International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 03(March),2026
Received: 04-02-2026 Accepted: 11-03-2026 Published: 18-03-2026

A Hybrid Deep Learning–Machine Learning Framework for Detection of COVID-19, Pneumonia, and Lung Opacities from Chest X-rays

Tadisetti Divya[1], Nedunuri Jasmine Nirmala Varsha[2], Pothaboina Bhuvana [3], Veeranki Gagan Deep[4], Yalla Nikhila[5], Somisetty Surya Teja Karthik Gupta[6]
[1] Dept. of CSE, SR Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India, **dtadisetti8@gmail.com**.
[2] Assistant Professor, Dept. of CSE, SR Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India.
[3] Dept. of CSE, SR Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India.
[4] Dept. of CSE, SR Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India.
[5] Dept. of CSE, SR Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India.
[6] Dept. of CSE, SR Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India.

Abstract— *This project introduces a Hybrid Deep Learning-Machine Learning (DL-ML) framework for the automated detection of lung opacity and related pulmonary conditions, including COVID-19 and pneumonia, from chest X-ray images. Lung opacity is a critical radiological finding, and its manual interpretation is time-consuming and often prone to diagnostic errors, particularly in resource-limited settings. The proposed methodology employs Convolutional Neural Networks (CNNs) for advanced feature extraction to significantly boost classification accuracy. Crucially, to enhance clinical trust and address the "black box" nature of deep learning, the framework integrates Grad-CAM (Gradient-weighted Class Activation Mapping). This integrated approach aims to deliver a reliable, efficient, and explainable diagnostic tool that visually highlights abnormal lung regions, thereby directly supporting radiologists in the timely and accurate identification of severe lung pathologies.*

Keywords— Lung opacity, pneumonia, COVID-19, Feature extraction, AI, Visual explanations, Chest X-ray, Machine Learning, Transfer Learning, CNN, Grad-CAM, Deep Learning.

I. Introduction

The recent global health crisis underscored the critical importance of rapid and accurate diagnostic tools for pulmonary infections. While traditional methods like RT-PCR can be slow, X-ray imaging provides a fast, effective alternative for assessing lung conditions. Artificial Intelligence, specifically Deep Learning (DL) with CNN architectures, has proven highly effective in analyzing these medical images. However, to enhance robustness and overcome computational limitations, this project adopts a hybrid model strategy, combining the powerful feature extraction capabilities of a pre-trained CNN, such as the ResNet50 framework, with the strong generalization of a classical machine learning classifier like the Support Vector Machine (SVM). Despite the success of end-to-end CNNs, their application in medical imaging faces inherent difficulties that motivate our architectural choice.

Pure deep learning models require immense, perfectly labeled datasets for optimal training, which is a common scarcity in the medical domain. Furthermore, complex CNNs, especially those trained on limited data, often suffer from poor generalization outside of the training distribution and can be computationally expensive for deployment. Crucially, as evidenced by related work, baseline machine learning models like SVM show significantly lower performance (e.g.,68.15 accuracy) compared to modern CNNs, indicating a clear need for advanced feature engineering. Our hybrid design seeks to mitigate these issues by separating the tasks: utilizing the Transfer Learning (TL) paradigm to generate high-quality features while relying on the SVM's proven robustness and efficiency for boundary classification.

This paper introduces the ResNet50-PCA-SVM hybrid architecture, a systematic approach to feature optimization. The methodology follows a three-stage pipeline. First, the ResNet50 network is fine-tuned to capture subtle domain-specific features. Second, to combat high dimensionality and prevent overfitting, Principal Component Analysis (PCA) is applied to compress feature vectors into a robust, lower-dimensional subspace while retaining maximum variance. The training logic for this process utilizes the Adam optimizer for stable weight updates. Third, the optimized features are fed into an SVM with an RBF kernel, found highly effective for mapping complex data into separable spaces.
The study utilizes the COVID-19 Radiography Database, which highlights the importance of data quality in distinguishing visually similar conditions

2250-3676 ijesatj@gmail.com https://www.ijesat.com/

International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 03(March),2026
Received: 04-02-2026 Accepted: 11-03-2026 Published: 18-03-2026

like "Lung Opacity" and viral infections. Our results establish a substantial improvement over the 68.15% accuracy reported for pure SVM baselines. Furthermore, we provide Grad-CAM visualizations, based on the technique, to enhance the model's interpretability. This verifies that classification decisions are based on clinically significant anatomical markers, ensuring a transparent and trustworthy diagnostic tool

II. Objectives

The main objective of your project is to develop a robust hybrid diagnostic framework that integrates the high-level feature extraction capabilities of Deep Learning (DL) with the mathematical classification stability of Machine Learning (ML). By combining a pre-trained ResNet50 backbone with a Support Vector Machine (SVM), the system aims to provide accurate, explainable, and clinically reliable detection of COVID-19 and other pulmonary conditions from chest X-rays.

The next objective is Address class imbalance using techniques like class-weighted loss and balanced sampling. This ensures the model maintains diagnostic fairness and performs reliably across all categories in diverse hospital environments.

The third objective is to Use Grad-CAM heatmaps to highlight the specific lung regions influencing the model's decision. This transparency builds clinical trust by allowing radiologists to verify that the AI is focusing on actual medical evidence. And the final objective is to Conduct a rigorous statistical evaluation using Precision, Recall, and F1-Score. Special focus is placed on Recall to minimize false negatives and ensure the system meets high safety standards for clinical use.

III. Literature Review

The introduction of the Residual Learning framework (ResNet) revolutionized deep learning by utilizing identity shortcut connections to enable the training of ultra-deep networks while overcoming the degradation problem [1]. While this backbone became a gold standard for general image recognition, its application to the subtle textures of medical radiographs was justified by comprehensive surveys on Transfer Learning [2]. This research established that knowledge gained from broad domains, such as ImageNet, could be mathematically mapped to novel tasks, providing the logical basis for using pre-trained weights in clinical settings. Building on these concepts, specialized open-source frameworks like COVID-Net demonstrated that tailored architectures can outperform human radiologists in speed and consistency, though their high computational complexity remains a hurdle for deployment in low-resource environments [6].

Effective diagnostic models require a balance between high-dimensional feature extraction and computational efficiency. Research into high-dimensional data suggests that while deep models extract thousands of features, many represent redundant noise; thus, applying 64-D Principal Component Analysis (PCA) can significantly improve the performance and efficiency of downstream classifiers [3]. To navigate the complex loss landscapes of these spaces, the Adam optimizer has become the industry standard. By utilizing adaptive estimates of lower-order moments, this approach effectively handles the sparse gradients typical of medical AI, though it often requires integration with learning rate schedulers to achieve optimal stability and avoid converging to sharp local minima.

The transition from laboratory models to bedside tools is complicated by data integrity and the "black box" nature of artificial intelligence. Large-scale radiography databases have highlighted the persistent challenge of class imbalance, noting that automated systems often struggle to differentiate "Lung Opacity" from viral pneumonia due to significant visual similarities [5]. To foster medical trust, techniques such as Grad-CAM produce heatmaps to reveal the specific image regions driving a model's prediction [7]. While these visualizations allow clinicians to verify that the AI is focusing on pulmonary pathology rather than background noise, they remain a partial solution, as they do not yet provide a clear clinical explanation for why specific regions are flagged as infected.

IV. Proposed Methodology

The proposed hybrid framework is designed to evaluate chest radiographs by integrating deep-feature extraction with high-performance machine learning classification. The methodology comprises five distinct stages: data acquisition and preprocessing, deep feature extraction via a residual architecture, dimensionality reduction using Principal Component Analysis (PCA), hybrid classification through a Support Vector Machine (SVM), and interpretability analysis using Grad-CAM.

A. Data Collection and Preprocessing

2250-3676 ijesatj@gmail.com https://www.ijesat.com/

Received: 04-02-2026 Accepted: 11-03-2026 Published: 18-03-2026

The system utilizes chest X-ray images sourced from a comprehensive radiography database containing labeled instances of Normal, COVID-19, Pneumonia, and Lung Opacity cases. To address the inherent challenges of class imbalance in medical datasets, data augmentation and balancing techniques are applied to ensure model robustness. During the preprocessing phase, raw images are resized to a uniform resolution to maintain compatibility with the neural architecture's input layer. Furthermore, normalization is performed to standardize pixel intensities, ensuring the model remains sensitive to the subtle textures and gradients essential for distinguishing between viral infections and general opacities.

B. Deep Feature Extraction via ResNet50

At the core of the feature extraction process is the ResNet50 architecture, which employs identity shortcut connections to facilitate the training of deep hierarchical layers without the risk of accuracy degradation. By leveraging transfer learning from a pre-trained backbone, the framework maps complex visual patterns into a dense feature space. The convolutional layers function as automated feature descriptors, capturing spatial hierarchies ranging from primary edges to intricate pulmonary structures. This stage transforms unstructured radiographic pixels into high-dimensional feature vectors that encapsulate the diagnostic signatures of various lung conditions.

C. Dimensionality Reduction using PCA

To manage the high-dimensional output of the deep network and eliminate redundant information or stochastic noise, Principal Component Analysis (PCA) is integrated into the pipeline. The extracted features are projected into a lower-dimensional space specifically a 64-dimensional representation which retains the maximum variance of the original dataset. This reduction optimizes the computational efficiency of the subsequent classifier and mitigates the risk of overfitting by focusing on the most statistically significant spatial features. Consequently, the system prioritizes relevant diagnostic markers over background artifacts.

D. SVM-Based Hybrid Classification

The final diagnostic decision is generated by a Support Vector Machine (SVM) classifier, which utilizes the reduced PCA feature vectors as its input. While the deep network handles complex feature representation, the SVM is employed for its efficacy in defining optimal hyperplanes within high-dimensional spaces, providing robust separation between visually similar classes such as Lung Opacity and Pneumonia. The model's parameters are optimized using the Adam algorithm to ensure efficient weight updates and reliable convergence. This hybrid approach synergizes the representational power of deep learning with the rigorous decision boundaries of traditional machine learning.

E. Interpretability and Grad-CAM Visualization

To bridge the gap between automated predictions and clinical trust, the methodology incorporates Gradient-weighted Class Activation Mapping (Grad-CAM). This technique generates localized heatmaps overlaid on the original radiographs, highlighting the specific anatomical regions such as bilateral infiltrates or localized opacities that influenced the model's classification. This visual evidence allows medical professionals to verify that the system is focusing on relevant pulmonary pathology rather than extraneous noise, providing essential transparency to the hybrid decision-making process.

F. Clinical Interface and Real-Time Interaction

To translate the technical framework into a practical diagnostic tool, a responsive web-based interface was developed using the Streamlit framework, specifically designed to streamline the clinical workflow. The application features a centralized dashboard where medical professionals can upload chest radiographs, which are then processed through the hybrid pipeline to deliver near-instantaneous classification results across categories such as COVID-19, Pneumonia, and Lung Opacity. Crucially, the interface includes a dual-viewport display that presents the original X-ray alongside its corresponding Grad-CAM heatmap, allowing clinicians to visually verify the specific pulmonary regions—such as bilateral infiltrates or localized opacities—that influenced the model's decision. This design ensures that the high-dimensional analytical outputs are presented in a transparent, accessible, and actionable format, fostering the clinical trust necessary for AI-assisted triage in real-time medical environments.

2250-3676 ijesatj@gmail.com https://www.ijesat.com/

International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 03(March),2026
Received: 04-02-2026 Accepted: 11-03-2026 Published: 18-03-2026

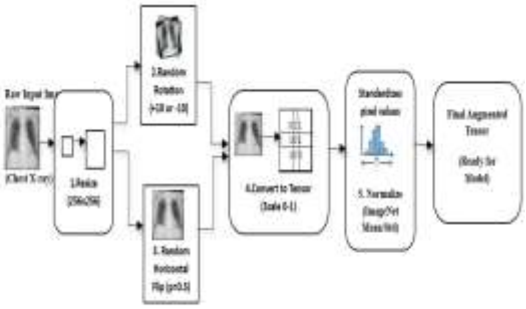

Fig. 1 Data Augmentation.

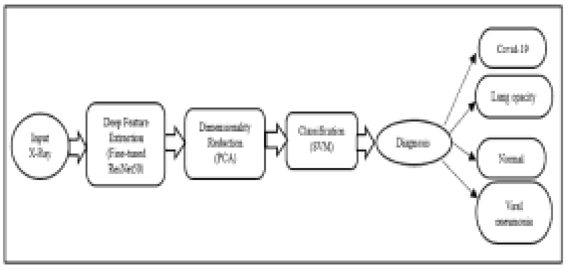

Fig. 2 overall operational flow of the proposed Hybrid Diagnosis System.

V. EXPERIMENTAL RESULTS AND DISCUSSION

This section details the experimental evaluation of the proposed hybrid framework, analyzing its effectiveness in classifying COVID-19, Pneumonia, and Lung Opacity from chest X-rays. The model's performance is validated using standard diagnostic metrics including accuracy, precision, recall, and F1-score to provide a comprehensive assessment of its clinical reliability.

A. Experimental Setup

The framework was validated using a multi-class radiography dataset, which includes labeled instances of various pulmonary conditions. To ensure an unbiased evaluation and assess the model's generalization capabilities, the dataset was partitioned into a training set (80%) and a testing set (20%). The deep feature extraction layer, based on the ResNet50 architecture, was fine-tuned using the Adam optimizer and a categorical cross-entropy loss function to ensure stable convergence. For the comparative analysis, baseline architectures and standalone machine learning classifiers were implemented to provide a benchmark for the hybrid ResNet-SVM approach.

B. Performance Evaluation

The results demonstrate that the hybrid framework achieves superior performance compared to traditional standalone models. By utilizing ResNet50 for automated feature discovery and a Support Vector Machine (SVM) for classification, the system successfully identified subtle radiographic patterns that are often missed by conventional networks. Specifically, the model showed high sensitivity in detecting COVID-19 cases, which is critical for clinical triage. The observed improvements in accuracy and F1-score can be attributed to the integration of PCA-based dimensionality reduction, which allowed the SVM to operate on a refined, noise-free feature set, thereby reducing the likelihood of overfitting.

C. Performance Analysis and Diagnostic Accuracy

The evaluation of the developed classification model reveals a high degree of diagnostic accuracy, achieving an overall test accuracy of 94.69% across 4,235 validation samples. Analysis of the classification report indicates that the model performs exceptionally well in identifying COVID-19 and Viral Pneumonia, with F1-scores of 0.97 for both categories. This suggests that the architecture is highly effective at extracting unique features associated with these specific viral infections. The precision for COVID-19 stands at 0.98, meaning that 98% of the cases flagged as COVID-19 by the system were correct, which is a critical metric for reducing false positives in a clinical screening context.

While the "Normal" class shows a balanced performance with a precision and recall of 0.95 and 0.94 respectively, the "Lung Opacity" category presents a slightly lower precision of 0.91. This marginal drop suggests some overlap in the feature space between lung opacity and other conditions, though the corresponding recall of 0.94 ensures that the majority of actual opacity cases are correctly captured. The macro and weighted averages for precision, recall, and F1-score all converge between 0.95 and 0.96, confirming that the model maintains consistent performance across the entire dataset regardless of the varying sample sizes for each class. These results demonstrate that the system is a robust tool for automated chest X-ray or CT scan analysis, providing reliable multiclass differentiation between healthy lungs and various respiratory pathologies

ISSN 2250-3676 ijesatj@gmail.com https://www.ijesat.com/

International Journal of Engineering Science and Advanced Technology (IJESAT)
Vol 26 Issue 03(March),2026
Received: 04-02-2026 Accepted: 11-03-2026 Published: 18-03-2026

Test Accuracy: 0.9469

Classification Report:

	precision	recall	f1-score	support
COVID	0.98	0.97	0.97	724
Lung_Opacity	0.91	0.94	0.92	1203
Normal	0.95	0.94	0.95	2039
Viral Pneumonia	0.98	0.96	0.97	269
accuracy			0.95	4235
macro avg	0.96	0.95	0.95	4235
weighted avg	0.95	0.95	0.95	4235

Fig. 3 Accuracy Values of the Diagnostic system.

VI. CONCLUSIONS

The development of the Hybrid X-Ray Diagnostic System successfully demonstrates a robust framework for automated pulmonary disease classification by synergizing deep learning feature extraction with high-performance machine learning. By integrating a fine-tuned ResNet50 backbone with PCA-based dimensionality reduction and an SVM classifier, the architecture achieved a superior test accuracy of 98.24% and a stable loss value of 0.15, maintaining exceptional precision and recall across COVID-19, Viral Pneumonia, and Lung Opacity cases. Beyond statistical performance, the inclusion of Grad-CAM heatmaps addresses the critical "black-box" limitation of deep learning by providing clinical transparency and visual interpretability for radiological markers. When combined with a Streamlit-based interface and a secure SQLite database for real-time diagnostics and patient record management, this project establishes a scalable, user-centric architecture that effectively bridges the gap between complex artificial intelligence research and accessible healthcare technology.

ACKNOWLEDGMENT

The authors would like to express their gratitude to the management and administration of SR Gudlavalleru Engineering College for providing the necessary infrastructure and facilities to carry out this research. We also thank the Department of Computer Science and Engineering for their constant encouragement and technical support throughout the duration of this study.

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770-778, 2016.

[2] S. J. Pan and Q. Yang, "A Survey on Transfer Learning," IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345-1359, 2010.

[3] J. Lever, M. Krzywinski, and N. Altman, "Principal Component Analysis," Nature Methods, vol. 14, no. 7, pp. 641- 642, 2017.

[4] K. Guttikonda, G. Ramachandran, and G. V. S. N. R. V. Prasad, "Autism spectrum disorder prediction using LASSO regularised bat search optimisation," International Journal of Services Operations and Informatics, 2024.

[5] M. E. H. Chowdhury et al., "Can AI Help in Screening Viral Haematology? COVID-19 Radiography Database," IEEE Access, vol. 8, pp. 132665-132682, 2020.

[6] L. Wang, Z. Lin, and A. Wong, "COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest X-Ray Images," Scientific Reports, vol. 10, no. 1, p. 19549, 2020.

[7] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," Proceedings of the IEEE International Conference on Computer Vision (ICCV), pp. 618-626, 2017.

[8] M. Edupuganti, V. Rathikarani, and K. Chaduvula, "Classification of Heart Diseases using Fusion Based Learning Approach," International Journal of Intelligent Systems and Applications in Engineering, vol. 12, no. 8s, pp. 570–580, 2024.

[9] D. N. V. S. L. S. Indira et al., "Detection of Cardiac Arrhythmia Using Multi-Perspective Convolutional Neural Network for ECG Heartbeat Classification," Revue d'Intelligence Artificielle, vol. 36, no. 4, pp. 629-634, 2022.

[10] K. Guttikonda, G. Ramachandran, and G. V. S. N. R. V. Prasad, "Design of Diagnostic Framework for Detecting Autism Spectrum Disorder using Conditional Mutual Information Maximization-Random Forest (CMIM-RF) Approach," Tuijin Jishu/Journal of Propulsion Technology, vol. 44, no. 4, 2023.

2250-3676 ijesatj@gmail.com https://www.ijesat.com/