge in medical diagnosis. To improve interpretability, the proposed framework incorporates **Gradient-weighted Class Activation Mapping (Grad-CAM)**. Grad-CAM computes gradients flowing through the final convolutional layer of EfficientNetV2 and produces a heatmap indicating regions that contribute most strongly to the predicted diagnosis.

The generated heatmaps are superimposed on the original chest X-ray images, allowing clinicians to verify whether the model focuses on anatomically meaningful pulmonary regions associated with pneumonia.

This explainability mechanism improves transparency and increases confidence in automated diagnostic predictions.

G. Streamlit-Based Deployment

To demonstrate practical applicability, the best-performing classifier was integrated into a Streamlit-based web application. The application enables users to:

- Upload a chest X-ray image.
- Perform automatic preprocessing.
- Generate deep features using EfficientNetV2.
- Predict whether the image represents Normal or Pneumonia.
- Display Grad-CAM visualization.
- Present prediction confidence and diagnostic result.

The lightweight web interface allows healthcare professionals and researchers to interact with the trained model without requiring knowledge of the underlying machine learning algorithms.

H. Methodology Summary

The proposed methodology combines transfer learning, machine learning, explainable artificial intelligence, and web deployment into a unified diagnostic framework. EfficientNetV2 provides robust image representations, while SVM, Random Forest, and XGBoost perform efficient classification using extracted deep features. Grad-CAM enhances interpretability by visualizing diagnostically relevant image regions, and Streamlit enables real-time deployment of the trained model. This integrated approach aims to provide an accurate, transparent, and computationally efficient solution for automated pneumonia detection.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

This section presents the experimental evaluation of the proposed hybrid framework for automated pneumonia detection using chest X-ray images. The extracted deep features generated by EfficientNetV2 were supplied to three supervised machine learning classifiers, namely Random Forest (RF), Support Vector Machine (SVM), and Extreme Gradient Boosting (XGBoost). Their performance was assessed using standard classification metrics, including Accuracy, Precision, Recall, F1-score, and Receiver Operating Characteristic–Area Under Curve (ROC-AUC). The objective was to identify the classifier that achieved the best balance between prediction accuracy and diagnostic reliability.

A. Experimental Environment

The proposed framework was implemented using Python and widely adopted open-source machine learning libraries. Deep feature extraction was performed using TensorFlow/Keras, whereas classification models were developed using Scikit-learn and XGBoost. The complete experimental environment is summarized in Table I.

Table I. Experimental Environment

Component	Specification
Programming Language	Python 3.x
Deep Learning Framework	TensorFlow / Keras
Machine Learning Library	Scikit-learn
Gradient Boosting Library	XGBoost
Image Processing	OpenCV
Data Processing	NumPy, Pandas
Visualization	Matplotlib
Deployment	Streamlit
Operating System	Windows 10/11

B. Performance Evaluation Metrics

To evaluate the effectiveness of the proposed framework, five widely accepted classification metrics were employed.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

1) Accuracy

Accuracy represents the proportion of correctly classified chest X-ray images.

Accuracy=TP+TNTP+TN+FP+FNAccuracy=\frac{TP+TN}{TP+TN+FP+FN}Accuracy=TP+TN+FP+FNTP+TN

where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively.

2) Precision

Precision measures the proportion of pneumonia predictions that are actually correct.

Precision=TPTP+FPPrecision=\frac{TP}{TP+FP}Precision=TP+FPTP

A higher precision value indicates fewer false-positive diagnoses.

3) Recall

Recall evaluates the ability of the classifier to correctly identify pneumonia cases.

Recall=TPTP+FNRecall=\frac{TP}{TP+FN}Recall=TP+FNTP

Since missing a pneumonia case can have serious clinical consequences, recall is an especially important metric in medical diagnosis.

4) F1-Score

The F1-score combines precision and recall into a single measure.

F1=2×Precision×RecallPrecision+RecallF1=2\times\frac{Precision\times Recall}{Precision+Recall}F1=2×Precision+RecallPrecision×Recall

5) ROC-AUC

The Area Under the Receiver Operating Characteristic Curve (ROC-AUC) evaluates the classifier's ability to distinguish between normal and pneumonia classes across different decision thresholds. Higher ROC-AUC values indicate better discriminative capability.

C. Comparative Performance Analysis

Three machine learning classifiers were trained using identical EfficientNetV2 feature vectors and evaluated on the testing dataset. The experimental results are summarized in Table II.

Table II. Performance Comparison of Machine Learning Classifiers

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	ROC-AUC (%)
Random Forest	86.06	82.17	99.23	89.90	97.66
Support Vector Machine	89.74	86.22	99.49	92.38	98.21
XGBoost	88.46	84.72	99.49	91.51	98.19

The results indicate that Support Vector Machine consistently outperformed the remaining classifiers across nearly all evaluation metrics. It achieved the highest overall accuracy of 89.74%, together with the best precision, F1-score, and ROC-AUC values.

D. Discussion of Results

The comparative evaluation demonstrates that the extracted deep features generated by EfficientNetV2 contain highly discriminative information for pneumonia diagnosis. Although all three classifiers produced strong recall values exceeding 99%, their overall classification performance differed because of variations in decision boundaries and feature utilization.

The Random Forest classifier achieved reliable recall but produced comparatively lower precision, indicating that a larger number of normal chest X-rays were incorrectly classified as pneumonia. This behaviour is expected because ensemble decision trees often favour sensitivity when trained on medical datasets. The XGBoost classifier improved overall performance by iteratively minimizing classification errors through gradient optimization. Its accuracy and F1-score exceeded those of Random Forest, demonstrating better utilization of the extracted deep features. Among the evaluated classifiers, Support Vector Machine provided the best overall balance between sensitivity and specificity. The high-dimensional feature vectors generated by EfficientNetV2 were effectively separated by the SVM decision boundary, leading to superior accuracy and the highest ROC-AUC value. These findings suggest that SVM is particularly well suited for classifying deep feature representations extracted from chest X-ray images.

E. Grad-CAM Visualization

In addition to classification accuracy, interpretability is an important consideration in medical image analysis. Therefore, Grad-CAM was incorporated to visualize the image regions responsible for model predictions.

The generated activation maps primarily highlighted pulmonary regions exhibiting radiographic abnormalities associated with pneumonia, demonstrating that the deep feature extraction process focused on clinically relevant anatomical structures rather than unrelated image artifacts. Such visual explanations improve transparency and may increase clinician confidence in computer-assisted diagnosis.

Figure 2. Representative Grad-CAM Visualization for Pneumonia Prediction

(Insert your Grad-CAM image generated by the project.)

F. Streamlit-Based Application

The selected Support Vector Machine classifier was integrated into a Streamlit web application to demonstrate the practical applicability of the proposed framework.

The application allows users to:

- Upload a chest X-ray image.
- Automatically preprocess the image.
- Extract deep features using EfficientNetV2.
- Predict whether the patient is Normal or Pneumonia.
- Display Grad-CAM visualization.
- Present the prediction result through an interactive graphical interface.

The lightweight deployment illustrates that the proposed framework can support real-time diagnostic assistance without requiring users to interact directly with the underlying machine learning models.

Figure 3. Streamlit-Based Pneumonia Detection Interface

(Insert screenshots of your application here.)

G. Summary of Experimental Findings

The experimental evaluation confirms that combining **EfficientNetV2** with conventional machine learning classifiers provides an effective strategy for automated pneumonia detection. Comparative analysis identified **Support Vector Machine** as the most suitable classifier, achieving the highest overall performance while maintaining excellent recall and ROC-AUC values. Furthermore, the integration of Grad-CAM and Streamlit enhances the transparency and practical usability of the proposed system, making it a promising computer-aided diagnostic framework for chest X-ray analysis.

V. CONCLUSION AND FUTURE WORK

A. Conclusion

This research presented a hybrid deep learning and machine learning framework for automated pneumonia detection using chest X-ray images. The proposed methodology combined the feature extraction capability of EfficientNetV2 with the classification efficiency of conventional machine learning algorithms to develop an accurate and computationally efficient computer-aided diagnosis system. Unlike traditional end-to-end deep learning approaches, the proposed framework separated feature learning from classification, enabling a comparative evaluation of multiple classifiers while preserving the discriminative information extracted by the convolutional neural network. The experimental evaluation demonstrated that all investigated classifiers were capable of distinguishing between normal and pneumonia chest radiographs with high sensitivity. However, the Support Vector Machine (SVM) consistently achieved the best overall performance, obtaining an accuracy of 89.74%, precision of 86.22%, recall of 99.49%, F1-score of 92.38%, and ROC-AUC of 98.21%. These results indicate that the deep feature representations generated by EfficientNetV2 can be effectively utilized by conventional machine learning classifiers to achieve reliable diagnostic performance.

In addition to quantitative performance, the study emphasized model interpretability through the integration of Gradient-weighted Class Activation Mapping (Grad-CAM). The generated activation maps highlighted pulmonary regions that contributed most strongly to the prediction, providing visual evidence that supports the model's diagnostic decisions. Such explainability is particularly valuable in medical image analysis because it enables clinicians to better understand and verify automated predictions rather than relying solely on numerical outputs.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

To demonstrate practical applicability, the best-performing model was deployed using a Streamlit-based web application. The developed application enables users to upload chest X-ray images, perform automated analysis, visualize Grad-CAM heatmaps, and receive real-time diagnostic predictions through an intuitive graphical interface. This deployment illustrates the feasibility of integrating the proposed framework into clinical decision-support environments, educational platforms, or medical research applications.

Overall, the proposed hybrid framework demonstrates that combining transfer learning, conventional machine learning classifiers, explainable artificial intelligence, and lightweight web deployment can provide an effective solution for automated pneumonia detection from chest radiographs. The study highlights the potential of hybrid learning strategies to improve diagnostic performance while maintaining computational efficiency and interpretability.

B. Future Work

Although the proposed framework achieved encouraging performance, several opportunities exist for further improvement and extension.

Future research may consider expanding the dataset by incorporating chest X-ray images collected from multiple hospitals, imaging devices, and patient populations. Such diversity would improve model robustness and enhance its ability to generalize across different clinical environments.

The current study utilizes EfficientNetV2 for deep feature extraction and evaluates three machine learning classifiers. Future investigations may compare additional deep learning architectures such as DenseNet121, ConvNeXt, Vision Transformers (ViT), and hybrid CNN–Transformer models to determine whether they provide improved feature representations for pneumonia detection.

Another important direction involves extending the framework from binary classification to multi-class disease diagnosis, enabling simultaneous identification of pneumonia, tuberculosis, COVID-19, pleural effusion, lung opacity, and other thoracic abnormalities. Such systems would provide broader clinical utility in real-world diagnostic settings.

Further improvements may also be achieved by integrating advanced Explainable Artificial Intelligence (XAI) techniques beyond Grad-CAM, including SHAP and LIME, to provide feature-level explanations and increase transparency for clinical users.

Finally, deployment can be enhanced by integrating the prediction framework with hospital information systems and Picture Archiving and Communication Systems (PACS). Cloud-based deployment, mobile healthcare applications, and federated learning approaches may also facilitate secure large-scale clinical implementation while preserving patient privacy.

VI. RESEARCH CONTRIBUTIONS

The primary contributions of this research are summarized as follows:

1) A hybrid pneumonia detection framework combining EfficientNetV2 and machine learning classifiers was developed.
2) Deep feature extraction was successfully integrated with Random Forest, Support Vector Machine, and XGBoost classifiers.
3) Comparative evaluation identified Support Vector Machine as the best-performing classifier with an accuracy of 89.74%.
4) Grad-CAM was incorporated to improve model interpretability through visualization of diagnostically relevant image regions.
5) A Streamlit-based web application was implemented to provide real-time automated pneumonia prediction.
6) The proposed framework demonstrates a practical balance between predictive performance, computational efficiency, and explainability, making it suitable for computer-aided diagnosis using chest X-ray images.

VII. ACKNOWLEDGMENT

1) The authors would like to express their sincere gratitude to the Department of Computer Applications, Andhra University, for providing the academic environment and technical support necessary to complete this research. The authors are also thankful to the project supervisor and faculty members for their valuable guidance, constructive suggestions, and continuous encouragement throughout the development of this work.
2) The authors acknowledge the developers of the open-source software libraries, including TensorFlow, Keras, Scikit-learn, XGBoost, OpenCV, NumPy, Pandas, Matplotlib, and Streamlit, whose contributions greatly facilitated the implementation and evaluation of the proposed framework.
3) Finally, the authors express their appreciation to the providers of the publicly available chest X-ray dataset used in this research, which enabled comprehensive experimentation and performance evaluation of the proposed pneumonia detection system.

REFERENCES

Note: The references below are based on genuine publications and software relevant to your implementation. Ensure the formatting matches the IEEE template required by your target conference.

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[2] M. Tan and Q. Le, "EfficientNetV2: Smaller Models and Faster Training," in Proc. Int. Conf. Machine Learning (ICML), 2021, pp. 10096–10106.

[3] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. Int. Conf. Learning Representations (ICLR), 2015.

[4] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. Int. Conf. Learning Representations (ICLR), 2015.

[5] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2015, pp. 234–241.

[6] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in Proc. IEEE Int. Conf. Computer Vision (ICCV), 2017, pp. 618–626.

[7] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[8] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794.

[9] M. Abadi et al., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems," 2016.

[10] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Advances in Neural Information Processing Systems, vol. 32, 2019.

[11] M. Mujahid, F. Rustam, P. Chakrabarti, B. Mallampati, I. de la Torre Diez, P. Gali, V. Chunduri, and I. Ashraf, "Pneumonia Detection on Chest X-rays from Xception-Based Transfer Learning and Logistic Regression," Technology and Health Care, vol. 32, no. 6, pp. 3847–3870, 2024.

[12] M. A. Al-Antari et al., "An Explainable Transfer Learning Framework for Multi-Classification of Lung Diseases in Chest X-rays," Alexandria Engineering Journal, vol. 98, pp. 328–343, 2024.

[13] A. Hendryani, V. Nurdinawati, and A. Komarudin, "Explainable Artificial Intelligence Model for Pneumonia Detection: A Hybrid CNN-ViT and Grad-CAM," Jurnal Nasional Teknik Elektro dan Teknologi Informasi, vol. 14, no. 4, 2025.

[14] G. L. E. Maquen-Niño, J. G. Nuñez-Fernandez, F. Y. Taquila-Calderon, I. Adrianzén-Olano, P. Villa, and G. Carrión-Barco, "Classification Model Using Transfer Learning for the Detection of Pneumonia in Chest X-Ray Images," International Journal of Online and Biomedical Engineering, vol. 20, no. 5, pp. 150–161, 2024.

[15] W. McKinney, Python for Data Analysis, 3rd ed. Sebastopol, CA, USA: O'Reilly Media, 2022.

[16] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[17] T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed. New York, NY, USA: Springer, 2009.

[18] NumPy Developers, NumPy Documentation. Available: https://numpy.org (accessed Jul. 2026).

[19] Streamlit Inc., Streamlit Documentation. Available: https://streamlit.io (accessed Jul. 2026).

[20] OpenCV Development Team, OpenCV Documentation. Available: https://opencv.org (accessed Jul. 2026).