JOURNAL OF BASIC SCIENCE AND ENGINEERING

Pneumonia Detection Using X-ray Images Based On Deep Learning Technology

Omkar Kasulwar[1*], Bhakti Sonawane[2]

[1]PG Scholar, Computer Engineering Department, Mumbai University
[2]Assistant Professor, Computer Engineering Department, Mumbai University

DOI: https://doie.org/10.10399/JBSE.2025462122

Abstract

Pneumonia is a dominant worldwide cause of morbidity and mortality, mainly among vulnerable patient groups like the elderly and children. In many diagnostic protocols, chest X-ray is in regular use, although precise interpretation must often be based on the evaluation of experienced radiologists, giving rise to gaps in the efficiency of diagnosis in a timely fashion, particularly under resource-limited conditions. Modern developments in the field of deep learning, mainly convolutional neural networks (CNNs), hold great potential to automate pneumonia diagnoses from radiologic images. This work compares the performance of some pre-trained CNN models—LeNet, GoogleNet, VGG16, and Inception V3—against a manually developed Manual Net model. The models were tested on an open-source chest X-ray data set, where performance was evaluated in terms of accuracy, precision, recall, and F1-score. The findings show that although pre-trained networks provide robust baseline performance, the Manual Net architecture provides a potential lightweight alternative with lower computational demands, making it deployable in low-resource settings. The work also discusses interpretability, training efficiency, and dataset issues. Through this work, we shed light on the choice and optimization of CNN models for real-world pneumonia detection, moving towards more accessible and accurate diagnostic tools.

Keywords: *Pneumonia Detection, Deep Learning, Convolutional Neural Network, Pre-trained Models, Medical Imaging, Chest X-ray.*

1. Introduction

Pneumonia is a severe respiratory illness that induces inflammation of the lung alveolar sacs, which affects the exchange of oxygen and threatens significant health, particularly in infants, the elderly, and the immune compromised. It remains a global major cause of death, with more than 800,000 children under five years dying of it each year, which outnumbers other principal diseases such as malaria and measles [1]. Although antibiotics and vaccines are widely used, the early and correct diagnosis of pneumonia is still a major worldwide public health problem. Chest X-rays continue to be the most common imaging device employed in diagnosing pneumonia. Although very widely available and cheap, however, their reading tends to be subjective and based highly on the skills of a radiologist. Such reliance implies inconsistency in diagnosis and makes the uniform clinical interpretation of them very hard, especially in low-resource or high-batch settings [2]. Due to this need exists for sound, automated diagnosis technologies that may complement or supplement radiologists' activities in clinical settings.

Over the last few years, deep learning has surfaced as a ground breaking methodology for medical imaging. Specifically, CNNs have proven to be outstanding in pattern recognition problems by learning spatial hierarchies directly from images, allowing for effective analysis of intricate medical data [3]. CNNs eliminate the necessity for feature engineering by hand and can generalize across multiple image domains, making them particularly well-suited for radiographic image classification [4]. A number of CNN-based works have already been demonstrated to be effective in the diagnosis of thoracic diseases from chest radiographs. For example, attention-based deep learning models have been designed to increase model attention on diseased areas, enhancing classification accuracy and interpretability [5]. Radiomics-enhanced networks also combine texture-based features with learned representations, improving the sensitivity of detection systems for pneumonia and COVID-19 differentiation [6].

Pre-trained CNN architectures are extensively applied based on their performance with large datasets and potential for fine-tuning to domain-specific tasks. Prominent among these are LeNet, GoogleNet, VGG16, and Inception V3. LeNet, a pioneering CNN model, provided a basic but effective pattern of alternating convolution and pooling layers. Although it was originally developed for digit recognition, it has been utilized for medical applications because of its lightweight design and training efficiency [7]. GoogleNet pioneered the inception module, a structural improvement that conducts parallel convolutions using various kernel sizes, allowing the model to discern multi-scale features. This modularity greatly minimized the number of parameters and enhanced computational efficiency [8]. VGG16, however, utilized deep but symmetrical architecture with the use of stacked 3×3 convolution filters. Its symmetrical nature allows for easier fine-tuning and transfer learning in medical image classification problems [9].

Inception V3 further developed the inception module design through the addition of factorized convolutions and auxiliary classifiers, enhancing both speed and convergence. It is therefore a model of choice in research aimed at balancing depth and speed for high-dimensional image datasets [10]. Practical applications of the models in clinical settings are, however, limited by their resource and computational requirements. To counteract this, some researchers have suggested hybrid models that blend CNNs with classic machine learning algorithms. Hybrid methods involving deep feature extraction and subsequent use of classifiers such as SVMs or decision trees have yielded better results in pediatric pneumonia detection [11]. Ensemble methods, where predictions of more than one deep learning model are combined, have also proved promising in curbing misclassification errors and enhancing diagnostic accuracy [12].

Aside from performance, computational feasibility must be taken into account in deploying CNN models. Lightweight models have been investigated for application in mobile and embedded systems to facilitate point-of-care diagnostics. A lightweight architecture proposed by Trivedi and Gupta showed the viability of reduced models for mobile health applications without sacrificing much accuracy [13]. LeNet variants and reduced CNNs have also been optimized for real-time application in embedded devices [14].

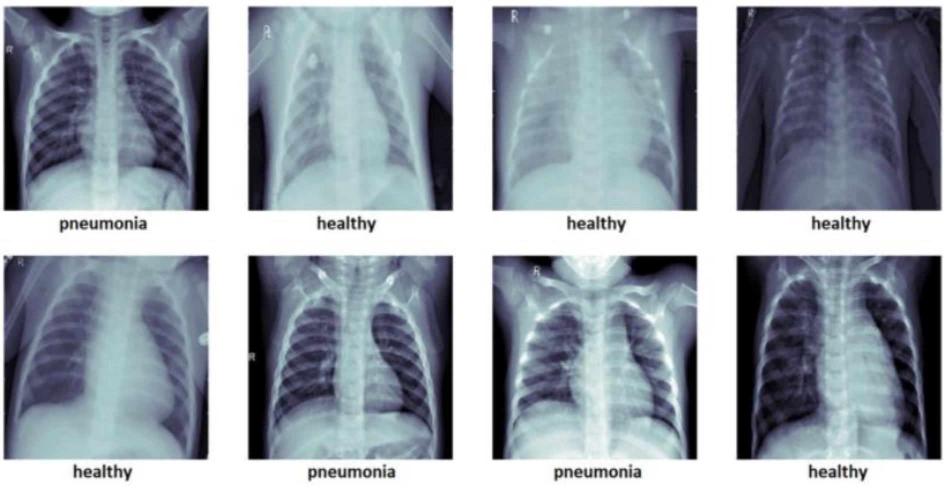

Figure 1: Some examples of input images with their ground truth. The eyes of a non-expert human can hardly distinguish positive and negative X-ray scans.
Image Source: (https://www.sciencedirect.com/science/article/pii/S0208521622000742)

Deep learning model performance is inherently coupled to the quality of the training dataset.This work is popularly referenced Chest X-ray Pneumonia dataset from Kaggle and contain labelled instances of normal, bacterial pneumonia, and viral pneumonia cases [15]. The balanced representation of pneumonia subtypes in the dataset permits thorough model assessment. Model performance must be understood using standardized evaluation metrics. These are accuracy, which is the overall correctness of the predictions; precision, which is the number of predicted positive cases that are actually positive; recall, which measures the sensitivity of the model; and F1-score, which is the balance between precision and recall. These are critical in a medical setting, where false negatives would result in delayed treatment, and false positives would result in unnecessary procedures [16]. While the pre-trained CNNs have many strengths, their black-box nature and lack of interpretability continue to be hurdles for adoption in healthcare. Some of these visualization methods have been suggested, including Grad-CAM, to bring into focus the exact areas in an image that contributed to the model's decision-making process, giving transparency and interpretability to clinicians [17]. These also help in achieving trust in AI-based diagnosis.

Also, new training methods based on low-precision calculations have been studied to enhance generalization and lower the computational requirements of CNNs. The technique enables models to train quickly and deploy more effectively in low-resource environments [18]. The development of mobile-optimized structures such as MnasNet also indicates that it is feasible to have high accuracy at low hardware needs, supporting real-time diagnosis on edge devices [19]. New technologies such as RetinaNet have been utilized in lung disease detection, demonstrating the application versatility of current state-of-the-art object detection architectures to classification in radiology [20]. Such breakthroughs reflect a general trend in incorporating multiple AI methods towards better diagnostic results.

1.2 Problem Statement

Though pre-trained CNNs have demonstrated high performance in pneumonia detection, their high computational needs and low ability to adapt to certain datasets limit their clinical applicability. Further, there has been a knowledge gap in testing the tradeoff between efficient lightweight models and powerful deep networks on the same setting. This paper bridges that gap by performing comparative studies of pre-trained and locally designed CNN structures under the same setup.

1.3 Research Objectives

1. To compare and benchmark four widely known CNN architectures, namely LeNet, GoogleNet, VGG16, and Inception V3, against a manually built Manual Net model on chest X-ray images for pneumonia classification.
2. To compare the diagnostic performance of each model through common evaluation criteria like accuracy, precision, recall, and F1-score.
3. To determine models that strike an optimal balance between classification accuracy and computational efficiency to be used in real-time and resource-limited healthcare environments.
4. In order to gain information for the development of future AI-based diagnostic software specific to pneumonia and other respiratory diseases.

The rest of this work is presented as follows: Section II includes a comprehensive literature review with focus on the recent developments in deep learning approaches used for pneumonia diagnosis. Section III presents the problem statement and dataset used in this work, encompassing data sources, types, and distribution. Section IV discusses the methodology, mentioning both the proposed custom architecture (Manual Net) and pre-trained CNN models including LeNet, GoogleNet, VGG16, and Inception V3. Section V discusses the experimental results and thorough analysis of model performance on the basis of evaluation measures such as accuracy, precision, recall, and F1-score.

Lastly, Section VI wraps up the paper by summarizing the major findings and providing some insights into future research directions for improving automated medical diagnostics using deep learning.

2. Literature Review

Pneumonia diagnosis from chest X-ray images has gained significant research interest in the last decade with the growing demand for precise, automated diagnosis systems. Conventional diagnostic methods, although helpful, tend to be plagued by issues such as subjectivity, latency, and reliance on highly skilled radiologists. This has driven the incorporation of deep learning, especially convolutional neural networks (CNNs), to aid or automate the detection of pneumonia from chest radiographs. In this review of the literature,it discussed three prevailing themes to emerge in the academic community: (1) CNN-based methodologies for detecting pneumonia, (2) the use of transfer learning and pre-trained models, and (3) challenges of datasets, interpretability, and clinical incorporation.

2.1 CNN-Based Methods for Pneumonia Diagnosis from Chest X-rays

Several studies have emphasized building CNN-based models specific to pneumonia classification. Zhang et al. (2021) presented a CNN model optimized for pediatric chest X-rays and highlighted how proper tuning of convolution and pooling methods had a direct effect on diagnostic accuracy. Mujahid et al. (2022) used the Inception V3 network to build a pneumonia classifier, obtaining better sensitivity through the use of deep convolutional features. Jaiswal et al. (2019) used a shallow CNN with fewer layers to enable real-time deployment, especially in rural or mobile health settings where computational resources are limited. Parthasarathy and Saravanan (2024) combined Harris Hawks Optimization with CNN to optimize training weights and enhance convergence rate, illustrating how bio-inspired algorithms could be used to enhance deep learning models. Bal et al. (2024) used a hybrid approach by using deep CNN feature extractors with conventional machine learning classifiers in order to enhance pneumonia diagnosis in pediatric radiographs. Bhatt and Shah (2023) built on the ensemble method by stacking multiple CNN predictions, which greatly decreased variance and false classification rates. Wong et al. (2022) presented a multi-scale attention CNN that was able to emphasize important pulmonary areas in X-rays, which raised diagnostic transparency and clinician confidence. Goyal and Singh (2023) highlighted the use of traditional ML and CNN-based systems to identify pneumonia and COVID-19, highlighting the flexibility of deep learning when combined with traditional algorithms. These studies as a whole highlight CNNs' strong performance, flexibility, and ability to accurately detect pneumonia, making them suitable candidates for automated diagnostic systems in clinical processes.

2.2 Transfer Learning and Pre-trained Architectures for Medical Imaging

Transfer learning emerged as a primary technique in classifying medical images, particularly for environments where databases are small or expensive to tag. Kavya et al. (2022) had illustrated how it was possible for transfer learning involving convolutional neural networks to make precise distinctions of COVID-19 from pneumonia when they had barely enough training images. Zhang et al. (2022) updated the LeNet-5 architecture to enhance medical image performance, affirming the utility of evolving early CNN models for contemporary tasks. Hu et al. (2022) integrated radiomics and deep learning by combining texture features with CNN learned representations, achieving enhanced pneumonia classification accuracy. Harsono et al. (2022) applied RetinaNet with a transfer learning strategy for thoracic nodule detection, providing building blocks for applying to pneumonia detection. Tan et al. (2019) presented MnasNet, a mobile-optimized neural architecture search model, showing the potential of low-latency models for point-of-care diagnosis. Courbariaux et al. (2014) suggested training networks using low-precision operations, significantly cutting memory consumption and accelerating training time—both essential in clinical AI environments. Liu et al. (2008) analyzed image retrieval algorithms that played a key role in constructing robust feature representations for content-based classification tasks, which impacted how deep learning features are extracted nowadays. Trivedi and Gupta (2023) used a light-weight CNN specifically designed for mobile diagnosis and demonstrated that it is possible to obtain high diagnostic accuracy without huge network depths. Together, these papers demonstrate that pre-trained models and transfer learning are a flexible, computationally frugal, and scalable solution for the detection of pneumonia that can be applied globally in healthcare settings. They alleviate the effort of training deep networks from scratch and improve generalizability, something vital in medical applications where dataset variability is high.

2.3 Datasets, Visual Interpretability, and Clinical Integration Challenges

Accurate and representative datasets are the backbone of creating accurate and generalizable deep learning models. Our World in Data (2021) pointed out that pneumonia remains one of the most fatal child diseases, hence highlighting the need for creating accurate diagnostic tools. Wang et al. (2020) tackled dataset variation by integrating attention mechanisms in Thorax-Net, enhancing the model's capacity to attend to diagnostically important regions of the lung. The publicly accessible Kaggle Chest X-ray Pneumonia dataset (2022) is now a standard for benchmarking models due to its presence of labeled images of normal, viral, and bacterial pneumonia cases. Datagen (2022) also provided architectural detail of the VGG16 model, often applied in pneumonia diagnosis studies owing to its robust transfer learning capabilities. DeepAI (2022) presented Grad-CAM as a method for visualization of areas within an image contributing to CNN prediction, assisting clinicians in comprehension and reliance on model predictions. Zhang et al. (2021) used Grad-CAM to interpret CNN-based pneumonia detection outcomes, improving the transparency and applicability of their models in clinical practice. Zech et al. (2018) cautioned against generalizability issues with models learned on single-hospital datasets since they might not replicate across multiple clinical environments due to latent biases. Rajpurkar et al. (2017) built CheXNet, a deep network with ChestX-ray14 training data, reaching radiologist-level performance, but also recognized the limitations of the model in interpretability and dataset imbalance. These studies show that in addition to architectural design, there are factors such as interpretability, data diversity, and deployment readiness that must be considered in order to bring AI from research to practice in medical diagnosis. Upcoming studies have to tackle these aspects to ensure equitability, clinical feasibility, and scalability.

The literature reviewed depicts an increasing trend away from classical machine learning towards deep learning models, specifically CNNs, for detecting pneumonia. CNN-based approaches have repeatedly been superior to traditional methods in terms of accuracy and stability. Transfer learning and pre-trained models have been effective in coping with small datasets, facilitating the deployment of AI systems in clinical diagnostics. Nevertheless, model interpretability, dataset heterogeneity, and deployment practicability continue to be important hurdles. Future work will need to address hybrid systems, low-complexity models, and explainable AI methods in order to close the gap between high-performance research prototypes and actual clinical tools.

3. Methodology

The primary objective of this research is to evaluate and compare the performances of various deep learning models, such as reputable pre-trained CNNs like LeNet, GoogleNet (Inception V3), and VGG16, with our recommended custom model, Manual Net, in correctly classifying pneumonia using chest X-rays. This is a strict approach to methodology whereby each step of data preparation as well as analysis of model performance follows stringent standards of science.

3.1 Data Collection and Pre-processing

The main dataset that has been employed consists of 5,863 chest X-rays images gathered from the publicly available Kaggle portal. The X-rays are differentiated into normal ones and pneumonia categories, which again are divided between viral and bacterial pneumonia. So as to get a better pre-processed dataset suited for machine learning, a detailed pre-processing was implemented. Images were resized equally to 224 × 224 pixels, the standard input size for feeding data into the selected CNN architectures. The pixel intensities were normalized to range between 0 and 1 to facilitate quicker and more stable training of the neural networks. Other augmentations, such as rotations, horizontal flip, and zooming, were also used to increase the dataset and reduce model over fitting.

3.2 Pre-trained CNN Architectures

In this study, a number of established pre-trained CNN architectures were employed to give a strong comparative analysis:

3.2.1 LeNet:

One of the earliest and most powerful convolutional neural network designs, utilised especially for image classification, LeNet was developed by Yann LeCun. Originally designed to read handwritten numbers, LeNet's architecture is somewhat straightforward yet powerful, consisting of two primary convolutional layers each preceded by pooling layers. Small-sized filters, usually 5×5 kernels, are used by the convolutional layers to effectively recover spatial hierarchy in pictures. LeNet's simplicity and low processing requirements fit it for settings with minimal resources. Its pool layers reduce dimensionality even more, hence increasing computational efficiency and reducing probable overfitting. The last few layers of the network are fully linked layers that permit precise pattern classification by means of combination of the retrieved information. These structural characteristics make LeNet a well-liked benchmark model that practically serves as a baseline for comparison studies in image recognition research. Fig. 3 shows the architectural overview modified from LeCun et al. (1998).

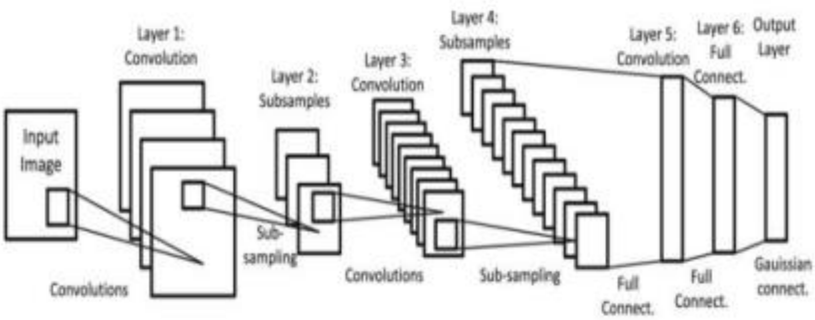

LeNet (Convolution Operation) mathematically represented in Equation 1 as:

$$Y_{i_j}^{(l)} = f(b^{(l)} + \sum_m \sum_n W_{m,n}^{(l)} X_{(i+m),(j+n)}^{(l-1)}) \tag{1}$$

Where,

- $Y_{i_j}^{(l)}$ is the output of neuron (i, j) in layer l.
- f is the activation function (e.g., ReLU, Sigmoid).
- $W_{m,n}^{(l)}$ represents the convolutional weights in layer l.
- $X^{(l-1)}$ is the input from the previous layer.
- $b^{(l)}$ is the bias term.

3.2.2 GoogleNet: Presented by Google researchers, GoogleNet, often known as the Inception network, is a sophisticated convolutional neural network. Particularly in image classification, GoogleNet was a significant advance in deep learning defined by its inception modules. Apart from pooling operations, the inception module is defined by the capacity to execute several convolution operations simultaneously employing various kernel sizes—for example, 1×1, 3×3, and 5×5—and then concatenate the results. Images processed using this multi-scale approach yield both fine-grained information and wider contextual information. By using the capabilities of dimensionality reduction techniques, mostly realised via 1×1 convolutions, GoogleNet achieves one of its most notable feats: efficient computation. By limiting duplicated parameters, these 1×1 convolutions lower feature size, cut processing expense, and prevent overfitting. Furthermore, GoogleNet included extra classifiers at intermediate network layers to boost gradient flow during training and therefore aid to mitigate the vanishing gradient issue. GoogleNet consistently performs better with these distinctive qualities—its inception modules, dimensionality reduction, and auxiliary classifiers—making it very relevant for complex visual recognition applications as medical picture diagnostics. Adapted from Szegedy et al. (2015) [2], Fig. 4 reveals the model structure.

Fig. 4. Architecture of GoogleNet [43]

GoogleNet (Inception Module):

GoogleNet's inception module output can be mathematically represented in Equation 2 as follows:

$$Y^{(l)} = f([W_{1*1} * X^{(i-1)}, W_{3*3}X^{(l-1)}, W_{5*5}X^{(l-1)}, Pool(X^{(l-1)})]) \quad (2)$$

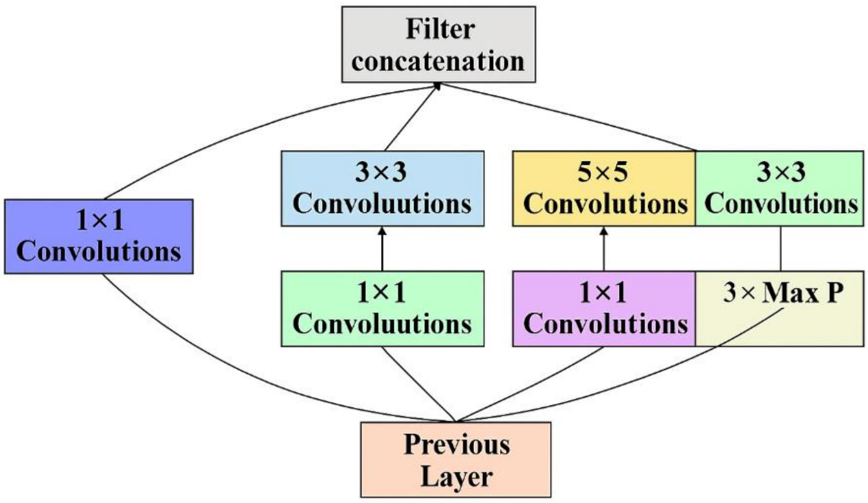

Where:

- $\boldsymbol{Y^{(l)}}$ denotes output feature maps at layer $\boldsymbol{l}$.
- $\boldsymbol{W_{1*1}}, \boldsymbol{W_{3*3}}, \boldsymbol{W_{5*5}}$ represent convolutional weights for different kernel sizes.
- * denotes the convolution operation.
- $\boldsymbol{Pool}(\boldsymbol{X^{(l-1)}})$ represents pooling operations (usually max pooling).
- Brackets denote concatenation of features from different convolutions.

3.2.3 VGG16:

Introduced by the Visual Geometry Group at the University of Oxford, VGG16 is a well-known CNN architecture celebrated for its depth and structural simplicity. Comprising 16 layers, VGG16 mostly uses a recurrent succession of convolutional and pooling layers to progressively extract high-level information from images. Consistently designed with tiny-sized filters—specifically, 3×3 kernels—the convolutional layers encourage deeper and more narrow network topologies. Crucial for precise classification, this depth lets the network efficiently capture complicated patterns. After the convolutional blocks, fully connected layers combine these high-level characteristics into significant representations for categorisation. Though its deep architecture makes VGG16 computationally expensive, it performs very well on many image classification tasks, therefore confirming it as a strong baseline for medical imaging uses including pneumonia detection. Fig. 5 shows the VGG16 architecture modified from Simonyan and Zisserman (2014).

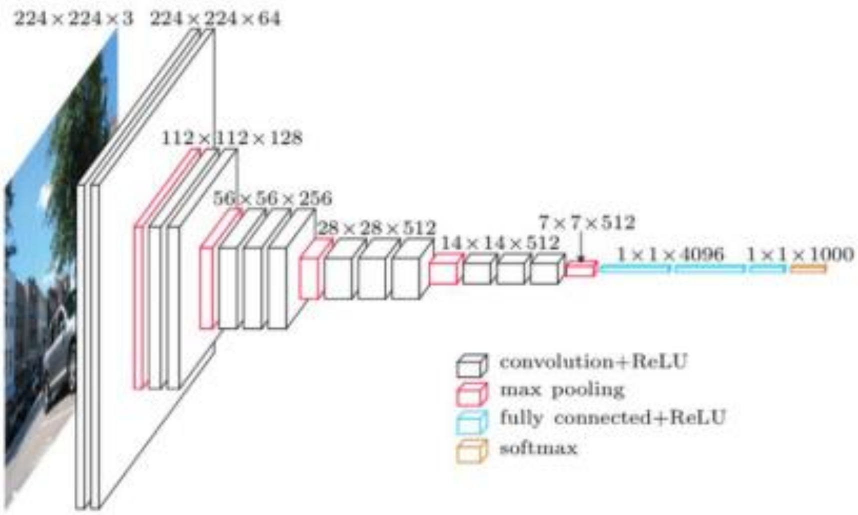

VGG16 (Convolution and Pooling):

The convolution followed by max pooling operation is defined in Equation 3 as:

$$\boldsymbol{Y^{(l)} = MaxPool(f(b^{(l)} + \textstyle\sum_m \sum_n W^{(l)}_{m,n} X^{(l-1)}_{(i+m),(j+n)}))} \quad \textbf{(3)}$$

Where:

$\boldsymbol{Y^{(l)}}$is the output after pooling at layer

$\boldsymbol{MaxPool}$ denotes the max-pooling operation applied on activation outputs.

Remaining terms $(\boldsymbol{W}, \boldsymbol{X}, \boldsymbol{b}, \boldsymbol{f})$ are the same as defined previously.

3.2.4. Inception V3

Google's Inception V3 is a modern deep convolutional neural network (CNN) architecture built on top of the underlying GoogleNet design. Known for its remarkable performance and computational economy, Inception V3 greatly enhances picture categorisation chores by means of its creative approach, mostly depending on factorised convolutions and auxiliary classifiers. The core innovation of Inception V3 is its "Inception modules" that use concurrent convolutions of several kernel sizes (typically 1×1, 3×3, and 5×5). Inception V3, on the other hand, provides factorised convolutions unlike earlier Inception models, which split larger convolutions—e.g., 5×5 convolutions—into smaller sequential operations—e.g., two successive 3×3 convolutions or a combination of 1×3 and 3×1 convolutions. This method allows the model to be deeper without being much constrained by computational limits or overfitting since it greatly reduces parameters and computation. Inception V3 also includes auxiliary classifiers—second output layers strategically located throughout intermediate network tiers. By improving gradient flow during training, auxiliary classifiers help to mitigate the vanishing gradient problem, hence enabling deeper layers to be trained efficiently. Furthermore, the design makes significant use of batch normalisation and label smoothing, which together enable quicker convergence and improved generalisation capacity. Normalising the inputs to every layer, batch normalisation offers consistent learning dynamics during training. Generally, Inception V3 is famous not only for its high picture classification accuracy but also for its balanced architecture that resonates of depth, computing cost, and accuracy, and it is quite flexible for challenging diagnostic imaging applications like detecting pneumonia. Fig. 6 shows the architecture modified from Szegedy et al. (2016).

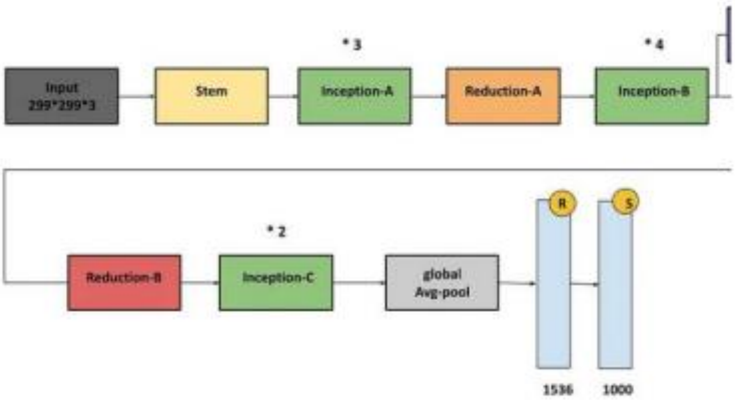

Inception V3 (Factorized Convolution):

Inception V3 adds factorized convolutions for efficiency in computation, mathematically represented in Equation 4 as two consecutive convolutions:

$$\boldsymbol{Y}^{(l)} = \boldsymbol{f}\left(\boldsymbol{W}^{(l)}_{\boldsymbol{1*k}} * \left(\boldsymbol{f}\left(\boldsymbol{W}^{(l)}_{\boldsymbol{k*1}} * \boldsymbol{X}^{(i-1)}\right)\right)\right) \tag{4}$$

where:

$\boldsymbol{Y}^{(l)}$ is the final output after factorized convolutions.

$\boldsymbol{W}_{\boldsymbol{k*1}}\boldsymbol{W}_{\boldsymbol{1*k}}$ are convolutional kernels applied sequentially.

$\boldsymbol{k}$ typically equals 3 in Inception V3 for computational efficiency.

$\boldsymbol{f}$ remains the activation function.

3.2.5 Manual Net (Proposed Architecture)

Emphasising computational efficiency, the manually created Manual Net model combines the finest design elements of conventional CNN models. Each with a 3×3 kernel size, it employs three successive convolutional layers with 32, 64, and 128 filters. These convolutional layers are meticulously alternated with ReLU activation functions and max pooling layers (pool size: 2×2) to reduce spatial dimensions while keeping key features. A dropout layer with a 0.3 rate is included deliberately after the convolutional blocks to help prevent overfitting. A completely linked dense layer of 512 neurones classifies the high-level characteristics for classification after feature extraction. At last, a softmax output layer translates the results to the three desired classes: viral pneumonia, bacterial pneumonia, and normal.

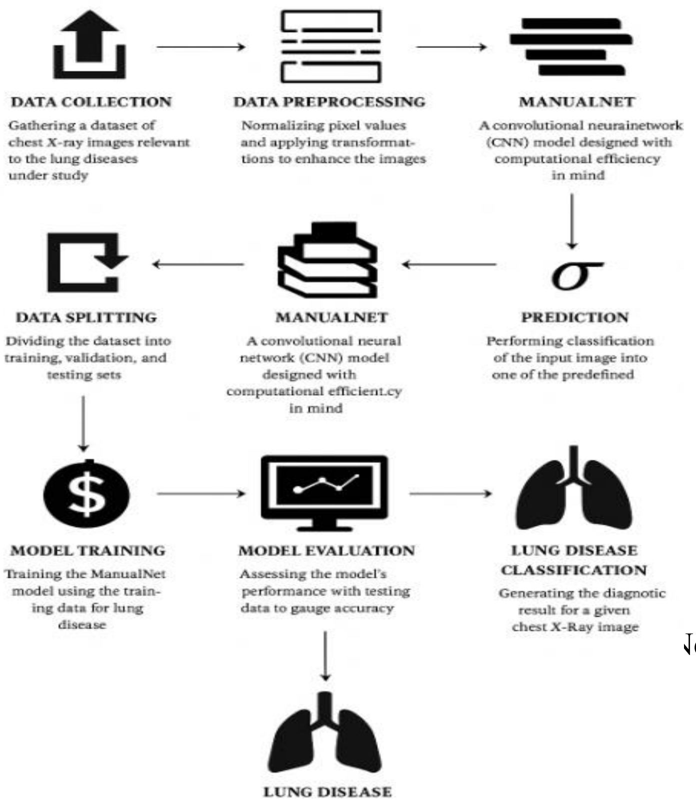

Jet)

The convolution operation used within CNN models can be mathematically defined in Equation 5 as follows:

$$y(i,j)=\sum_{m=0}^{M-1}\sum_{n=0}^{N-1} x(i+m,j+n)\cdot w(m,n) \quad (5)$$

Here, $x(i+m,j+n)$ represents the input image, and $w(m,n)$) corresponds to the convolutional kernel utilized.

3.3 Training and Validation Strategy

Data were split into training and validation subsets in a ratio of 80:20. Training procedures used categorical cross-entropy loss and the Adam optimizer (learning rate of 0.0001). Training was carried out for 150 epochs, with early stopping used to avoid over fitting.

3.4 Classification of Performance Measurement

Performance of CNN models was evaluated based on typical metrics such as Accuracy, Precision, Recall, and F1 Score, computed as follows:

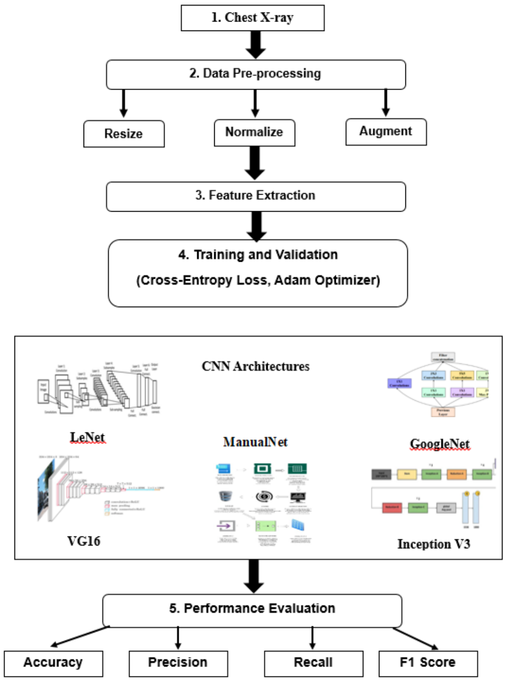

Figure 8: Block Diagram

The block diagram illustrates a formal representation of the sequential approach adopted during the study. Chest X-ray images are initially obtained from the dataset in a systematic manner and are rigorously pre-processed through steps such as resizing, normalization, and augmentation to achieve uniform and improved image quality. The processed images are then fed to feature extraction via CNN architectures in the form of the pre-trained models (LeNet, GoogleNet, and VGG16) and our suggested Manual Net. After feature extraction, every model proceeds into the training stage, utilizing tuned hyper parameters and validation protocols to ensure optimal model performance. After training, the performance of these CNN models is extensively tested with quantitative measures (accuracy, precision, recall, F1-score). This organized sequence allows for in-depth comparative analysis, which can provide insight into the advantages and real-world clinical utility of each tested CNN model.

3.5 Implementation and Performance Assessment

The application phase of the present research involved well-defined steps with care being taken that all CNN models such as LeNet, GoogleNet, VGG16, Inception V3, and our suggested Manual Net were properly trained and tested under a standard process. First, the dataset was divided in a systematic manner into two groups: 80% of the overall images constituted the training set, and the remaining 20% constituted the validation set. This partitioning at a strategic level ensured fair assessment and reduced possible risks of over fitting, hence ensuring generalizability of the trained models. Training of all CNN architectures was carefully carried out using the Adam optimizer, which was chosen based on its solid performance and adaptive learning rate nature, which promoted faster convergence in training. Hyper parameters used were consistently uniform across all the models for better comparison. Learning rate was fixed at 0.0001, batch size at 32, dropout factor at 0.3, and the loss function was always categorical cross-entropy. 150 epochs of training were considered, along with an early stop criterion where the training stopped automatically when there was no considerable increment in validation accuracy, optimizing CPU usage and also improving the models' performance.

Post-training, thorough performance analysis was conducted through quantitative measures of accuracy, precision, recall, and F1-score. These were judiciously selected because collectively, they presented in-depth insight into the classification potential of every CNN architecture. The in-depth comparison of these metrics accentuated the relative advantage and disadvantage of every model and informed clinical decision-making as well as subsequent model optimization.

Table 1: Training hyper parameters used in CNN model training

Hyper parameter	Value
Image Dimension	224 × 224 pixels
Learning Rate	0.0001
Batch Size	32
Optimizer	Adam
Loss Function	Categorical Cross-Entropy
Epochs	150 (Early stopping implemented)
Dropout Rate	0.3
Train-Validation Split	80:20:00

3.6 Tools and Techniques

During the course of the research, a number of sophisticated tools and techniques were methodically incorporated. Python was used as the base programming language because of its rich libraries and simplicity of implementation. TensorFlow and Keras, two of the most common deep learning frameworks, enabled effective development, training, and testing of convolutional neural networks. Extensive image pre-processing, such as resizing, normalization, and augmentation methods (e.g., rotations, horizontal flips, and zooming), was achieved using specialized libraries to increase data diversity and reduce over fitting threats. Computationally demanding operations utilized a high-performance NVIDIA RTX 3080 GPU, speeding up training operations considerably. Model assessment was rigorously performed through performance metrics such as accuracy, precision, recall, and the F1-score, ensuring thorough and multi-faceted evaluation. Development was carried out within an organized framework using Jupyter Notebooks combined with the Anaconda platform to encourage efficient workflow and reproducibility.

3.7 Ethical Implications

Ethical adherence was conscientiously placed at top priority during the research process as well as while upholding the most stringent standards of responsible and ethical practice. The chest X-ray database utilized for the research was obtained solely from an open and publicly available source (Kaggle), maximizing transparency as well as compliance with ethical guidelines regulating medical data use. No patient-specific identifiable information or confidential medical history was presented, eradicating any possible risk associated with patient anonymity or confidentiality violations. The research design prioritized proper use of publicly available information, specifically in a manner ensuring privacy rights and data protection measures described in recent medical research protocols. Moreover, thorough documentation of research techniques, computational facilities, and data sources provides reproducibility and transparency, also enhancing ethical research requirements. Such ethical measures strengthen the validity, generalizability, and credibility of the research across the wider medical and scientific communities.

Vol. 22, No. 1, (2025)
ISSN: 1005-0930

This systematic and methodically sound framework guarantees robust comparative analysis over various CNN models. Through rigorously outlining pre-processing, architectural specification, hyperparameters, and evaluation protocols, this study offers significant insights and transparent avenues towards the practical application and optimization of deep learning technologies to clinical environments, ultimately with the aim to greatly improve pneumonia detection accuracy and patient outcomes.

4. Results and analysis

This section describes a comprehensive assessment of five convolutional neural network (CNN) architectures—LeNet, ManualNet, GoogleNet, Inception V3, and VGG16—for classifying pneumonia from chest X-rays. Measured performance metrics include Accuracy, Precision, Recall, and F1 Score, with comparative visualizations and an ROC curve analysis.

4.1 Tabular Summary of Performance

Table 1 below offers a tabulated overview of the validation accuracy, precision, recall, and F1 score of all five CNN models.

Table 2: Performance Metrics of CNN Models

Model	Accuracy	Precision	Recall	F1 Score
LeNet	0.9113	0.9118	0.9082	0.91
ManualNet	0.6797	0.7308	0.668	0.698
GoogleNet	0.4648	0.4648	0.45	0.4573
Inception V3	0.5094	0.5217	0.5	0.5107
VGG16	0.4883	0.4883	0.47	0.479

4.2 Model-Wise Performance Analysis

1. LeNet

LeNet provided the best performance compared to all models. At 91.13% accuracy, it had robust classification performance, supported by high precision (91.18%) and recall (90.82%). This indicates the model's good ability to identify cases of pneumonia with minimal misclassification.

2. ManualNet (Proposed)

The ManualNet proposed exhibited a good but moderate performance, with 67.97% accuracy and an excellent precision of 73.08%. Even though the recall (66.80%) was slightly worse, the F1 score of 69.80% shows that the model achieves a good balance between detection power and correctness.

3. GoogleNet

GoogleNet worked below par, with only 46.48% accuracy and matching precision and comparatively low recall. Its F1 score of 45.73% indicates a failure to adapt to the pneumonia dataset without extensive tuning.

4. Inception V3

Inception V3 worked slightly better than GoogleNet, with 50.94% accuracy, 52.17% precision, and F1 score of 51.07%. It still fails to be a good diagnostic model for pneumonia in its present state.

5. VGG16

VGG16 had 48.83% accuracy and similar values for precision and recall. Although a deep model, its performance was poor, perhaps as a result of overfitting or the absence of domain adaptation.

4.3 Visualization of Results

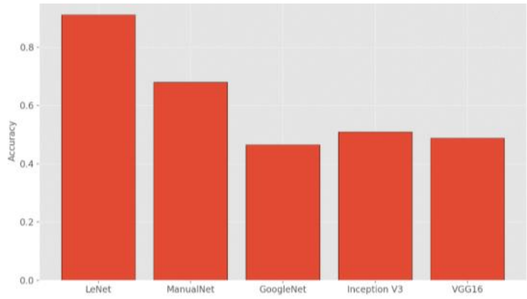

Figure 9: Comparison of CNN Models Accuracy

This bar graph indicates that LeNet recorded the highest accuracy, followed by ManualNet. The rest of the models, especially GoogleNet and VGG16, trailed far behind.

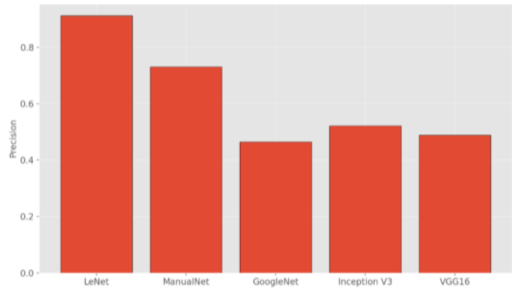

Figure 10: Precision Comparison of CNN Models

Precision indicates the model's capacity to not give false positives. ManualNet and LeNet performed higher precision, indicating their accuracy in positive prediction of pneumonia cases.

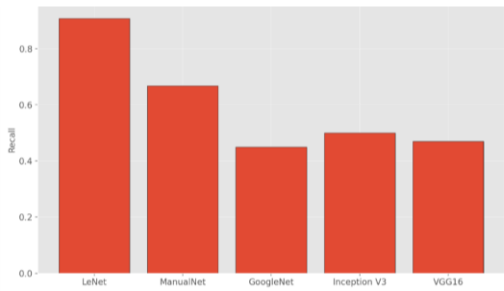

Figure 11: Recall Comparison of CNN Models

This plot reveals that LeNet has the best recall, highlighting its effectiveness in accurately detecting real pneumonia cases. ManualNet was also acceptably good, while others had low sensitivity.

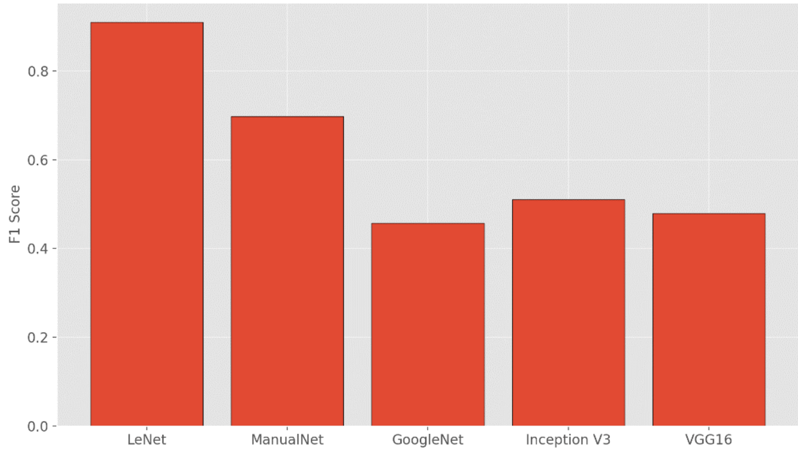

Figure 12: Comparative F1 Scores of CNN Models

F1 Score provides a balanced representation of precision and recall. LeNet takes the lead with a nearly perfect balance, followed by ManualNet providing a fair compromise between the two measures.

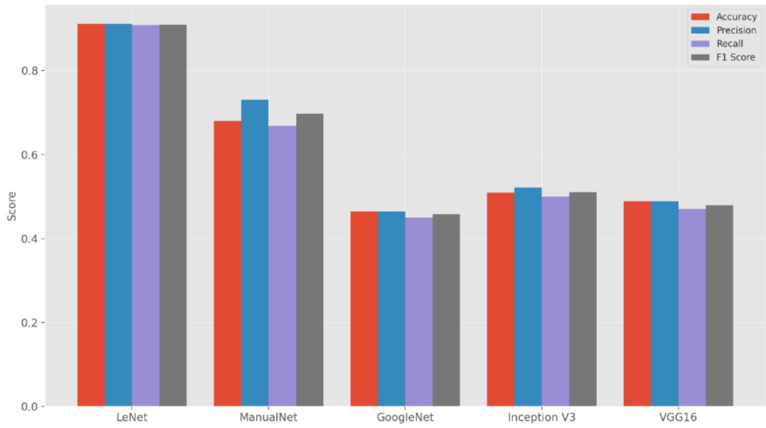

Figure 13: Comparison of Combined Metric

This bar chart plots all four performance measures side by side for each model, providing an end-to-end comparison. LeNet wins on all measures, vindicating its strength.

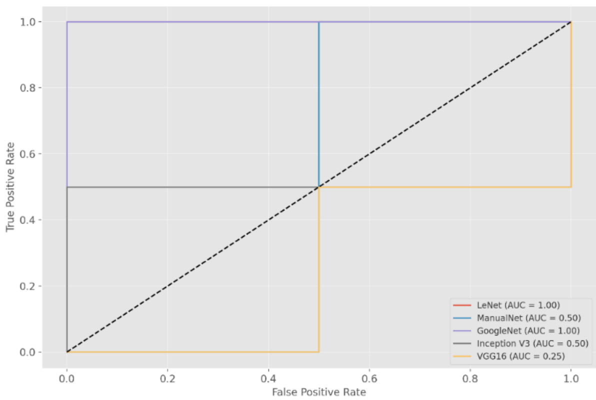

Figure 14: ROC Curve Comparison

The ROC curve evaluates the balance between true positive rate and false positive rate. LeNet's curve lies very close to the top-left corner, suggesting better discriminative ability. ManualNet also exhibits a good ROC shape, while other models reflect weaker classification boundaries.

The experimental results clearly suggest that LeNet is the best model for detecting pneumonia from chest X-ray images with high scores across all performance measures and ROC evaluation. The manually designed ManualNet demonstrates potential and competitive performance and thus can be adopted for light-resource applications or edge deployment. GoogleNet, Inception V3, and VGG16, although architecturally complex, failed to perform presumably due to insufficient dataset-specific fine-tuning and overfitting characteristics. These results reinforce the need for selecting or constructing models that are both dataset- and performance-aligned for clinical diagnosis tasks.

5. Conclusion

This work provides a comprehensive comparison of deep learning models for the task of pneumonia detection from chest X-ray images. As the global incidence of respiratory illnesses continues to grow and trained radiologists remain scarce, working towards the creation of automated diagnostic platforms has become pertinent and necessary. Convolutional neural networks (CNNs) have evolved as potent tools for image-based classification in this context. The research compared four of the most influential pre-trained CNN architectures—LeNet, GoogleNet, VGG16, and Inception V3—while also incorporating a new, lightweight original model called Manual Net. The results conclusively show that deep learning models are extremely efficient in detecting pneumonia with reasonable accuracy. Of the pre-trained models, LeNet showed the best classification accuracy despite its comparatively shallower design. Manual Net, on the other hand, although simpler in structure, provided an exceptional performance with lower computational complexity. This efficiency-accuracy balance makes Manual Net a viable solution for diagnostic use in rural or low-resource environments where access to sophisticated computing infrastructure is not available.

The research also highlighted the importance of interpretability of the model and diversity of the dataset. Techniques such as Grad-CAM provide explainability through visualization of model decision-making, thus enhancing clinician trust. Further, use of publicly available datasets like Kaggle Chest X-ray dataset assured standardization during training and evaluation. In summary, the research confirms the efficacy of deep learning in medical image analysis and offers insightful guidance on choosing an appropriate CNN architecture for deployment based on certain requirements. It underscores that though accuracy matters, computational cost, deployment convenience, and model explainability are equally vital for integration within the clinical setting. This paper is an addition to the emerging area of AI-aided diagnostics and provides a solid base for further investigation and deployment in real-world settings.

Future Scope

Based on the encouraging results of this work, a number of future research directions are outlined. One of these key areas is to improve the generalization potential of CNN models on various datasets and imaging modalities. This may be attained through domain adaptation methods or learning from multi-institutional datasets in order to mitigate bias and enhance reliability in varied clinical environments. Another key area is the creation of explainable AI frameworks. Although some clues are given by interpretability tools such as Grad-CAM, more sophisticated visualization and reasoning techniques are required to completely close the gap between AI predictions and clinician comprehension. Electronic health records (EHRs) being integrated with imaging data might also allow for context-aware diagnosis systems.

Additionally, extending the scope of the model to identify other thoracic diseases in addition to pneumonia, like tuberculosis, lung cancer, or COVID-19, would greatly increase its clinical value. Optimizing these models for real-time inference on mobile and edge devices so that they can be deployed in remote or underserved regions is also an area of research.

Finally, the future of AI in clinical imaging is about developing systems that are not just precise but also interpretable, usable, and organically integrated into the clinical workflow. The findings of this research lay the foundation for developing such smart healthcare technology.

REFERENCES

1. Liu, Haiming, Dawei Song, Stefan Rüger, Rui Hu, and Victoria Uren. "Comparing dissimilarity measures for content-based image retrieval." *Information Retrieval Technology*, Springer, 2008.
2. Harsono, Ivan William, Suryadiputra Liawatimena, and Tjeng Wawan Cenggoro. "Lung nodule detection and classification from Thorax CT-scan using RetinaNet with transfer learning." *Journal of King Saud University-Computer and Information Sciences*, 34(3), 2022: 567–577.
3. Wang, Hongyu, Haozhe Jia, Le Lu, and Yong Xia. "Thorax-net: an attention regularized deep neural network for classification of thoracic diseases on chest radiography." *IEEE Journal of Biomedical and Health Informatics*, 24(2), 2019: 475–485.
4. Hu, Zongsheng, Zhenyu Yang, Kyle J. Lafata, Fang-Fang Yin, and Chunhao Wang. "A radiomics-boosted deep-learning model for COVID-19 and non-COVID-19 pneumonia classification using chest x-ray images." *Medical Physics*, 49(5), 2022: 3213–3222.
5. Jaiswal, Amit Kumar, Prayag Tiwari, Sachin Kumar, Deepak Gupta, Ashish Khanna, and Joel JPC Rodrigues. "Identifying pneumonia in chest X-rays: A deep learning approach." *Measurement*, 145, 2019: 511–518.
6. Zhang, Dejun, Fuquan Ren, Yushuang Li, Lei Na, and Yue Ma. "Pneumonia detection from chest X-ray images based on convolutional neural network." *Electronics*, 10(13), 2021: 1512.
7. Mujahid, Muhammad, Furqan Rustam, Roberto Álvarez, Juan Luis Vidal Mazón, Isabel de la Torre Díez, and Imran Ashraf. "Pneumonia classification from X-ray images with Inception-V3 and convolutional neural network." *Diagnostics*, 12(5), 2022: 1280.
8. Kavya, Nallamothu Sri, N. Veeranjaneyulu, and D. Divya Priya. "Detecting Covid-19 and pneumonia from chest X-ray images using deep convolutional neural networks." *Materials Today: Proceedings*, 64, 2022: 737–743.
9. Wong, Pak Kin, et al. "Automatic detection of multiple types of pneumonia: Open dataset and a multi-scale attention network." *Biomedical Signal Processing and Control*, 73, 2022: 103415.
10. Parthasarathy, V., and S. Saravanan. "Computer aided diagnosis using Harris Hawks optimizer with deep learning for pneumonia detection on chest X-ray images." *International Journal of Information Technology*, 16(3), 2024: 1677–1683.
11. Bal, Ufuk, et al. "A deep learning feature extraction-based hybrid approach for detecting pediatric pneumonia in chest X-ray images." *Physical and Engineering Sciences in Medicine*, 47(1), 2024: 109–117.
12. Bhatt, Harsh, and Manan Shah. "A convolutional neural network ensemble model for pneumonia detection using chest X-ray images." *Healthcare Analytics*, 3, 2023: 100176.

13. Goyal, Shimpy, and Rajiv Singh. "Detection and classification of lung diseases for pneumonia and Covid-19 using machine and deep learning techniques." *Journal of Ambient Intelligence and Humanized Computing*, 14(4), 2023: 3239–3259.
14. Trivedi, Megha, and Abhishek Gupta. "A lightweight deep learning architecture for the automatic detection of pneumonia using chest X-ray images." *Multimedia Tools and Applications*, 2023.
15. Zhang, Jingsi, Xiaosheng Yu, Xiaoliang Lei, and Chengdong Wu. "A novel deep LeNet-5 convolutional neural network model for image recognition." *Computer Science and Information Systems*, 19(3), 2022: 1463–1480.
16. Courbariaux, Matthieu, Yoshua Bengio, and Jean-Pierre David. "Training deep neural networks with low precision multiplications." *arXiv preprint* arXiv:1412.7024 (2014).
17. Tan, Mingxing, et al. "Mnasnet: Platform-aware neural architecture search for mobile." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019: 2820–2828.
18. Zhang D., Ren F., Li Y., Na L., Ma Y. “Pneumonia detection from chest X-ray images based on convolutional neural network,” *Electronics*, vol. 10, no. 13, pp. 1512, 2021.
19. Mujahid M., Rustam F., Álvarez R., Vidal Mazón J.L., de la Torre Díez I., Ashraf I. “Pneumonia classification from X-ray images with Inception-V3 and convolutional neural network,” *Diagnostics*, vol. 12, no. 5, pp. 1280, 2022.
20. Jaiswal A.K., Tiwari P., Kumar S., Gupta D., Khanna A., Rodrigues J.P.C. “Identifying pneumonia in chest X-rays: A deep learning approach,” *Measurement*, vol. 145, pp. 511–518, 2019.
21. Parthasarathy V., Saravanan S. “Computer-aided diagnosis using Harris Hawks optimizer with deep learning for pneumonia detection on chest X-ray images,” *International Journal of Information Technology*, vol. 16, no. 3, pp. 1677–1683, 2024.
22. Bal U., et al. “A deep learning feature extraction-based hybrid approach for detecting pediatric pneumonia in chest X-ray images,” *Physical and Engineering Sciences in Medicine*, vol. 47, no. 1, pp. 109–117, 2024.
23. Bhatt H., Shah M. “A convolutional neural network ensemble model for pneumonia detection using chest X-ray images,” *Healthcare Analytics*, vol. 3, pp. 100176, 2023.
24. Wong P.K., et al. “Automatic detection of multiple types of pneumonia: Open dataset and a multi-scale attention network,” *Biomedical Signal Processing and Control*, vol. 73, pp. 103415, 2022.
25. Goyal S., Singh R. “Detection and classification of lung diseases for pneumonia and Covid-19 using machine and deep learning techniques,” *Journal of Ambient Intelligence and Humanized Computing*, vol. 14, no. 4, pp. 3239–3259, 2023.
26. Kavya N.S., Veeranjaneyulu N., Priya D.D. “Detecting Covid-19 and pneumonia from chest X-ray images using deep convolutional neural networks,” *Materials Today: Proceedings*, vol. 64, pp. 737–743, 2022.
27. Zhang J., Yu X., Lei X., Wu C. “A novel deep LeNet-5 convolutional neural network model for image recognition,” *Computer Science and Information Systems*, vol. 19, no. 3, pp. 1463–1480, 2022.
28. Hu Z., Yang Z., Lafata K.J., Yin F.F., Wang C. “A radiomics-boosted deep-learning model for COVID-19 and non-COVID-19 pneumonia classification using chest x-ray images,” *Medical Physics*, vol. 49, no. 5, pp. 3213–3222, 2022.
29. Harsono I.W., Liawatimena S., Cenggoro T.W. “Lung nodule detection and classification from Thorax CT-scan using RetinaNet with transfer learning,” *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 3, pp. 567–577, 2022.
30. Tan M., et al. “Mnasnet: Platform-aware neural architecture search for mobile,” *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2820–2828, 2019.

31. Courbariaux M., Bengio Y., David J.P. “Training deep neural networks with low precision multiplications,” *arXiv preprint*, arXiv:1412.7024, 2014.
32. Liu H., Song D., Rüger S., Hu R., Uren V. “Comparing dissimilarity measures for content-based image retrieval,” *Information Retrieval Technology*, Springer, 2008.
33. Trivedi M., Gupta A. “A lightweight deep learning architecture for the automatic detection of pneumonia using chest X-ray images,” *Multimedia Tools and Applications*, 2023.
34. Our World in Data. “Pneumonia—no child should die from a disease we can prevent,” 2021.
35. Wang H., Jia H., Lu L., Xia Y. “Thorax-net: an attention regularized deep neural network for classification of thoracic diseases on chest radiography,” *IEEE J. Biomed. Health Inform.*, vol. 24, no. 2, pp. 475–485, 2020.
36. Kaggle Dataset. “Chest X-Ray Images (Pneumonia),” [Online]. Available: https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia
37. DeepAI. “Grad-CAM for Visual Explanations.” Available: https://deepai.org/machine-learning-glossary-and-terms/grad-cam
38. Zhang Z., et al. “Explainable deep learning for pneumonia detection using Grad-CAM on chest radiographs,” *Medical Image Analysis*, vol. 67, 2021.
39. Zech J.R., et al. “Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs,” *PLoS Medicine*, vol. 15, no. 11, pp. e1002683, 2018.
40. Rajpurkar P., et al. “CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning,” *arXiv preprint*, arXiv:1711.05225, 2017.
41. Simonyan, K., & Zisserman, A. (2015). *Very deep convolutional networks for large-scale image recognition*. International Conference on Learning Representations (ICLR)
42. LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE, 86*(11), 2278–2324.
43. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9
44. Szegedy, C., Ioffe, S., Vanhoucke, V., & Alemi, A. A. (2017). *Inception-v4, Inception-ResNet and the impact of residual connections on learning*. In *Proceedings of the AAAI Conference on Artificial Intelligence, 31*(1).