VER
Vascular & Endovascular Review

Detection of Pneumonia Utilising Deep learning-based Feature Extraction

Ankur Jain[1], Deepankar Bharadwaj[2]

[1,2]IFTM University, Delhi Road Lodhipur Rajput Moradabad
ankur1101@gmail.com – ORCID ID: https://orcid.org/0000-0002-3246-7612
deepankar.bharadwaj@gmail.com - ORCID ID: https://orcid.org/0009-0006-0035-680X

ABSTRACT

Pneumonia is a potentially fatal infectious disease that impacts one or both lungs in humans, primarily caused by the bacterium Streptococcus pneumoniae. According to the World Health Organisation (WHO), pneumonia is responsible for one in three deaths in India. Chest X-rays utilised for pneumonia diagnosis require assessment by specialist radiologists. Consequently, creating an automated method for pneumonia detection will facilitate timely treatment of the condition, especially in distant regions. Convolutional Neural Networks (CNNs) have garnered significant attention for illness classification due to the efficacy of deep learning algorithms in analysing medical imagery. Moreover, characteristics acquired by pre-trained CNN models on extensive datasets are highly beneficial for image classification tasks. This study evaluates the efficacy of pre-trained CNN models employed as feature extractors in conjunction with various classifiers for the classification of abnormal and normal chest X-rays. We analytically ascertain the best CNN model for the objective. The statistical results obtained indicate that pretrained CNN models, when utilised with supervised classifier methods, are highly advantageous for analysing chest X-ray pictures, particularly for the detection of pneumonia.

KEYWORDS: Deep Convolutional Neural Networks, Support Vector Machines, Transfer Learning, Random Forest, Naive Bayes, K-Nearest Neighbours, Feature Extraction.

How to Cite: Ankur Jain, Deepankar Bharadwaj., (2025) Detection of Pneumonia Utilising Deep learning-based Feature Extraction, Vascular and Endovascular Review, Vol.8, No.12s, 34-42.

INTRODUCTION

In recent years, Computer Aided Designs (CAD) have emerged as a primary research subject in machine learning. The existing CAD systems have demonstrated efficacy in the medical field, particularly in the diagnosis of breast cancer, mammography, and lung nodules. In the application of Machine Learning (ML) techniques to medical imaging, significant features are of paramount importance. Consequently, the majority of earlier methods employed manually designed characteristics for the development of CAD systems using image analysis [1], [2], [3]. Nonetheless, the handcrafted features, constrained by task-specific restrictions, were insufficient in providing significant attributes. The utilisation of Deep Learning (DL) models, namely Convolutional Neural Networks (CNNs), has demonstrated their inherent capability to extract valuable features in picture classification problems [4], [5]. This feature-extraction procedure necessitates transfer learning techniques, wherein pre-trained CNN models acquire general features from large-scale datasets like as ImageNet, which are subsequently applied to the specific job. The availability of pre-trained CNN models such as AlexNet, VGGNet, Xception, ResNet, and DenseNet significantly facilitates the process of feature extraction. Within

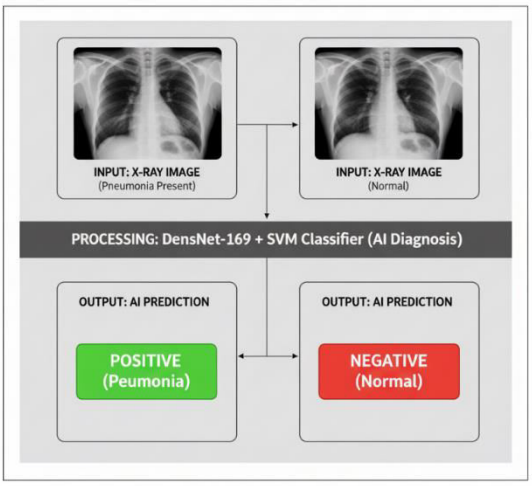

Figure. 1. An illustration of a Normal CXR (right) and a Pneumonia CXR (left) from the ChestX-ray14 dataset. The pathology in the left chest X-ray is not immediately differentiable from that in the right chest X-ray.

Furthermore, the classification employed with high-quality extracted features demonstrates enhanced efficacy in picture classification [6], [7], [8].

Chest screening subroutines primarily utilised for detecting lung nodules can also facilitate the diagnosis of other conditions, including pneumonia, effusion, and cardiomegaly. Pneumonia is a lethal infectious disease that affects millions, particularly individuals over 65 years old with chronic conditions such as asthma or diabetes. Chest X-rays are regarded as the most effective way for diagnosing pneumonia, as they determine the extent and location of the infected area in the lungs. Nonetheless, analysing chest radiographs is not a trivial undertaking for radiotherapists. Pneumonia in chest X-ray pictures may present as blurry and can be misinterpreted as other conditions. The assessment of chest X-rays in cases of pneumonia can be deceptive, as various other conditions, such as congestive heart failure and pulmonary scarring, may present similarly to pneumonia. The primary cause of the misclassification of the X-ray pictures in the dataset is this. Consequently, the task is arduous and the creation of an algorithm for identifying thoracic disorders such as pneumonia might enhance the accessibility of clinical services in remote regions. This study assessed the efficacy of various pre-trained CNN model variations in conjunction with different classifiers for the classification of abnormal and normal chest X-rays. This study makes significant contributions: (a) a comparative analytical examination of various pre-trained CNN models as feature extractors for chest X-ray analysis, (b) the presentation of these models alongside different classifiers to identify the most suitable classifier in the classification domain, and (c) the assessment of the optimal pre-trained CNN model with hyperparameter tuning of the best-performing classifier to enhance performance further. This paper's format is outlined as follows: Section 2 presents a review of pertinent research conducted in the same domain. Section 3 provides a comprehensive summary of all pertinent details regarding the utilised dataset. Section 4 presents a detailed description of the employed approach, which is categorised into many steps. Section 5 delineates the experimental configuration utilised for the trials conducted on several pre-trained CNN model variants, as well as the outcomes derived from the application of distinct classifiers. Section 6 encompasses the results and debates pertaining to the final AUC scores achieved.

RELATED WORK

Recently, the investigation of machine learning (ML) algorithms for the detection of thoracic disorders has garnered attention in the field of medical picture classification. Lakhani and Sundaram (2017) [9] proposed a method for identifying pulmonary tuberculosis utilising the architectures of two distinct Deep Convolutional Neural Networks, AlexNet and GoogleNet. The classification of pulmonary nodules for lung cancer diagnosis, proposed by Huang et al. [10], also utilised deep learning techniques. Islam et al. [11] presented the evaluation of various Convolutional Neural Network (CNN) variations for abnormality identification in chest X-rays utilising the publically accessible OpenI dataset [12]. To enhance the investigation of machine learning in chest screening, Wang et al. (2017)[13] published an extensive dataset of frontal chest X-rays. Recently, Pranav Rajpurkar, Jeremy Irvin, and others (2017) [[14] investigated this dataset for pneumonia detection, surpassing radiologists' performance. They named their model ChexNet, which used a DenseNet-121 architecture to identify all 14 diseases from a total of 112,200 images in the dataset. Following the CheXNet model, Benjamin Antin et al. (2017) [15]developed a logistic regression model for pneumonia detection utilising the same dataset. Pulkit Kumar and Monika Grewal (2017) [16] utilised cascading convolutional networks to advance their research on multilabel classification of thoracic illnesses. Zhe Li (2018) A convolutional network model was recently proposed for illness identification and localisation.

Dataset Overview

The dataset employed is ChestX-ray14, published by Wang et al. (2017) [13], and is also accessible on the Kaggle [13], [17] platform, comprising 112,120 frontal chest X-ray pictures from 30,085 patients. Each radiography image in the dataset is annotated with one or more of 14 distinct thoracic illnesses. These labels were derived by Natural Language Processing (NLP) by text-mining disease classification from the corresponding radiological reports and are anticipated to exceed 90% accuracy. In this study, we adopt previous methodologies [14] and consider the labels as the definitive reference for pneumonia detection. Before the publication of this dataset, the most extensive publicly accessible dataset of chest radiographs was Openi comprising approximately 4,143 X-ray pictures.

All radiographic pictures in the dataset possess a resolution of 1024 by 1024. Among the 112,120 photos, 1,431 are identified as labelled with pneumonia. To equilibrate the dataset for binary classification, 1431 normal X-ray images (designated as 'No Findings') have been extracted from the dataset. The final dataset utilised for the classification job comprises a subset of the original dataset, consisting of 1431 positive image samples labelled 'Pneumonia' and 1431 negative image samples labelled 'No Findings.' The dataset was subsequently partitioned into two segments, from which 573 photos were randomly chosen for testing from the complete final dataset. The photos were reduced from a resolution of 1024 by 1024 to 224 by 224 prior to being entered into the network.

METHODOLOGY

This section provides a comprehensive discussion of the employed methodology. The pneumonia detection system utilising the 'Densely Connected Convolutional Neural Network' (DenseNet-169) is illustrated in Figure 2. The architecture of the suggested model is divided into three distinct stages: the preprocessing stage, the feature-extraction stage, and the classification step.

The Preprocessing Phase

The principal objective of employing Convolutional Neural Networks in image classification tasks is to diminish the computational complexity of the model, which tends to escalate with picture inputs. The original 3-channel pictures were downscaled from 1024×1024 to 224×224 pixels to mitigate computational demands and enhance processing speed. All subsequent procedures have been applied to these reduced photos.

The Feature Extraction Phase

Despite utilising various pre-trained CNN models for feature extraction, the statistical findings indicated that DenseNet-169 was the most effective model for this stage. This step addresses the architecture of the DenseNet-169 model and its role in feature extraction.

The architecture of DenseNet-169: Deep Convolutional Neural Networks (DCNNs) have emerged as the most effective frameworks for image recognition because to the unique characteristics of their convolutional and pooling layers. However, as the network deepens, the input information or gradient diminishes significantly by the time it reaches the last layer of the network. DenseNets mitigate the issue of gradient vanishing by establishing direct connections among all layers with uniform feature sizes. The primary reason for employing DenseNet architecture as a feature extractor is that a deeper network yields more generalised features. The pre-trained 169-layer Densely Connected Convolutional Neural Network.

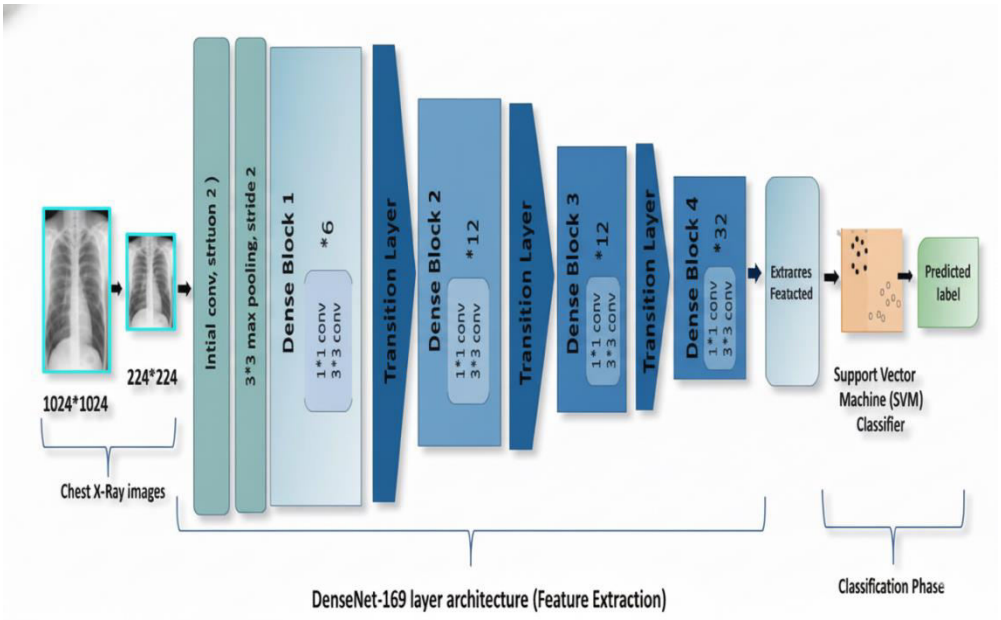

Figure. 2 Illustrates a flowchart of our applied methodology.

DenseNet-169 has been employed for the feature extraction process. This model was introduced by Huang et al. (2016) [18]. The variation employed in our investigation is trained on the extensive publicly accessible ImageNet dataset. The DenseNet-169 architecture consists of one initial convolutional and pooling layer, three transition layers, and four dense blocks. The final layer, known as the classification layer, follows these preceding levels. The initial convolutional layer executes 7×7 convolutions with a stride of 2, succeeded by 3×3 max pooling with a stride of 2. The network comprises a dense block succeeded by three sets, each consisting of a transition layer followed by a dense block. The deep connectivity advocated by Huang et al. [18] in DenseNets is achieved by establishing direct connections between any layer and any other layer in the network. The lth layer in the network acquires the feature maps from all preceding layers, hence enhancing the gradient flow throughout the whole network. This necessitates the concatenation of feature maps from preceding layers, which is only feasible if all feature maps are of uniform dimensions. However, since Convolutional Neural Networks primarily aim to downsample the size of feature maps, the DenseNet architecture is segmented into multiple densely connected dense blocks. The layers situated between these dense blocks are termed transition layers. Each transition layer in the network comprises a batch normalisation layer, a 1×1 convolutional layer, and a 2×2 average pooling layer with a specified stride.

As previously stated, there are four dense blocks, each including two convolutional layers, the first of size 1 × 1, followed by a 3 × 3 layer. The dimensions of the four dense blocks in the DenseNet169 architecture, pretrained on ImageNet, are 6 and 12.
Thirty-two and thirty-two. Adjacent to this is the final layer, the classification layer, which executes global average pooling of 7×7, followed by a concluding fully-connected layer that employs 'softmax' as the activation function.

Feature Extraction: The feature extraction procedure described in section 4.2.1 encompasses all layers of the network, excluding the final classification layer. The final feature representation was read as a 50176×1 dimensional vector, which was subsequently provided as input to various classifiers.

The Classification Phase

Subsequent to feature extraction, various classifiers, including Random Forest and Support Vector Machine, were employed for the classification task. The optimal results were achieved when the Support Vector Machine was employed as the classifier for the task. In the optimal suggested model, features derived from DenseNet-169 were utilised with an SVM classifier to achieve superior outcomes. The parameters and kernel utilised with SVM are described as follows: Consider a training dataset represented as (x1,y1), (x2,y2), ..., (xn,yn). The data must be divided into two sets of classes, where xi ε Fd denotes the feature vector and yi ε (0,1) signifies the label class. A Support Vector Machine employed for binary classification.

Classification can identify the optimal hyperplane for the aforementioned training data, specifically the one that maximises the margin between the classes and effectively separates the data points of one class from another. The efficacy of SVM is significantly contingent upon the choice of kernel and parameters. We employed the Gaussian radial basis function (RBF) kernel [19]. The gamma and C parameters of the RBF kernel significantly influence the performance of SVM. The gamma parameter intuitively determines the degree of influence a single training example exerts, with lower values indicating 'distant' and higher values indicating 'proximate.' The gamma parameter indicates the inverse of the radius of influence of the samples designated as support vectors by the model. Conversely, the C parameter mitigates the misclassification of training samples. A low C yields a smooth decision boundary, whereas a high C aims to accurately categorise all training samples by allowing the model to designate more examples as support vectors.

Experimental Setup

This section describes many tests conducted to offer the ideal model for pneumonia detection.

Feature Extractor and Classifier

For pre-trained convolutional neural network models, including Xception [20], VGG16 [21], VGG-19 [21], ResNet-50 [19], and DenseNet-121 [18] We assessed the performance of DenseNet-169 [18] using various classifiers, including Random Forest, K-nearest neighbours, Naive Bayes, and Support Vector Machine (SVM). Table-1 enumerates the performance of all models in the classification of abnormal and normal chest X-rays. The ResNet-50 CNN model, with a depth of 168, in conjunction with an SVM classifier, surpassed all previous pretrained CNN models, with an AUC score of 0.7749. We noted that DenseNets achieved performance comparable to ResNet50, attaining an approximate AUC of 0.75. Table-2 presents the outcomes achieved by DenseNet-121 and DenseNet-169. Statistical findings indicated that ResNet-50 and DenseNets (DenseNet-121 and DenseNet-169) are the most effective pre-trained CNN models for the feature-extraction phase, while SVM (with RBF kernel) serves as the classifier for the classification phase. Figure 3 illustrates the performance of ResNet50 and DenseNets (DenseNet-121 and DenseNet-169) in conjunction with various classifiers, highlighting the SVM classifier as the most effective for achieving superior AUC values across all three pretrained CNN models. During the assessment of the ideal CNN model, it was observed that VGGNets (VGG16 and VGG19) achieved the lowest scores among all the utilised pretrained models.

TABLE 1

RESULTS OBTAINED BY VARIOUS PRE-TRAINED MODELS WITH DIFFERENT CLASSIFIERS

Feature Extractor	Classifier	AUC
XCeption	SVM(rbf kernel)	0.7034
XCeption	Naïve Bayes	0.6362
XCeption	k-nearest neighbors	0.6867
XCeption	Random Forest	0.6406
VGG-16	SVM(rbf kernel)	0.5
VGG-16	Naïve Bayes	0.6193
VGG-16	k-nearest neighbors	0.6847
VGG-16	Random Forest	0.6563
VGG-19	SVM(rbf kernel)	0.5
VGG-19	Naïve Bayes	0.5952
VGG-19	k-nearest neighbors	0.68502
VGG-19	Random Forest	0.6481
ResNet-50	**SVM(rbf kernel)**	**0.7749**
ResNet-50	Naïve Bayes	0.6891
ResNet-50	k-nearest neighbors	0.7298
ResNet-50	Random Forest	0.5793

TABLE 2

RESULTS OBTAINED BY DENSENET-121 AND DENSENET-169 PRE-TRAINED MODELS WITH DIFFERENT CLASSIFIERS

Feature Extractor	Classifier	AUC
DenseNet-121	SVM(rbf kernel)	0.7577
DenseNet-121	Naïve Bayes	0.6691
DenseNet-121	k-nearest neighbors	0.6981
DenseNet-121	Random Forest	0.6771
DenseNet-169	SVM(rbf kernel)	0.7476
DenseNet-169	Naïve Bayes	0.6758
DenseNet-169	k-nearest neighbors	0.6835
DenseNet-169	Random Forest	0.6733

Optimal Hyperparameter Optimisation

To enhance model performance, we conducted hyper-parameter tuning using the SVM classifier with an RBF kernel in each instance. The procedure significantly influenced the statistical outcomes, achieving the highest AUC score to date. We executed around 350 distinct combinations of C and gamma with each recommended CNN model to attain improved outcomes. However, most optimised hyper-parameter settings exhibited no significant enhancement in performance. Tables 3, 4, and 5 enumerate solely the significant combinations of C and Gamma for all three models: (a) ResNet-50 followed by SVM, (b) DenseNet-121 followed by SVM, (c) DenseNet-169 followed by SVM. The optimisation of hyper-parameters for the SVM classifier resulted in a notable enhancement for both DenseNet-121 and DenseNet-169; however, no statistically significant increase was noted during hyper-parameter tweaking with features taken from ResNet-50. Figure 4 illustrates the fluctuation of AUC scores corresponding to various combinations of C and gamma in the context of DenseNet-121.

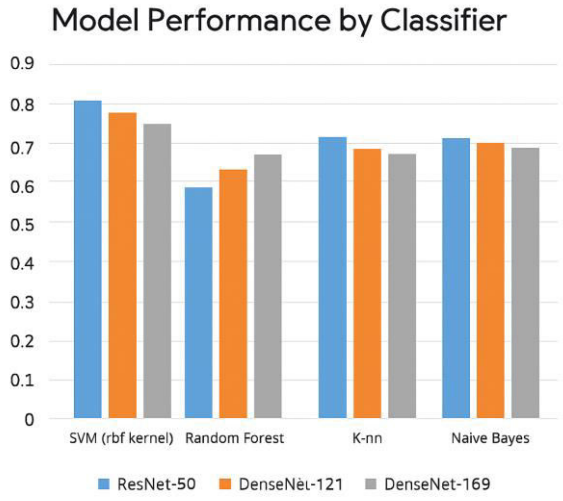

Figure. 3 Illustrates a bar graph of AUC scores derived from pre-trained CNN models employing various classifiers.

TABLE 3

RESULTS OBTAINED BY PARAMETER TUNING WHEN PRE-TRAINED RESNET-50 MODEL IS USED AS FEATURE EXTRACTOR

Technique	C	gamma	AUC
ResNet-50 + SVM	1.5	1.9e-05	0.7859
ResNet-50 + SVM	1.5	0.9e-05	0.7841
ResNet-50 + SVM	1.5	2.5e-05	0.7840
ResNet-50 + SVM	2	1.9e-05	0.7842
ResNet-50 + SVM	3	1.9e-05	0.7841

TABLE 4

RESULTS OBTAINED BY PARAMETER TUNING WHEN PRE-TRAINED DENSENET-121 MODEL IS USED AS FEATURE EXTRACTOR

Technique	C	gamma	AUC
DenseNet-121 + SVM	1.5	1.9e-05	0.7296
DenseNet-121 + SVM	2.0	1.9e-05	0.7634
DenseNet-121 + SVM	3	1.9e-05	0.7669
DenseNet-121 + SVM	3	0.9e-05	0.7699
DenseNet-121 + SVM	**3**	**0.85e-05**	**0.7717**
DenseNet-121 + SVM	3	0.8e-05	0.7681
DenseNet-121 + SVM	3.5	1.9e-05	0.7652

TABLE 5

RESULTS OBTAINED BY PARAMETER TUNING WHEN PRE-TRAINED DENSENET-169 MODEL IS USED AS FEATURE EXTRACTOR

Technique	C	gamma	AUC
DenseNet-169 + SVM	1.5	1.9e-05	0.7791
DenseNet-169 + SVM	2.0	1.9e-05	0.7901
DenseNet-169 + SVM	3	1.9e-05	0.7969
DenseNet-169 + SVM	3	0.9e-05	0.7966
DenseNet-169 + SVM	3.5	0.85e-05	0.7912
DenseNet-169 + SVM	**3.5**	**1.9e-05**	**0.8002**
DenseNet-169 + SVM	3.5	0.9e-05	0.7999
DenseNet-169 + SVM	3.5	2e-05	0.7904
DenseNet-169 + SVM	4	1.9e-05	0.7984

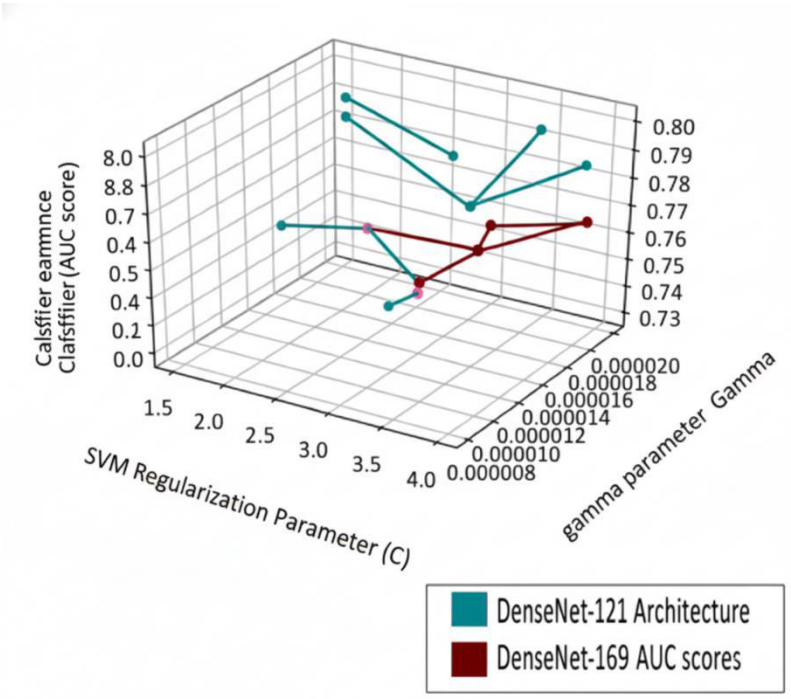

Figure. 4. Illustrates the impact of varying C and gamma parameter values on the AUC score for DenseNet121 and DenseNet169. We conducted the aforementioned experimental study to select the optimal model for classifying Chest X-Rays. In evaluating the ideal feature extractor among available pre-trained CNN models, ResNet50 surpassed the performance of all other models, followed by the SVM classifier at default hyperparameter settings. The ideal hyperparameter optimisation approach indicated that DenseNet-169 should be utilised for enhanced feature representation.

FINDINGS AND ANALYSIS

The tailored model, which integrates CNN-based feature extraction with a supervised classifier algorithm, yielded an optimal solution for classifying abnormal (Pneumonia-labeled) and normal chest X-ray images, primarily due to the significant features supplied by DenseNets, complemented by the optimal hyper-parameter values of the SVM classifier. Literary studies elucidate the contribution

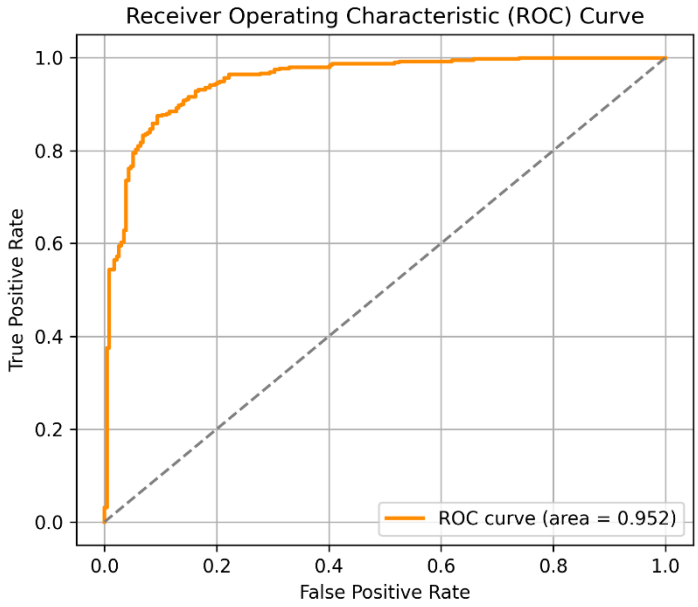

Figure. 5. Illustrates the test ROC curve for DenseNet-169.

Transfer learning methodologies, encompassing feature extraction, applied to visual recognition problems [4]. Consequently, we retrieved features from multiple variants of pre-trained CNN models, including VGGNets [21], Xception [20], ResNet-50 [19], and DenseNets [18]. Literature studies indicate the utilisation of classifiers in conjunction with CNN-based feature extraction

mostly in medical image analysis help enhance the efficacy of models. In accordance with the previously discussed methodologies, we assessed each of the pre-trained models using several classifiers to identify the optimal model for the objective. The comparative experimental results in Tables 1 and 2 indicate that ResNet50 surpassed the performance of all other pre-trained CNN models when utilised with the default parameter settings of the SVM classifier. Furthermore, DenseNets had performance comparable to ResNet50. Literature indicates that DenseNets surpassed all pre-trained CNNs on the ImageNet dataset (Huang et al., 2017 [18]). Consequently, we selected ResNet50, DenseNet-121, and DenseNet169 as the preferred CNN models for the feature-extraction phase, and SVM as the ideal classifier for the classification phase in further tests of the study. The choice of the SVM classifier with an RBF kernel, based on the statistical results depicted in Figure 3, then prompted the search for appropriate hyperparameter values (C and gamma). During hyper-parameter optimisation, we executed over 350 combinations of C and gamma, with the significant combinations detailed in Tables 3, 4, and 5. In this process, we found that DenseNet-169 surpassed all other customised models and was therefore selected as the optimal feature extractor for the final customised model, accompanied with the ideal hyper-parameter values of the SVM RBF kernel. The superior outcomes obtained with the DenseNet169 architecture as feature extractors can be attributed to its capacity to retrieve feature maps from all preceding levels. Literature on DenseNets indicates that the information flow from the initial layer to the final layers, along with the elimination of duplicate features by transition layers, are the principal factors contributing to high-feature representations. No research has been identified that examines the combination of CNN-based feature extraction and supervised classifier techniques for the specified problem. Within We have suggested a model architecture for identifying pneumonia from frontal chest X-ray images, employing Densenet as feature extractors and SVM as classifiers. In the course of enhancing model performance, we discovered that our tailored model surpasses the results reported in the recent study by Benjamin Antin et al. [15] about pneumonia diagnosis.

TABLE 6

RESULTS OF THE PROPOSED MODEL

Feature Extractor	Classifier	C	gamma	AUC
DenseNet-169	SVM(rbf kernel)	3.5	2e-05	0.7904
DenseNet-169	**SVM(rbf kernel)**	**3.5**	**1..9e-05**	**0.8002**
DenseNet-169	SVM(rbf kernel)	3.5	0.9e-05	0.7999

TABLE 7

RESULTS OBTAINED BY BENJAMIN ANTIN ET AL. (2017) [18] IN PNEUMONIA DETECTION PROBLEM.

Author Name	Technique Used	AUC
Benjamin Antin [18]	Logistic Regression	0.60
Benjamin Antin [18]	DenseNet-121	0.609

Table 6 and Table 7 present the statistical findings of our customised model and the model proposed by Benjamin et al., respectively. The proposed model in our study attains an AUC of 0.8002, with the corresponding ROC curve illustrated in Figure 5.

Constraints

Despite the results being compelling, our model exhibited certain limitations that we deem essential to acknowledge. The primary restriction is the absence of historical data regarding the associated patient in our evaluation model. Secondly, only frontal chest X-rays were utilised; nevertheless, it has been demonstrated that lateral view chest X-rays are as beneficial for diagnosis [22]. Thirdly, because to the extensive use of convolutional layers, the model requires substantial computer capacity; otherwise, it will consume excessive time in calculations.

CONCLUSION

The presence of skilled radiologists is essential for the accurate diagnosis of thoracic diseases. This research primarily seeks to enhance medical proficiency in regions where the presence of radiotherapists remains scarce. Our study enables the early detection of pneumonia to avert detrimental outcomes, including mortality, in distant regions. To date, limited efforts have been made to specifically diagnose pneumonia using the aforementioned information. The advancement of algorithms in this field can significantly enhance healthcare services. In this classifier, We evaluated the performance of several pretrained CNN models in

conjunction with different classifiers, subsequently selecting DenseNet-169 based on the statistical outcomes. Utilising feature extraction for the initial phase and support vector machines for the classification phase. We also showed that performing hyper-parameter optimization in the classification stage ameliorated the model performance. With the series of experiments conducted, we aim to provide the dominating pre-trained CNN model and classifier for the future work in the similar research domain. Our study will likely lead to the development of better algorithms for detecting Pneumonia in the foreseeable future.

REFERENCES

1. D. K. Das, M. Ghosh, M. Pal, A. K. Maiti, and C. Chakraborty, "Machine learning approach for automated screening of malaria parasite using light microscopic images," *Micron*, vol. 45, pp. 97–106, 2013.
2. M. Poostchi, K. Silamut, R. J. Maude, S. Jaeger, and G. Thoma, "Image analysis and machine learning for detecting malaria," *Translational Research*, vol. 194, pp. 36–55, 2018.
3. N. E. Ross, C. J. Pritchard, D. M. Rubin, and A. G. Duse, "Automated image processing method for the diagnosis and classification of malaria on thin blood smears," *Med Biol Eng Comput*, vol. 44, no. 5, pp. 427–436, 2006.
4. A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN features off-the-shelf: an astounding baseline for recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2014, pp. 806–813.
5. R. Nijhawan, J. Das, and R. Balasubramanian, "A hybrid CNN+ random forest approach to delineate debris covered glaciers using deep features," *Journal of the Indian Society of Remote Sensing*, vol. 46, no. 6, pp. 981–989, 2018.
6. H. Mohsen, E.-S. A. El-Dahshan, E.-S. M. El-Horbaty, and A.-B. M. Salem, "Classification using deep learning neural networks for brain tumors," *Future Computing and Informatics Journal*, vol. 3, no. 1, pp. 68–71, 2018.
7. R. Nijhawan, M. Rishi, A. Tiwari, and R. Dua, "A novel deep learning framework approach for natural calamities detection," in *Information and Communication Technology for Competitive Strategies: Proceedings of Third International Conference on ICTCS 2017*, 2018, pp. 561–569.
8. R. Nijhawan, R. Verma, S. Bhushan, R. Dua, A. Mittal, and others, "An integrated deep learning framework approach for nail disease identification," in *2017 13th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*, 2017, pp. 197–202.
9. P. Lakhani and B. Sundaram, "Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks," *Radiology*, vol. 284, no. 2, pp. 574–582, 2017.
10. K.-L. Hua, C.-H. Hsu, S. C. Hidayati, W.-H. Cheng, and Y.-J. Chen, "Computer-aided classification of lung nodules on computed tomography images via deep learning technique," *Onco Targets Ther*, pp. 2015–2022, 2015.
11. M. T. Islam, M. A. Aowal, A. T. Minhaz, and K. Ashraf, "Abnormality detection and localization in chest x-rays using deep convolutional neural networks," *arXiv preprint arXiv:1705.09850*, 2017.
12. R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, "Deep learning for healthcare: Review, opportunities and challenges," *Brief Bioinform*, vol. 19, no. 6, pp. 1236–1246, May 2017, doi: 10.1093/bib/bbx044.
13. 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2097–2106.
14. P. Rajpurkar *et al.*, "Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning," *arXiv preprint arXiv:1711.05225*, 2017.
15. B. Antin, J. Kravitz, and E. Martayan, "Detecting pneumonia in chest X-Rays with supervised learning," *Semanticscholar. org*, vol. 2017, 2017.
16. P. Kumar, M. Grewal, and M. M. Srivastava, "Boosted cascaded convnets for multilabel classification of thoracic diseases in chest radiographs," in *International conference image analysis and recognition*, 2018, pp. 546–552.
17. X. Yang, Z. Zeng, S. G. Teo, L. Wang, V. Chandrasekhar, and S. Hoi, "Deep Learning for Practical Image Recognition," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, New York, NY, USA: ACM, Jul. 2018, pp. 923–931. doi: 10.1145/3219819.3219907.
18. G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
19. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
20. F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1251–1258.
21. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
22. S. Raoof, D. Feigin, A. Sung, S. Raoof, L. Irugulpati, and E. C. Rosenow III, "Interpretation of plain chest roentgenogram," *Chest*, vol. 141, no. 2, pp. 545–558, 2012.