



# **iJRASET**

International Journal For Research in  
Applied Science and Engineering Technology



---

# **INTERNATIONAL JOURNAL FOR RESEARCH**

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

---

**Volume: 12    Issue: IV    Month of publication: April 2024**

**DOI: <https://doi.org/10.22214/ijraset.2024.61172>**

**[www.ijraset.com](http://www.ijraset.com)**

**Call:  08813907089**

**E-mail ID: [ijraset@gmail.com](mailto:ijraset@gmail.com)**

# Prediction of Chest X-Ray Abnormalities using Ensemble Learning

Sunil Wankhade<sup>1</sup>, Shubham Mahale<sup>2</sup>, Jeet Mane<sup>3</sup>, Nupoor Shetye<sup>4</sup>, Tejasvi Mahalunge<sup>5</sup>

<sup>1</sup>Head of the Department, <sup>2, 3, 4</sup>Student, Department of Information Technology, Rajiv Gandhi Institute of Technology

**Abstract:** Chest X-rays are a cornerstone of medical imaging, enabling rapid and cost-effective diagnosis of numerous chest pathologies. However, limitations in radiologist availability and long wait times can significantly delay diagnoses, potentially impacting patient outcomes. This study explores the potential of deep learning to address these challenges. Deep learning algorithms can analyse medical images and identify patterns associated with various diseases. In this study, various deep learning models were trained and evaluated, including DenseNet121, EfficientNetB1, Xception, and an ensemble model combining EfficientNetB1 and Xception, using a random sample of 5,606 chest X-ray images from the National Institutes of Health (NIH) chest X-ray dataset. The ensemble model achieved the most promising results, demonstrating an accuracy of 72.66% and an AUC-ROC (Area Under the Receiver Operating Characteristic Curve) of 76.05%.

**Keywords:** Chest-X-Rays, Convolutional Neural Network, Deep Learning, Medical Image Processing, Ensemble Learning, Efficient Net, Xception

## I. INTRODUCTION

Chest X-rays remain an indispensable tool in the medical field. Their rapid acquisition, low cost, and ability to visualise a widerange of lung and chest pathologies make them a cornerstone of initial patient evaluation [1]. However, accurate diagnosis hinges on the expertise of radiologists, and limitations in their availability can lead to significant delays. These delays can have a cascading effect, potentially impacting patient outcomes and leading to increased healthcare costs [2].

Deep learning (DL), a subfield of artificial intelligence, has emerged as a promising solution to address these challenges. DL algorithms excel at analysing vast amounts of medical image data and identifying subtle patterns associated with various diseases. This ability to extract meaningful information from complex medical images has the potential to revolutionize chest X-ray analysis[3].

This study investigates the potential of DL for chest X-ray analysis. We aim to leverage the power of DL models to develop a system that can assist radiologists in interpreting chest X-rays. Our goal is to achieve faster and more efficient diagnoses, ultimately improving patient care. This research aligns with the growing body of work exploring DL for chest X-ray analysis.

In recent years, numerous studies have explored the application of deep learning models for chest X-ray analysis. A widely used chest X-ray application was developed by Wang et al. Their contribution is a large dataset named ChestX-Ray8, containing 108,948 frontal X-ray images from 32,717 patients. Leveraging natural language processing and convolutional neural networks, they achieved good results in detecting 14 different chest diseases from radiology reports [4]. Jeremy Irvin et al (2017) developed a deep learning model for automated pneumonia detection from X-ray images, showcasing the potential of these techniques in improving diagnostic accuracy. The CheXNet model utilises a 121-layer DenseNet architecture pre-trained on the ImageNet dataset. In CheXNet, the final fully connected layer of the pre-trained DenseNet is replaced with a single output layer with a sigmoid activation function, enabling the model to predict the probability of various pathologies being present in a chest X-ray [5].

The availability of benchmark datasets, such as the NIH chest X-ray dataset, has facilitated progress in this field. This dataset contains many chest X-ray images with associated labels for various pathologies, making it suitable for training and evaluating deep learning models.

In addition to single-model approaches, ensemble learning techniques have gained attention for improving predictive performance. Ensemble methods, such as bagging, boosting, and stacking, combine multiple models to enhance accuracy and robustness [6]. Despite the progress made in the field, there remains a need for robust models capable of accurately detecting multiple pathologies simultaneously.

This study aims to address this gap by proposing an ensemble deep learning approach using the NIH chest X-ray dataset. By combining robust deep learning models, EfficientNetB1 and Xception, we seek to improve the accuracy and reliability of chest X-ray analysis for diagnosing various chest-related pathologies.

## II. CONVOLUTIONAL NEURAL NETWORKS

Convolutional Neural Networks (CNNs) are a class of deep learning models specifically designed for processing structured grid-like data, such as images. CNNs employ a hierarchical architecture inspired by the visual cortex of animals. The basic building blocks of a CNN include convolutional layers, pooling layers, and fully connected layers.

In the architecture, convolutional layers apply a set of learnable filters to the input image, extracting features such as edges, textures, and patterns. Pooling layers then down sample the feature maps obtained from the convolutional layers, reducing spatial dimensions while preserving important information. This hierarchical feature extraction process is repeated multiple times, with each subsequent layer learning increasingly complex features. Finally, fully connected layers interpret the extracted features and produce the final output, such as class probabilities in image classification tasks. Overall, CNNs have revolutionized computer vision tasks by automatically learning hierarchical representations from raw input data, making them widely used in image recognition, object detection, and more.




Fig. 1 Architecture of Convolutional Neural Networks

### A. DenseNet121

DenseNet121 is a convolutional neural network architecture that utilizes dense convolutional connections. This approach is a core characteristic of the DenseNet architecture, and it differs from traditional convolutional neural networks (CNNs) by connecting each layer to all subsequent layers in the network. DenseNet121 refers to a specific variant of the DenseNet architecture that has approximately 121 layers. It follows the principle of dense convolutional blocks stacked together with intervening pooling layers to gradually reduce the size of feature maps and introduce spatial invariance. This architecture utilizes dense convolutional connections, where each layer receives information from all preceding layers. The network also incorporates global average pooling at the end to transform the remaining feature maps into a fixed-size vector before feeding it into the final classification layer. This approach is believed to improve feature propagation and potentially alleviate the vanishing gradient problem.

### B. EfficientNet

EfficientNet is a convolutional neural network architecture developed to achieve state-of-the-art performance with significantly fewer parameters compared to traditional models. It introduces a novel compound scaling method that uniformly scales the network's depth, width, and resolution, enabling better model efficiency and accuracy. By optimizing these scaling factors, EfficientNet achieves superior performance across various tasks while maintaining computational efficiency. The architecture of EfficientNet incorporates efficient building blocks, such as inverted residual blocks with squeeze-and-excitation, to efficiently capture complex patterns in the data. These blocks enable the model to extract essential features from the input data while minimizing computational overhead. Moreover, EfficientNet's compound scaling method dynamically adjusts the network's depth, width, and resolution to achieve optimal performance on different tasks and resource constraints. EfficientNet's design makes it particularly suitable for deployment on mobile and edge devices where computational resources are limited. Despite these constraints, EfficientNet maintains high accuracy in tasks like image classification, object detection, and segmentation. Its efficiency and effectiveness make it a preferred choice for various real-world applications, where achieving the right balance between performance and computational cost is crucial.

C. Xception

Xception builds upon the idea of efficient feature extraction. It utilizes depth-wise separable convolutions, a two-step process. First, a separate filter is applied to each colour channel in the image, capturing individual channel information. Then, a 1x1 convolution combines these features into a more compact representation. This approach reduces computational cost compared to standard convolutions, making Xception well-suited for resource-constrained environments. Xception employs these depth-wise separable convolutions in a specific architecture. It starts with an "entry flow" using regular convolutions for initial feature extraction. The core of the model lies in the "middle flow," where repeated blocks stack depth-wise separable convolutions with batch normalization and activation layers. The number of channels progressively increases, allowing the network to capture increasingly complex features. Finally, the "exit flow" further refines the features and prepares them for classification. This combination of efficient building blocks and a well-designed architecture allows Xception to achieve high accuracy while maintaining computational efficiency.

III. METHODOLOGY

A. Dataset Preparation and Preprocessing

Our study utilized a random sample of the National Institutes of Health (NIH) chest X-ray dataset from Kaggle, containing 5606 images. We cleaned the data, transforming labels into a format suitable for the model (one-hot encoding), and linked image paths to their corresponding labels.

To enhance the dataset's robustness, we employed data splitting (training, validation, testing) and augmentation techniques. Splitting ensures the model is trained on a representative sample while reserving data for evaluation. A common split ratio of 80% for training, 10% for validation, and 10% for testing was used. Data augmentation techniques were then applied to artificially expand the dataset size and improve model generalization. Data augmentation involves creating modified versions of existing images, such as applying random flips, rotations, or scaling. This process helps the model learn features that are robust to small variations in image appearance, ultimately reducing the risk of overfitting on the training data. Finally, normalization techniques were applied to improve model training stability. By implementing these data preparation steps, we aimed to create a high-quality dataset for training our deep learning models for chest X-ray analysis.

B. Classification of Chest X-rays

We have trained four deep learning models on the given dataset: DenseNet121, EfficientNetB1, Xception. Additionally, our study introduces an ensemble model EffXception consisting of EfficientNetB1 and Xception to utilize the improved performance provided by ensemble techniques and combine the strengths of EfficientNetB1 and Xception, leveraging their complementary features. The models leverage pre-trained convolutional neural networks (CNNs) on the massive ImageNet dataset for feature extraction, followed by fine-tuning for the specific task of multi-label classification of chest X-ray pathologies. Transfer learning allows us to leverage knowledge gained from pre-trained models on large datasets for new tasks, significantly improving training efficiency and potentially enhancing performance compared to training models from scratch.

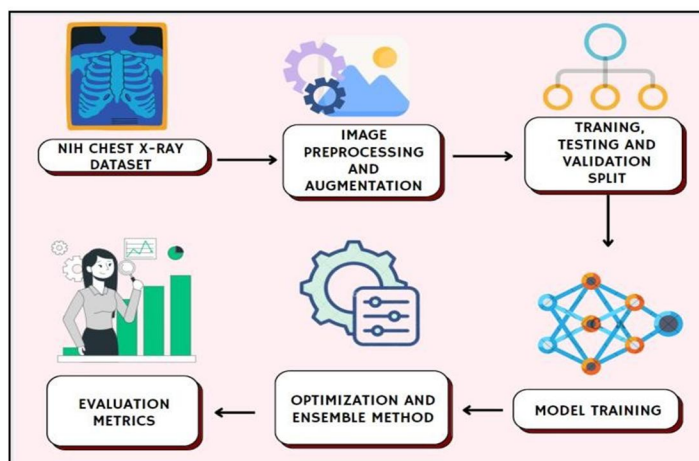


Fig. 2 Block Diagram of the System



Fig. 3 Flowchart of Chest X-Ray abnormality detection using EffXception ensemble model

The architecture of the models is as follows:

- 1) *Input Layer*: The model starts with an input layer that takes a chest X-ray image as input. The specific image size (height,width, and number of channels) may be defined based on the pre-trained model's requirements. These images were fed via a pandas dataframe using Tensorflow's inbuilt functions.
- 2) *Pre-trained Base Model*: The core of the architecture is a pre-trained convolutional neural network (CNN) loaded with weights learned on a large image dataset like ImageNet. This pre-trained model acts as a feature extractor, automatically identifying and encoding high-level visual features relevant to image classification tasks.
- 3) *Global Average Pooling (GAP)*: A Global Average Pooling (GAP) layer is applied after the pre-trained model. This layer reduces the spatial dimensions of the feature maps extracted by the CNN, summarizing the information into a single feature vector for each channel.
- 4) *Dense Layer*: A fully connected (dense) layer with a size equal to the number of target pathologies is added after the GAP layer. This layer transforms the reduced feature vector into a set of output logits (unnormalized probabilities) representing the likelihood of each pathology being present in the chest X-ray.
- 5) *Activation Function*: Finally, a sigmoid activation function is applied to the logits from the dense layer. The sigmoid function normalizes the output values between 0 and 1, representing the probability of each pathology.
- 6) *Loss Function*: The model is compiled with a custom loss function that incorporates class weights to address potential class imbalance in the dataset. The specific loss function is a variant of binary cross-entropy incorporating the calculated class weights.
- 7) *Optimizer*: An optimizer like Adam is used to update the model weights during training based on the calculated loss values.
- 8) *Metrics*: Metrics like accuracy, AUC-ROC, precision, recall, and F1 score are tracked to evaluate model performance during training and on unseen test data.

Ensemble learning is a powerful technique that leverages the predictions from multiple models to potentially achieve better performance compared to a single model.

There are several ensemble techniques including Averaging, Bagging, Boosting and Stacking. In our work, we employed an ensemble model: EffXception that combined the predictions from two of our pre-trained models: EfficientNet-B1 and Xception. Each chest X-ray image is fed forward separately through both the EfficientNetB1 and Xception models. These models independently process the image and generate individual probability predictions for each pathology. The corresponding probability predictions for each pathology, obtained from both EfficientNetB1 and Xception, are then averaged. The resulting averaged probabilities represent the ensemble model's final prediction for the likelihood of each pathology being present in the chest X-ray. These models have different architectures (EfficientNet-B1 focuses on efficiency while Xception is deeper). By combining them, we aimed to potentially leverage their complementary strengths and capture a more diverse set of features for improved prediction accuracy. Also, Ensemble learning can help reduce the overall prediction error by averaging out individual model errors. This can lead to a more robust prediction compared to relying solely on a single model.

#### IV. RESULTS AND DISCUSSION

In our study, we trained four models—DenseNet121, EfficientNetB1, Xception, and an ensemble model—on a random sample of Chest X-rays dataset for 20 epochs. Among these models, the ensemble model, which combines predictions from EfficientNetB1 and Xception, outperformed others in terms of accuracy, AUC-ROC, and recall. This demonstrates the effectiveness of leveraging multiple models to enhance predictive performance in chest X-ray abnormality detection. The ensemble model achieved a mean accuracy of 72.66%, mean AUC-ROC of 76.05%, and mean recall of 19.03%.

The superior performance of the ensemble model highlights the benefits of combining the strengths of different architectures. By leveraging the diverse features learned by EfficientNetB1 and Xception, the ensemble model provides a more robust and accurate prediction. This approach is particularly advantageous in medical imaging, where accurate diagnosis is crucial. Our findings underscore the potential impact of utilizing ensemble models for chest X-ray analysis, paving the way for more reliable and efficient diagnosis of chest-related pathologies, ultimately leading to improved patient outcomes.

Fig. 4 AUC ROC curve of EffXception ensemble model




Fig. 4 shows the AUC ROC curve of the ensemble model demonstrates its superior performance, showcasing higher true positive rates and lower false positive rates across various thresholds, validating its effectiveness in chest X-ray abnormality detection.

TABLE I  
Comparison of performance metrics of all models

Model	Accuracy	AUC-ROC	Recall
DenseNet121	0.574911	0.693157	0.141612
EfficientNetB1	0.700357	0.749863	0.174414
Xception	0.684814	0.723187	0.178519
EffXception (Ensemble)	0.726601	0.760543	0.190395

Table I compares the performance metrics across different models, including DenseNet121, EfficientNetB1, Xception, and the ensemble model (EffXception). Metrics such as accuracy, AUC-ROC, and recall are shown, indicating the effectiveness of each model in classifying chest X-ray abnormalities. The ensemble model (EffXception) demonstrates the highest accuracy, AUC-ROC, and recall among the models, indicating its superior performance in detecting abnormalities.




Fig. 5 Visualization of predicted vs actual label for chest x-ray images

Fig 5 displays a chest X-ray along with its ground truth label and the predicted values for all 14 labels. In this example, the ground truth label is "Atelectasis," and the predicted label with the highest value is also "Atelectasis," indicating that the model's prediction aligns well with the actual label. This suggests that the model can effectively predict abnormalities in chest X-rays.

## V. CONCLUSIONS

In this project, we delve into the critical domain of chest abnormality detection using chest X-ray images, with a unique approach that harnesses the power of multiple pretrained deep learning models. The ability to detect chest abnormalities with high precision has substantial implications for early diagnosis, treatment planning, and healthcare resource optimization. Our goal is to address the significant challenge of improving the accuracy and reliability of detecting chest abnormalities, which is vital for enhancing patient care and diagnostic accuracy in the healthcare sector.

By leveraging models like DenseNet121, EfficientNetB1, and Xception, the project showcases significant advancements in accurately identifying various chest pathologies from X-ray images. The ensemble technique further enhances predictive performance, underscoring the importance of leveraging diverse models for improved diagnosis. Moreover, addressing class imbalance through weighted loss functions ensures robustness in handling real-world datasets. Overall, this project not only contributes to the advancement of medical diagnostics but also highlights the potential of neural networks to revolutionize healthcare by providing efficient, accurate, and scalable solutions for diagnosing diseases from medical imaging data. Moving forward, continued research and development in this field will be essential for maximizing the potential impact of deep learning in healthcare and improving patient outcomes globally.



## REFERENCES

- [1] Feng, Y., Teh, H. S., & Cai, Y. "Deep learning for chest radiology: A review." *Current Radiology Reports*, vol. 7, no. 8, 2019, p. 24. [DOI: 10.1007/s40134-019-0333-9]
- [2] Fatihoglu, E., et al. "X-ray use in chest imaging in emergency department on the basis of cost and effectiveness." *Academic Radiology*, vol. 23, no. 10, 2016, pp. 1239-1245. [DOI: 10.1016/j.acra.2016.05.008]
- [3] Litjens, G., et al. "A survey on deep learning in medical image analysis." *Medical Image Analysis*, vol. 42, 2017, pp. 60-88. [DOI: 10.1016/j.media.2017.07.005]
- [4] Wang, X., et al. "ChestX-Ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases." *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017, pp. 3462-3471. [DOI: 10.1109/CVPR.2017.369]
- [5] Irvin, J., et al. "CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison." *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 4992-5000.
- [6] Caruana, R., et al. "Ensemble selection from libraries of models." *Proceedings of the Twenty-First International Conference on Machine Learning (ICML'04)*, ACM, 2004, pp. 18-23. [DOI: 10.1145/1015330.1015432]





10.22214/IJRASET



45.98



IMPACT FACTOR:  
7.129



IMPACT FACTOR:  
7.429



# INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24\*7 Support on Whatsapp)