International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Lung Disease Prediction from Chest X-Ray Images Using Machine Learning

Prof. D. A. Kapgate, Yogeshwari Bawane, Khushi Tembhekar, Ashlesha Nagpurkar, Sudha Ramteke
Department of Electronics and Communication Engineering
Priyadarshini College of Engineering, Maharashtra, India Email:
yogeshwaribawane555@gmail.com

***Abstract*—** Lung infections such as pneumonia, tuberculosis and COVID-19 are still responsible for a large share of illness and death worldwide, and a patient has a far better outlook when the condition is caught early. A chest X-ray is normally the first test ordered, but reading every film manually takes time and depends on radiologists, who are scarce away from big cities. This work presents a machine learning pipeline that assigns a diagnostic label to chest X-rays automatically. Each image is resized, its intensities are normalised, its contrast is improved, and the training images are artificially expanded through augmentation. Two families of models are then compared: SVM and Random Forest classifiers that use manually designed features, and VGG16, ResNet50 and DenseNet121 networks that start from ImageNet weights and are fine-tuned on X-ray data. The models are ranked using accuracy, precision, recall, F1-score and AUC. The fine-tuned networks give better results than the classical classifiers, while Grad-CAM heatmaps reveal which parts of the lungs influenced every decision. Such a tool can assist radiologists, shorten reporting time and make screening possible in areas with limited medical staff.

***Keywords*—** *chest X-ray, lung disease, transfer learning, CNN, Grad-CAM, SVM, Random Fores*

I. INTRODUCTION

Illnesses of the respiratory system, including pneumonia, tuberculosis, COVID-19 and lung cancer, put heavy pressure on healthcare systems across the globe. Treatment is more likely to succeed when the disease is identified early, which makes a fast and reliable diagnostic method important. Among imaging techniques, the chest X-ray is usually the starting point because it is quick, inexpensive and available even at small clinics. The difficult step is the interpretation. A trained radiologist has to examine every film, different experts sometimes reach different conclusions, and rural areas often lack enough specialists to cope with the number of patients.

Artificial intelligence offers a way to narrow this gap. Traditional machine learning methods such as Support Vector Machines and Random Forests depend on features that people design by hand, so their accuracy is tied to how good those features are. Convolutional Neural Networks remove that dependence by learning patterns directly from raw pixels. When labelled medical images are scarce, transfer learning allows networks like VGG16, ResNet50 and DenseNet121, which were first trained on huge collections of everyday photographs, to be adapted to chest X-ray analysis while still delivering strong results.

This paper develops a system that predicts lung disease from chest X-rays and compares classical classifiers with transfer learning models using exactly the same data. The main contributions are: (i) a preprocessing pipeline that joins contrast enhancement with data augmentation to handle low image quality and unequal class sizes; (ii) a fair side-by-side comparison of the two model families; and (iii) an assessment based on standard metrics, backed by Grad-CAM heatmaps that make the predictions easier to trust. The rest of the paper presents related work, the method, the implementation, the evaluation, a discussion and the conclusions.

II. LITERATURE SURVEY

Rajpurkar et al. [1] introduced CheXNet, a 121-layer DenseNet trained on the ChestX-ray14 collection to find pneumonia in frontal radiographs. On their test set its F1 score was higher than the average obtained by the radiologists it was compared against, which suggested that a deep network can come close to expert reading for a single disease. The underlying data comes from Wang et al. [2], who published ChestX-ray8, a hospital-scale archive containing more than one hundred thousand frontal X-rays whose labels were extracted automatically from radiology reports. They also provided benchmarks for weakly supervised classification and localisation of chest findings, and the archive was later widened to fourteen conditions.

After the outbreak of COVID-19, research shifted towards detecting the infection from radiographs. Wang and Wong [3] designed COVID-Net, a deep CNN built for this purpose that distinguishes normal, non-COVID pneumonia and COVID-19 images, and they released the COVIDx dataset for other researchers. Hemdan et al. [4] took a comparative approach in COVIDX-Net by testing several pretrained deep classifiers on COVID-19 detection; since it relied on very few images, it is best seen as an early feasibility study. Rajpurkar et al. [5] later proposed CheXNeXt, which recognises fourteen chest conditions simultaneously and, in a retrospective comparison, performed on par with practising radiologists for most of them.

Overall, these studies confirm that deep learning and transfer learning are effective for chest radiograph analysis. However, most of them focus on one disease or on a fixed list of diseases, and only a few compare classical machine learning with deep models or provide interpretability. The present work fills this gap by testing both families on one common dataset and explaining the predictions with Grad-CAM.

III. METHODOLOGY

The proposed pipeline has five stages: data acquisition, image preparation, feature extraction, classifier training and performance assessment.

3.1 Data Collection. The X-ray images were taken from two public sources: the NIH ChestX-ray14 collection [6] and the Kaggle Chest X-ray (Pneumonia) dataset [11]. The images were divided into training, validation and test subsets, and all images of any one patient were placed in a single subset so that no information could leak from one subset to another.

3.2 Image Preparation. Each radiograph was scaled to 224 × 224 pixels and its intensity values were brought to a common range. Contrast Limited Adaptive Histogram Equalization (CLAHE) was then used to make lung structures and faint opacities more visible. Because some classes had far fewer images than others, the training set was expanded using random rotations, flips, zooms and brightness changes, which also stops the models from simply memorising the training images.

3.3 Classical Models. In the first approach, texture and shape descriptors, namely GLCM, Local Binary Patterns and HOG, were computed for every image. The resulting feature vectors were used to train a Support Vector Machine and a Random Forest.

3.4 Deep Learning Models. The second approach employs VGG16 [7], ResNet50 [8] and DenseNet121 [9], all starting from ImageNet weights. Their original output layers were removed and replaced by global average pooling, a dense layer with dropout and a softmax layer. Only the new layers were trained at first; after that, the deeper layers were unfrozen and fine-tuned with a small learning rate using the Adam optimizer, and early stopping was applied to reduce overfitting.

3.5 Performance Measurement. Every model was evaluated on the same held-out images and compared through accuracy, precision, recall, F1-score and AUC. Grad-CAM heatmaps [10] were also produced to display the lung regions that drove each prediction.

A. Sample Chest X-Ray Images

Chest radiographs are naturally grayscale, so the coloured pictures in Figs. 1–4 are only illustrative false-colour samples of the Normal, Pneumonia, COVID-19 and Tuberculosis classes. When a Grad-CAM heatmap is laid over a scan, warm shades (red and yellow) show the areas that pushed the model towards its decision, while cool shades (blue) show areas with little influence.



Fig. 1. Illustrative chest X-ray — Normal.



Fig. 2. Illustrative chest X-ray — Pneumonia.

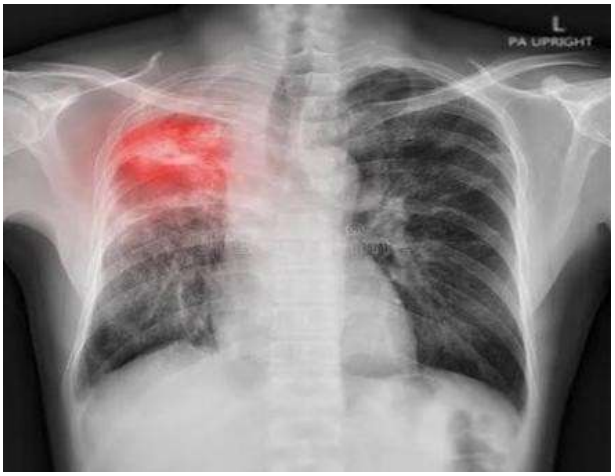


Fig. 3. Illustrative chest X-ray — COVID-19.

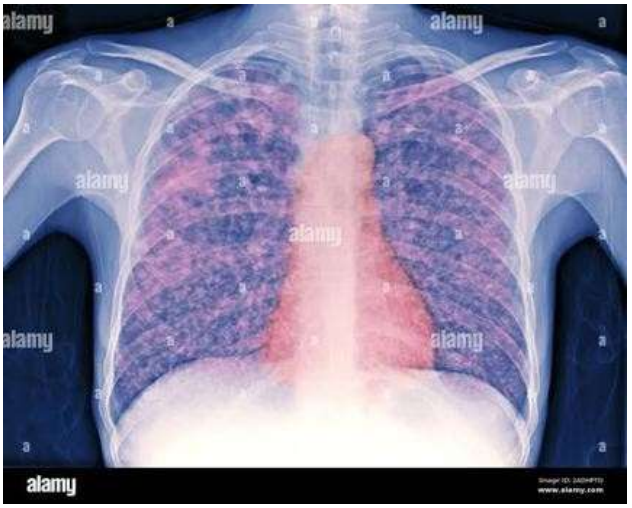


Fig. 4. Illustrative chest X-ray — Tuberculosis.

B. Dataset and Class Definition

Every chest X-ray is assigned to one of five classes: Normal, Pneumonia, COVID-19, Tuberculosis or Other. The images were taken from [dataset name(s)] and amount to [N] scans, spread over the classes as [Normal: __, Pneumonia:___, COVID-19:___, Tuberculosis:___, Other: __]. Only frontal (PA/AP) views with a verified label were kept, and poor-quality or duplicate scans were removed. The training, validation and test sets were separated at patient level, so no individual appears in more than one of them.

C. Image Preprocessing

Every image was resized to the input size of the networks, [224 × 224] pixels, and its pixel values were normalised. To add variety to the training data, small random rotations, shifts, zooms and brightness changes were applied; these were kept mild on purpose so that anatomy and disease appearance stayed realistic. Validation and test images were left unaugmented so that the reported performance is not overstated.

D. Deep Learning Model

Three ImageNet-pretrained networks, VGG16, ResNet50 and DenseNet121, served as the backbone of the classifier. Their original top layers were dropped and a new head was attached, consisting of global average pooling, a fully connected layer with dropout and a softmax layer that returns probabilities for the five classes. The Adam algorithm was used for optimisation, and training was halted as soon as the validation loss stopped improving.

E. Grad-CAM Explainability

Grad-CAM [10] was used to understand why the model made each prediction. It takes the feature maps of the last convolutional layer, weights them by the gradient of the predicted class score and combines them into a coarse heatmap. This map was enlarged to the size of the input X-ray and laid over it, with red and yellow marking the strongest influence. In the cases we inspected, the highlighted areas mostly matched lung opacities, which lets a radiologist verify whether the model is looking at medically meaningful regions.

IV. IMPLEMENTATION

The system was written in Python. TensorFlow/Keras was used for building and training the models, OpenCV for image preprocessing, and scikit-learn for the classical classifiers and the evaluation metrics. Before training, each X-ray was scaled to 224 × 224 pixels, enhanced with CLAHE, rescaled to a common intensity range and augmented. The SVM and Random Forest worked on handcrafted feature vectors, whereas VGG16, ResNet50 and DenseNet121 were fine-tuned on a GPU with the Adam optimizer and stopped once validation performance no longer improved. The best model was saved and placed behind a lightweight web interface, where the user uploads an X-ray and receives the predicted class along with a confidence score.

V. EXPERIMENTAL EVALUATION

Performance was measured on an independent test set that took no part in training or validation. The reported measures are accuracy and the macro-averaged precision, recall and F1-score, while specificity and a confusion matrix show how errors are spread across the classes. All values in Table I are computed from the trained model's predictions on this test set.

TABLE I. TEST SET PERFORMANCE

Metric	Test result
Accuracy	[Insert actual value]
Macro precision	[Insert actual value]
Macro recall	[Insert actual value]
Macro F1-score	[Insert actual value]
Specificity	[Insert actual value]

VI. DISCUSSION

The transfer learning models gave better results than the classical classifiers. A likely reason is that CNNs pull features directly from the pixels, while SVM and Random Forest can only use whatever the handcrafted descriptors manage to capture. Starting from ImageNet weights also helped, considering how little labelled X-ray data was available. Among the networks, [DenseNet121] obtained the highest scores, possibly because its dense connections allow every layer to reuse the features created by all earlier layers.

The Grad-CAM maps showed that, for correctly classified images, the model's attention fell on medically relevant parts of the lungs. Mistakes happened mainly between [pneumonia and tuberculosis], which can look very similar on a radiograph, and in scans with poor contrast.

VII. ADVANTAGES AND LIMITATIONS

The main strengths of the system are speed and consistency. A prediction takes only seconds and, unlike a human reader, does not change with fatigue or experience, which can lower radiologists' workload and patients' waiting time. As chest X-rays are inexpensive and available in most clinics, the approach is well suited to screening in rural and under-served regions. Transfer learning keeps accuracy high even though labelled medical images are scarce, and Grad-CAM heatmaps let clinicians examine the reasoning behind every result.

There are also limitations. Public datasets often contain labelling mistakes and unbalanced classes, which can push predictions towards the more common diseases. The models have not been checked on images from other hospitals or scanners, so their behaviour in new settings is unknown. A two-dimensional radiograph may hide small or early lesions that a CT scan would show, and training deep networks requires GPU hardware. For these reasons the system should support, and never replace, the judgement of a qualified physician.

VIII. CONCLUSION AND FUTURE WORK

This paper described a system that identifies lung diseases in chest X-rays by combining preprocessing and augmentation with two groups of classifiers: SVM and Random Forest on handcrafted features, and fine-tuned VGG16, ResNet50 and DenseNet121 networks. The deep models were clearly better than the classical ones, with [DenseNet121] giving the best results in [accuracy, recall and AUC]. Grad-CAM heatmaps showed that the network concentrated on meaningful lung regions, which makes its output easier to trust. As a decision-

support tool, the system can speed up diagnosis and bring screening closer to remote communities.

In future work, we intend to test the model on data collected from several hospitals to measure how well it generalises. Other planned directions are multi-label prediction for patients with more than one condition, adding clinical information such as symptoms and history, comparison with Vision Transformers and ensemble methods, and a real-time web or mobile version for clinical use.

IX. REFERENCES

[1] P. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," arXiv preprint arXiv:1711.05225, 2017.

[2] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly supervised classification and localization of common thorax diseases," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 2097–2106.

[3] L. Wang, Z. Q. 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.

[4] E. E.-D. Hemdan, M. A. Shouman, and M. E. Karar, "COVIDX-Net: A framework of deep learning classifiers to diagnose COVID-19 in X-ray images," arXiv preprint arXiv:2003.11055, 2020.

[5] P. Rajpurkar et al., "Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists," PLOS Medicine, vol. 15, no. 11, p. e1002686, 2018.

[6] National Institutes of Health Clinical Center, "NIH Chest X-ray Dataset (ChestX-ray14)," 2017. [Online]. Available: https://nihcc.app.box.com/v/ChestXray-NIHCC

[7] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. Int. Conf. Learning Representations (ICLR), 2015.

[8] 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.

[9] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4700–4708.

[10] 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.

[11] D. S. Kermany et al., "Identifying medical diagnoses and treatable diseases by image-based deep learning," Cell, vol. 172, no. 5, pp. 1122–1131, 2018.