International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2121-2127
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0282

AI – Based Pneumonia Detection Using CNN and Transfer Learning

Vinodhini B[1], Akshay R[2], Gowtham B[3], Haripraneshwaran S[4], Jayavarman R[5]
[1]Assistant professor, Dept. of AI&DS, Jai Shriram Engineering College, Tirupur, Tamilnadu, India
[2,3,4,5]UG Scholar, Dept. of AI&DS, Jai Shriram Engineering College, Tirupur, Tamilnadu, India
Emails:*vinodhiniaids@jayshriram.edu.in1,akshayrr04@gmail.com2,gowthamx1208@gmail.com3, haripranesh861@gmail.com4, jayavarman.ravivarman@gmail.com[5]*

Abstract

Pneumonia still remains one of the primary causes of health issues and fatalities worldwide since it requires immediate treatment in healthcare institutions that do not always have powerful diagnostic instruments. The process of reading the chest X-ray images is not only difficult, but also time-consuming, and results offer varied outcomes as observed by different observers. This paper will discuss a deep learning model that will only consume little computing power to automatically discover pneumonia using convolutional neural networks, transfer learning, and model ensemble methods. To make new models with better generalization and overfitting resistance the researchers used several lightweight pre-trained models including AlexNet, ResNet-18, DenseNet-201 and SqueezeNet. The experiment results indicate good results in terms of accuracy, sensitivity and F1-score and less processing need. These were improvements authenticated on openly available chest X-rays.

Keywords: *Pneumonia Detection, Chest X-ray Imaging, Deep Learning, Transfer Learning, Model Ensembling.*

1. Introduction

Pneumonia is a severe respiratory disease and a major cause of death to lots of people in the world. Much of the pneumonia mortality rate can be avoided in the event of diagnosing and treating it in time and in an accurate way. The inexpensive way of screening is Chest X -rays, which require a specialist radiologist to interpret them, and the interpretation may differ widely among and within physicians, which is particularly problematic in busy or low-resource environments [12, 13, 17]. Lately, AI has been promising a great deal- deep learning and more specifically convolutional neural networks (CNNs) can be utilized in image processing, such as classifying and even detecting illnesses [1,3]. CNNs are trained with the images as the direct source of the layered features [2, 3]. Nevertheless, it is not easy to start fresh on medical data due to the small size of the datasets [4, 10]. Transfer learning emerged to address the issue of data by borrowing information on big natural-image datasets and applying it to medical tasks [4, 10, 11]. Leveraging fine-tuning pre-trained CNNs provides superior convergence and accuracy in the detection of pneumonia in chest X-rays [12, 18]. Regrettably, to date, the major majority of studies are done on heavyweight architectures and thus require a long time to incur the architecture and are difficult to implement on real-time devices or on edge devices [6, 7]. Besides, single-model methods may not be robust and may also be biased which means that the performance will not be consistent [14, 15]. Ensemble learning is used to assist in stabilizing and increasing the reliability of results by aggregating predictions of various diverse models [14]. Thus, on this, the proposal is a completely automated pneumonia detector based on the transfer learning method with CNNs on lightweight pretrained networks and follows up by fusing with the ensemble method. The system can be used to reduce the constraints of single-model transfer learning by incorporating multiple fined tuned lightweight CNNs, appropriate preprocessing, and data augmentation. Testing it on public chest X-ray datasets, we found that the model worked better and had more variability than the baseline ones.

2. Literature Review

In the early automated medical image analysis, the paradigms used mostly were classical machine-learning and rule based. Such methodologies commonly used handcrafted techniques of feature-extraction, e.g., histogram of oriented gradients, local binary patterns, and gray-level co-occurrence matrices, the result of which was later sent through traditional classifiers, e.g. support-vector machines, random forests, and k-nearest neighbors. The fact that none of these pipelines could be later replaced by a holistically validated experiment explicit framework meant that their operational effectiveness strongly depended on hand-crafted feature engineering and understanding of the domain. This raised scalability and heterogeneity questions: at that point, even the handcrafted feature sets could not capture complex pathological patterns to scale up to different circumstances of chest-x-ray acquisition, and their generalization to different imageries was not optimal. Deep learning resulted in a critical paradigm shift in medical-image analysis, classifying convolutional neural networks (CNNs) as the front modality to visual feature learning. Theoretical foundation and practical abilities of CNNs have been built, showing that they are able to autonomously project hierarchical spatial representations of raw image content [1, 2]. The high accuracy of deep CNNs in the classification of large collections of images also helped support the superiority of deep representations with training on large datasets [3]. Later studies used CNNs in form factorial of AlexNet and VGG to classify chest-x-rays, demonstrating improvements over conventional methods. Nevertheless, such gains were often thwarted by scarce supply of labelled medical data that resulted in overfitting and reduced generalization in clinical practice [4]. Transfer learning developed as an effective fix to address the shortage of annotated medical images, harnessing based model of massive natural-image datasets. Empirical comparisons showed that the systematic fine-tuning of the method over systematic training of the neuralization resulted in more consistent visualization of anatomical systems and structures of medical images [4]. The thoracic imaging-specific method of multisource transfer learning was shown to be able to transfer knowledge across datasets and increase the feature robustness and generalization [5]. These findings were later to be supported in subsequent studies that reported lower training times, accelerated convergence to detect pneumonia in chest-x-ray data and were also able to maintain competitive classification accuracy [12]. However, there was still disparity in performance on datasets even so, this highlights deep challenges of dataset bias and domain shift even still. Alongside the development of transfer-learning, the means of the CNN architecture improved the abilities of the chest-x-ray study as well. U-Net based encoder-decoder models with multidimensional features enabled use of anatomical target localization and segmentation of bodies, in order to facilitate the future diagnostic processes [5]. They were used to solve network degradation in Residual-learning models to achieve successful training of deeper architectures [6], and to promote feature re-use and relieve gradient propagation in densely-connected models [7]. Although they enjoyed high representational power, these advanced models level led to high parameter counts and computational complexity, and thus hinder their applications in fast-changing or resource-limited clinical settings. In an attempt to balance performance and computational efficiency, scholars came up with lightweight CNN architectures like SqueezeNet and MobileNet. These networks significantly decreased the number of parameters used and cost incurred in the calculus without significantly impacting predictivity in a significant way [8, 9]. They demonstrated the practicality of edge based medical imaging in real-time. However, the reduced accuracy of representation of lightweight models sometimes meant much worse performance than that of heavyweight models, continuing the long-running compromise of functionality and accuracy. Generalizations and robustness are still crucial issues in automated pneumonia mechanisms. Massive benchmarking initiatives and large-scale chest-x-ray datasets have revealed the drawbacks around class imbalance, label noise and internal dataset bias [13, 17]. These results highlight that strong performance in one dataset does not warrant the same to diverse clinical groups. Ensemble learning has also been considered as a solution, combining forecasts of complementary models to be more stable and more diagnostic [14, 15]. Furthermore, the interpretability methods, such as gradient-based visualization, have strengthened the clinical confidence as the decision procedure of the models can be explained [16, 20]. Taken together, these investigations indicate a pressing need to establish frameworks that maximize precision, stability and computational efficiency.

3. Proposed Methodology

Here, the system design and pathway of functioning of the proposed pneumonia detection system involving ensemble learning, lightweight model development, and CNN-based transfer learning guided. In essence, it is configured to scale easily and can be implemented in actual clinical environments and still provide an excellent diagnostic strength and remain efficient on the computation side. Figure 1 presents the entire framework and brings out the modular flow of data collection to final categorization.Shows Figure 1 Block diagram illustrating the end-to-end architecture of the proposed pneumonia detection framework.

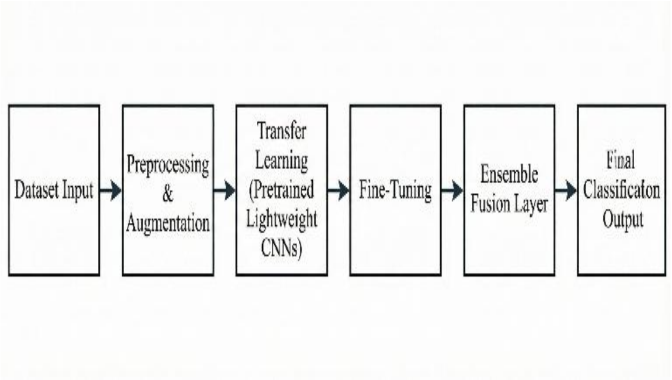

Figure 1 **Block diagram illustrating the end-to-end architecture of the proposed pneumonia detection framework**

3.1. Dataset Description

Our dataset is publicly available chest X-ray one that is usually utilized in the researches concerning pneumonia. The images are categorized as normal or pneumonia. In order to make our work reproducing, we simply concentrate on the basic features of data sets here and do not venture into extravagant statistics which do not belong to this part.

3.2. Data Preprocessing

To provide stability to the process of learning and maintain consistency at the same time, all chest X-rays are subjected to the same sequence of preparing steps prior to training the model. Due to the fixed input size on the pretrained CNNs we adjust the size of the images. We also do intensity normalization to overcome the disparities associated with various machines and settings. Background artifacts are suppressed and contrast improvement and noise reduction are used to highlight possibilities of clarity of the clinically important features.

In training, we perform data augmentation both to equalize the classes and enhance generalization. Augmentations can have both intensity adjustments such as scaling and brightness adjustments as well as controlled geometric transformations such as rotation, horizontal flipping, translation, and random crops. It selects the augmentations carefully so as to add variety without loss of diagnostic meaning. This preprocessing step aids in avoiding overfitting, enhances resistance to new samples, and the amount of predictive power of the model to new clinical conditions.

3.3. Model Architecture

The system is based on the core of pretrained CNNs of varying representational power and computational efficiency. Lightweight CNNs are the learners set as the foundation in order to reduce the parameters and inference time. Every network launches weights that have been taught on huge sets of natural-images, which means that features can be retired. Our approach involves a transfer-learning approach: all the lower convolutional layers remain fixed to extract such generic visual structures as edges and textures, whereas the higher-level layers are fine-tuned to respond to pneumonia-related information in the chest X-rays. Task specific fully connected layers are exchanged with the original classification heads that aim at binary classification. The depth of fine-tuning is controlled by us to balance between adaptability and stability in training so that we can still perform well even with a limited amount of data.

Transfer Learning Fine-Tuning Objective This equation formalizes how pretrained CNN parameters are adapted to the pneumonia classification task.

$$\mathcal{L}(\theta) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c}\log\left(\hat{y}_{i,c}\right)$$

Where

- N= number of training samples
- C= number of classes (Pneumonia, Normal)
- y_(i,c)= ground truth label
- yˆ_(i,c)= predicted probability from the fine-tuned CNN

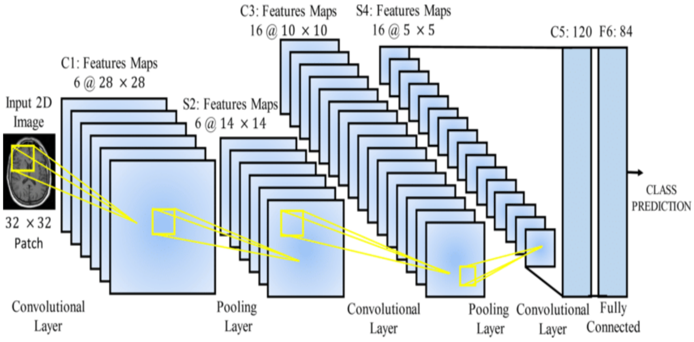

Figure 2 Overview of the proposed CNN architecture

The architectural design emphasizes modularity, allowing individual base models to be trained independently while maintaining compatibility with the ensemble framework. Figure 3 presents the architectural flow, illustrating how pretrained lightweight CNNs are integrated into the overall pipeline.Shows Figure 2 Overview of the proposed CNN architecture.

3.4.Model Ensembling Strategy

In order to address the issues associated with the single-model learning, the proposed system makes use of the ensemble learning technique, where multiple models of the lightweight fine-tuned CNNs are combined to improve the performance of the system. In the decision level, the combination of the models is performed through the soft voting technique, where the probabilities of the classes are averaged to obtain the final output. The advantage of the proposed technique is that it improves the robustness of the system through the reduction of the variance and the individual models' bias. The ensemble technique improves the generalization and robustness of the system through the varied datasets by learning the feature representation through the combined models, where the performance is improved without significantly increasing the computational cost to ensure the clinical deployment of the system in real-time.

Ensemble Decision Fusion

This equation mathematically defines how predictions from multiple CNN models are combined.

$$\hat{y}_{ensemble} = \frac{1}{M}\sum_{m=1}^{M}\hat{y}_m|$$

Where:

- M= number of CNN models in the ensemble
- yˆ_m= output probability from the m^thmodel
- yˆ_ensemble= final ensemble prediction

3.5. Training Configuration

Training of the model is done with the use of conventional optimization strategies that can be used on medical image analysis based on deep learning. Adaptive optimization algorithm is utilized to maintain the stability of convergence in the process of fine-tuning. Binary classification is guided with the help of a cross-entropy loss function. Mini-batch learning with fixed batch size through several epochs is used to perform training and learning rates are closely monitored to avoid overfitting and allow even convergence. To improve model generalization, regularization methods, such as dropout and weight decay are added. All base models have similar training configurations, which are designed to guarantee reasonable integration of the ensembles and reproducibility. At this point, no task-specific heuristics or data-dependent tuning are added to ensure the methodological rigor.

4. Model Evaluation Metrics

In an effort to be honest in determining how effective and clinically reliable this model of pneumonia detection truly is, we draw in numerous metrics of quantitative evaluation. Such measures allow us to view the overall performance of classification, their susceptibility to pneumonia cases, their ability to stand up to the class imbalance, and diagnostic reliability in the medical imaging environment.

4.1. Confusion Matrix – Based Metrics

We begin with the confusion matrix which includes true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN). Based on them, we get some critical measure: Accuracy informs us about the proportion of all the samples on which the classification was correct.Precision is the percentage of accurately detected cases of pneumonia among all samples as the flag of pneumonia cases. High precision ensures that there are fewer false positives and this is important to prevent unnecessary clinical responses. Recall is the measure of the ability of the model to detect pneumonia. It is basically sensitivity - the lack of a negative case may lead to the delay of the treatment and therefore the recall is vital. The F1-score is the harmonic mean of precision and recall that provides us with one figure that trades-off between the two, this is particularly useful when working with imbalanced medical imaging data.

$F1=(2\cdot Precision\cdot Recall)/(Precision+Recall)$

4.2. Receiver Operating Characteristic And Auc

We run ROC analysis in order to see the model in more various decision thresholds. The AUC (Area Under the Curve) is a threshold-free measure which informs us of the closeness of the model to classifying the two classes. Increase in AUC implies a stronger discriminative strength and increased confidence in diagnostics[21].

4.3. Loss Function Monitoring

Categorical cross-entropy loss is monitored during training. It comes in handy to monitor convergence, as well as the understanding of whether the model is holding steady. We can tell at early stages, particularly in fine-tuning pretrained models, that we are over- or underfitting, by observing loss trends as well as the classification metrics[22].

4.4. Model Robustness and Generalization

To have a complete analysis, we calculate measures both on an individual base model and the ensemble measure. A comparison between these values allows us to understand the extent to which the ensemble decreases the predictor variance, enhances the stability, and NSE with other datasets[23].

5. Results And Performance Analysis

This section will provide the results of the experiment that we achieved in identifying pneumonia in chest X-ray images. We compared the individual pretrained lightweight CNN models with the ensemble of them using conventional medical image classification metrics.Shows Figure 4 Performance evaluation of the suggested pneumonia detection framework[24].

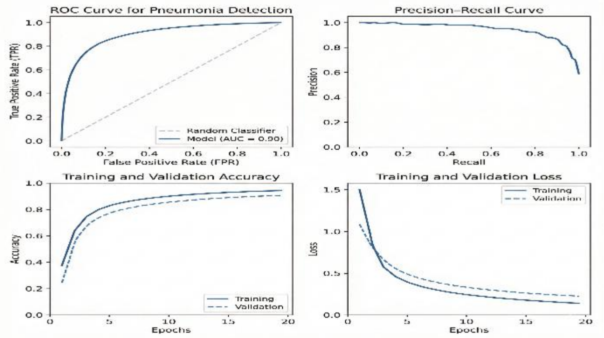

Figure 4 Performance evaluation of the suggested pneumonia detection framework

To obtain a comprehensive evaluation of diagnostic performance, we measured accuracy, precision, recall, specificity, F1 -score, and AUC. The findings indicate that transfer learning improves upon training performance when starting afresh, thus supporting their applicability during the situation where scarce labeled data are available. Light CNNs possess very good discriminative power and their computational requirement is low which is excellent in terms of efficient deployment. The ensemble achieved superior results over all the base models in all the evaluation measures. The enhancement of recall and F1 -score, in particular, indicate that the ensemble detects and removes more pneumonia cases and minimizes false negativity. These benefits arise out of the combination of complementary feature representations that are learnt by different pretrained models, which lowers model bias and variance. The ROC analysis again denotes the strength of the ensemble and the larger AUC values imply the greater power of class separability. Our fine-tuning and regularization methods also prove successful as seen in the training and validation loss curves showing stable convergence with low regularization.

6. Conclusion

The main idea of this paper is to demonstrate how to erect a robust and high-performance deep-learning pipeline to detect pneumonia in chest X-rays, based on CNN transfer learning and certain ensemble techniques. It addresses issues such as insufficient labeled data, high computational expenses and overfitting by connecting small pretrained CNNs to it. The accuracy, robustness and generalization improve significantly when several fine-tuned lightweight models are combined in an ensemble than when an individual model is used or when typical transfer learning is used. It is also made to match the good balance between performance and efficiency as it is also designed to be deployable. The entire architecture is scalable, adjustable and fits perfectly well into computer-aided diagnosis systems, particularly when the resources are limited. In general, the conclusion is that a combination of transfer learning and lightweight CNNs with ensemble learning presents a reliable, clinically applicable technology to screen pneumonia by analyzing chest X-rays and can help move the deep-learning diagnostics into the real-life medical environments.

Reference

[1]. Albawi, S., Mohammed, T. A., & Al-Zawi, S. (2017). Understanding of a Convolutional Neural Network. https://arxiv.org/pdf/1511.08458.pdf

[2]. LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional Networks and Applications in Vision. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=5537907

[3]. Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). ImageNet Classification with Deep Convolutional Neural Networks. https://papers.nips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf

[4]. Tajbakhsh, N., Shin, J. Y., Gurudu, S. R., et al. (2016). Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning https://arxiv.org/pdf/1706.03465.pdf

[5]. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. https://arxiv.org/pdf/1505.04597.pdf

[6]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. https://arxiv.org/pdf/1512.03385.pdf

[7]. Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely Connected Convolutional Networks. https://arxiv.org/pdf/1608.06993.pdf

[8]. Iandola, F. N., Han, S., Moskewicz, M. W., et al. (2016). SqueezeNet: AlexNet-level accuracy with 50x fewer parameters. https://arxiv.org/pdf/1602.07360.pdf

[9]. Howard, A. G., Zhu, M., Chen, B., et al. (2017). MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. https://arxiv.org/pdf/1704.04861.pdf

[10]. Pan, S. J., & Yang, Q. (2009). A Survey on Transfer Learning. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=5288526

[11]. Akçay, S., Kundegorski, M. E., Devereux, M., & Breckon, T. P. (2016). Transfer Learning Using CNNs for Object Classification within X-ray Baggage Security Imagery.https://arxiv.org/pdf/1606.01371.pdf

[12]. Rahman, T., Chowdhury, M. E. H., Khandakar, A., et al. (2020). Transfer Learning with Deep CNN for Pneumonia Detection using Chest X-ray. https://www.mdpi.com/2076-3417/10/9/3233/pdf

[13]. Rajpurkar, P., Irvin, J., Ball, R. L., et al. (2018). Deep Learning for Chest Radiograph Diagnosis: CheXNeXt. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6211269/pdf/pmed.1002686.pdf

[14]. Sirazitdinov, I., Kholiavchenko, M., Mustafaev, T., et al. (2019). Deep Neural Network Ensemble for Pneumonia Localization from Chest X-rays. https://arxiv.org/pdf/1905.02149.pdf

[15]. Jung, H., Kim, B., Lee, I., et al. (2018).

Classification of Lung Nodules in CT scans Using 3D CNN with Checkpoint Ensemble. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5848308/pdf/12880_2018_Article_297.pdf

[16]. Selvaraju, R. R., Cogswell, M., Das, A., et al. (2017). Grad-CAM: Visual Explanations from Deep Networks. https://arxiv.org/pdf/1610.02391.pdf

[17]. Wang, X., Peng, Y., Lu, L., et al. (2017). NIH ChestX-ray14 Dataset and Benchmarks. https://arxiv.org/pdf/1705.02315.pdf

[18]. Xu, S., Wu, H., & Bie, R. (2018). CXNet-m1: Anomaly Detection on Chest X-rays. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8260326

[19]. Souza, J. C., Diniz, J. O. B., Ferreira, J. L., et al. (2019). Automatic Lung Segmentation in Chest X-ray Using Deep Neural Networks. https://www.sciencedirect.com/science/article/pii/S0169260719302149

[20]. Rajaraman, S., Candemir, S., Kim, I., et al. (2018). Visualization and Interpretation of CNN Predictions in Pediatric Pneumonia. https://www.mdpi.com/20763417/8/10/1715/pdf

[21]. Christodoulides, S., Xanthopoulos, M., Ebner, L., et al. (2016). Multisource Transfer Learning with CNNs for Lung Pattern Analysis. https://arxiv.org/pdf/1609.06101.pdf

[22]. P. Mooney, Chest X-Ray Images (Pneumonia), Kaggle Dataset, 2018. [Online]. Available: https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia

[23]. P. Mooney, Chest X-Ray Images, Kaggle Dataset, 2018. [Online]. Available: https://www.kaggle.com/datasets/paulti/chest-xray-images

[24]. T. Dincer, Labelled Chest X-Ray Images, Kaggle Dataset, 2019. [Online]. Available: https://www.kaggle.com/datasets/tolgadincer/labeled-chest-xray-images