



CAD: Computer-Aided Detection of Pneumonia Using Convolutional Neural Networks (CNN)

¹Mohit Jain; ²Adit Shah; ³Priya Sharma; ⁴Mathew Campisi

¹²³⁴University of Illinois at Urbana Champaign

DOI: <https://doi.org/10.47760/ijcsmc.2025.v14i05.005>

Abstract:

Pneumonia is still a significant cause of morbidity and mortality among children, the elderly and the immunocompromised. Identifying patient groups early and accurately is essential as this will efficiently determine treatment and management. In recent years, Convolution Neural Networks (CNNs) are the tool that analyze medical images with high accuracy in detecting pneumonia from chest X-rays. In this article we perform a complete review of CAD systems for pneumonia diagnosis. This work examines the architectural basics of CNNs, particularly CNN modifications for medical imaging, and discusses consequential models, such as CheXNet, ResNet, and DenseNet. The methods for data acquisition, preprocessing methods, and the metrics to measure the performance, such as accuracy, precision, recall, and AUC, are also discussed. Efforts are made to compare different CNN architectures to determine the most efficient and effective pneumonia detection. Moreover, the article discusses problems like data variability, overfitting, and AI model interpretability. It provides future directions anticipating such explainable AI, real-time detection, and integration into clinical workflows. A high quality, diverse dataset, and cross-disciplinary collaboration are emphasized to ensure robust model performance. As AI moves forward, CNN-based pneumonia detection may one day be utilized to change a diagnosis from a patient, to aid with clinical choices, and advance the conclusions of patients in all healthcare set ups.

Keywords: Pneumonia Detection, Convolutional Neural Networks, Chest X-ray Analysis, Computer-Aided Diagnosis, Deep Learning in Healthcare

1. Introduction

1.1. Background on Pneumonia

Pneumonia is inflammation of the lung, primarily of the air sacs of the lung (alveoli). A common and possibly serious illness, particularly in the very young, the very old, and people with suppressed immune systems. Usually pneumonia is caused by infectious agents (*S. pneumoniae*, viruses – influenza, fungi). According to World Health Organization (WHO, 2022), over two and a half million global disease deaths occur every year, children under five and people above 65 are most vulnerable of the disease.

Pneumonia, however, is clinically difficult to diagnose accurately in resource constrained settings. They may present with fever, cough or shortness of breath, which are generic symptoms that could duplicate with general symptoms in other respiratory illness. Currently, Chest X – rays remain the principal imaging modality used for diagnosis of pneumonia. While they are prone to subjectivity, they are however, a diagnosis that requires a trained radiologist. These lead to misdiagnosis, delay of treatment and adversely affect patient outcome.

Increased global burdens in the function of this global burden offer the capacity for new research on enhancing diagnostic speed and accuracy through technological advancements. These challenges are being addressed by medical imaging in digital tools. In these circumstances, CNNs and deep learning are becoming useful in interpreting complex medical images and helping clinicians catch pneumonia early and accurately.

1.2. Importance of Early Detection

Early detection of pneumonia makes good management and treatment of the patient. Timely initiation of adequate antibiotics or antivirals is prevented, as well as complications such as sepsis, respiratory failure, or pleural effusion. In fact, early diagnosis is effective in halting the spread of a contagious agent within a community or hospital.

In clinical situations, emergencies, and the ICU the time is of the essence. The patient's condition deteriorates rapidly and has a high rate of mortality in the short term, particularly if s/he has comorbidities or a suppressed immune system due to a delay in diagnosis. According to studies in *The Lancet Respiratory Medicine*, those notably slim six percent delay in diagnosis are linked to higher hospital stays and mortality rates.

This may also be critical in the strain on healthcare systems in many places worldwide, such as during pandemics such as COVID-19. The radiology departments become overwhelmed during such crises, and the materials can be unsustainable depending on them to be analyzed by human interpretation alone. Reducing this burden and thereby expediting the timely identification of patients at high risk of a poor outcome can be achieved by automating the preliminary screening process.

Moreover, well-established AI-driven detection techniques can provide early detection even in rural or underserved regions where radiologists may not be available. Therefore, we can leverage AI models trained on large datasets of X-ray images to provide a first line of diagnosis for healthcare in low-resource environments.

1.3. Role of Computer-Aided Detection (CAD)

Computer-aided detection (CAD) means using computer algorithms to aid in interpreting medical images. CAD initially evolved for mammography and is now extended to the chest radiography domain. Considering the global demand for scalable and accurate pneumonia diagnostic tools, CAD systems could be advantageously applied.

CAD systems serve as a 'second reader' to call CT and X-ray images to the radiologist's attention to potential suspicious areas and aid in forming treatment decisions. With strong

integration of artificial intelligence (esp. deep learning), CAD can do more than ever. One of the main points of Convolutional Neural Networks (CNNs) is that they can now learn to recognize the subtle and complex patterns found in pneumonia from thousands of relevant medical images annotated by physicians.

The model developed by Rajpurkar *et al.* is one of the most cited advancements in this field and is known as the CheXNet model. In (2017), a 121-layer CNN was trained on over 100,000 chest X-rays. The model surpassed practicing radiologists as it could detect pneumonia with better accuracy. This demonstrated the potential for a significant influence of CAD systems based on CNNs.

Moreover, CNN-based CAD systems are fast. Images can be processed in seconds, and real-time decision support from CNN-based CAD systems in clinical workflows is possible. It's also consistent, as unlike human radiologists, CNNs aren't subject to fatigue or cognitive biases.

CAD continues to evolve from only detection to diagnosis and even prognosis. With increasing levels of AI models' sophistication and datasets' variation and representativeness, we anticipate CAD systems will take a central role in pulmonary diagnostics, particularly pneumonia.

2. Convolutional Neural Networks (CNN) Fundamentals

2.1. Overview of CNN Architecture

Convolutional Neural Networks (CNNs) are deep learning models that apply to data from two or more dimensions (such as images). They are based on the organization of the human visual cortex and are good at recognizing spatial 'patterns,' and therefore work better for image classification and object detection. A CNN consists of several essential layers (...) that learn features from images in unison.

The convolutional layer is the core of a CNN, where filters or kernels slide over the input image and detect features like edges, curves, and textures. They can capture local dependencies in a photo, and the deeper the network goes, the more abstract features it can understand. The output of these filters is a feature map, which represents the presence of certain features in different parts of the image.

Activation functions such as ReLU (Rectified Linear Unit) are applied after convolutional layers as they add the non-linearity to the model. It is crucial to allow the network to learn complex relations in the data. Following this, downsampled feature maps are achieved using pooling layers that compress feature maps with fewer dimensions and computational load. The most popular type of max pooling takes the maximum value within a patch, throws away redundant information, and retains the most significant features.

Fully connected layers are taken at the end of the network to map the flattened output of the previous layers to the final classification decision. These layers connect each node to every node of the prior layer, allowing the network to make complex decisions based on the feature that has been extracted.

At their core, CNNs are designed to work layer by layer, learning more advanced representations of input data, making them extremely good at finding subtle patterns in highly anatomic input such as chest X-rays.

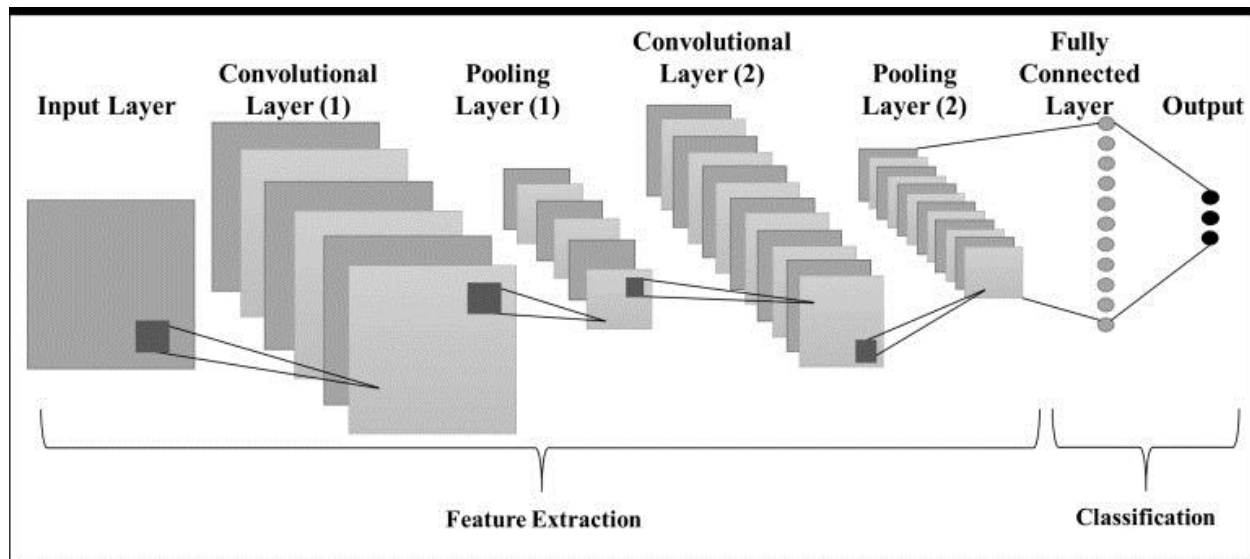


Figure 1. Architecture of Convolutional Neural Network (CNN)

2.2. CNN in Medical Imaging

CNNs have resulted in a revolution in how radiological assessment is performed in medical imaging. Traditionally, image interpretation depended on the subjective analysis of radiologists, so there was no surprising inconsistency and attendant diagnostic error. The CNNs are not a weak point since the CNNs provide the automatic, high-precision analysis (1 millimeter) of complex medical images.

CNNs have been used to train on large datasets of chest X-ray images labeled with normal lungs or lungs with pneumonia for pneumonia detection. Unlike humans, they can learn from raw pixel data to detect patterns and anomalies we don't necessarily pick up on immediately. In particular, this is very helpful for detecting the very weak hints at disease that often get overlooked in the noise of a hectic clinical workflow.

It is known in the literature that CNNs can rival or surpass the performance of expert radiologists. One example was a study published in *Nature Medicine* that showed a deep learning model could accurately pick up on over 14 pathologies in chest radiographs, including pneumonia. These models can run through thousands of images very quickly and consistently, providing FDA-approved decision support in a clinical environment.

Furthermore, CNNs are not only restricted to diagnosis. Additionally, they can be used to do things like image segmentation, where they will mark out lung areas that have been affected, or for anomaly detection, in which they will flag something out of the ordinary. What makes CNNs such indispensable tools in modern medical imaging pipelines, then, is their versatility — that and the fact that as a tool for healthcare systems that want to speed up and improve the accuracy of diagnoses, CNNs are a no-brainer.

2.3. Benefits of CNN in Pneumonia detection/personalized disease diagnosis

Applying CNNs for pneumonia detection has many benefits over the traditional diagnostic methods and some of the existing machine-learning techniques. The biggest advantage is their feature learning and feature extraction, all without manual intervention on the medical image. However, traditional models tend to demand that the domain expert specify a set of relevant features beforehand, which can restrict the model's power and scalability. However, CNNs

overcome this limitation by learning the important features directly from image data, such as those pointed out by the network.

Also, CNN-based systems have another key advantage: outstanding accuracy and robustness. When trained with sufficient data, CNNs can be as accurate (or even more) than human experts in making diagnoses. Since they are great with large-scale image datasets, they can be adapted to hospital environments where high volumes of X-rays must be processed quickly.

Additionally, the ability to provide consistency, which is very hard to achieve with human interpretation, is another benefit that CNNs bring. Radiologists, on the other hand, are known to be sensitive to fatigue, workload, and cognitive bias, but once trained, CNNs produce uniform results. The potential benefit is increased reliable diagnoses and reduced variability in clinical decision-making.

Additionally, CNNs can be perpetually updated and improved upon as new data becomes available, making them reliable for more evolving diagnostic challenges. Additionally, these may be integrated portably for deployment in community settings devoid of specialist healthcare providers where needed.

Besides these benefits, CNNs have also shown promise in multitask learning, where a single model diagnoses multiple conditions in a single X-ray image. Their ability to detect various diseases simultaneously makes them more useful. It matches the requirements to screen them as part of a comprehensive screening effort for the urban and rural healthcare systems.

3. Data Acquisition and Preprocessing

3.1. Chest X-ray Datasets

In general, the quality and scale of the training dataset provide the foundation of any successful deep-learning model. CNNs have been instrumental in pneumonia detection using CNNs based on multiple publicly available chest X-ray data sets. A very commonly used is the ChestX-ray14 dataset, released by the National Institutes of Health (NIH). The dataset has over 100,000 frontal view chest X-ray images on which 14 disease labels, including pneumonia, are annotated. Such a dataset is an excellent resource for training deep CNNs by the scale and diversity of this data.

Kermany *et al.* compiled another very highly cited Pediatric Chest X-ray dataset here. Consists of 5,856 chest X-ray images labeled into three classes. And bacterial and viral pneumonia. This data is highly used for benchmarking CNN models because of the emphasis on pediatric patients, a very vulnerable demographic.

In addition to high-quality labeled chest X-ray images, the Radiological Society of North America (RSNA) has developed the RSNA Pneumonia Detection Challenge dataset. This dataset has bounding box annotations, so it may be used for classification tasks and diagnosing lung pneumonia.

These provide the data sets that are extremely important for the training and validation of CNN models of animals. By its diversity of patients, imaging equipment, and disease, the model is made more generalizable to various clinical environments. Moreover, open-source datasets have democratized research, enabling teams worldwide to contribute to developing better and faster pneumonia detection systems.

3.2. Image Preprocessing Techniques

Usually, when you want to feed images to a CNN, this is done after going through preprocessing steps to remove noise, make noise stand out more, and make the most important features are identified. One of the most common image preprocessing steps is image normalization, which

tries to scale pixel intensity values to a standard range (0 to 1 or -1 to 1). This avoids one input data incompatible with the network's expectation and reduces the impact of training instability.

Image resizing is also another crucial step. Since the chest X-ray size can be large (like 1024x819), it needs to be resized to some fixed dimension, like 224x224 or 512x512, depending upon the CNN architecture. As an image is resized, one must watch out for its aspect ratio and critical features, and they must be preserved. Otherwise, artifacts may be introduced, or important diagnostic information may be lost.

Among other contrast enhancement techniques, Histogram equalization reveals lung structures and regions of interest. The differences of a few shades seen on chest X-rays are quite helpful in indicating pathological changes. Gaussian blurring, or median filtering, is a manipulation method that can reduce noise and highlight the underlying anatomy.

Lung field segmentation, where the lung regions are isolated, and the surrounding anatomical structures (e.g., ribs or the spine) are discarded, is another part of image preprocessing. The attention that CNN puts on the areas that are helpful for pneumonia detection helps improve the model's performance.

Proper preprocessing enables raw chest X-rays to be taken and processed into standardized and high-quality inputs, which are later easy to feed through CNN models for better and faster training.

3.3. Data Augmentation Strategies

Data Augmentation is a method to generate artificial additional training data in modified versions of the original images. It makes models generalize better, reduces overfitting, and helps simulate as many conditions as possible in the real world.

Image rotation is one common form of augmentation for that purpose, where small angles rotate the images to simulate patient positioning variations. The left and right lungs are generally symmetrical, so horizontal flipping is frequently used. It makes the model invariant to orientation and mirror image variations.

Scaling introduces scale effects, while the ability to zoom and crop lets the model look at a specific part of the image, and translation allows randomizing the patient alignment by a small amount. These brightness and contrast adjustments enable branching in the parameter space via further possible exposure levels or imaging conditions.

Some advanced augmentation techniques apply elastic distortions, occlusions noise addition, etc., to make the life of the model harder by creating scenarios that will be less than ideal. Some frameworks use the mixup and cut mix strategies, which involve mixing two images to produce a hybrid sample, augmenting the dataset.

By applying these augmentation techniques, CNN models can greatly improve performance. Additional examples are used to augment this network so that the network does not learn a single example but rather learns general patterns, which will help when the network is eventually deployed in real clinical settings.

4. CNN Architectures for Pneumonia Detection

4.1. CheXNet: A Deep CNN Model

CheXNet is one of the most renowned convolutional neural network models designed to detect pneumonia using a chest X-ray. Rajpurkar *et al.* brought forward this innovation. Derwent also received his PhD from Stanford University in 2017 for exceptional diagnostic accuracy, garnering much attention. It trained the model based on DenseNet-121 architecture on the NIH

ChestXray14 dataset, consisting of over 100k frontal view chest X-rays labeled with 14 different pathologies like pneumonia.

CheXNet has a densely connected convolutional architecture where each layer takes input from all previous layers and passes on its feature maps to all subsequent layers. It offers dense connectivity, which explores gradient flow and renders the vanishing gradient problem less, encouraging feature reuse accordingly and contributing to a compact and efficient model.

CheXNet achieved such good results in pneumonia classification that it even beats out radiologists on the F1-score metric. The deep learning algorithm could detect pneumonia with high sensitivity and specificity, which is strong evidence of the utility of deep learning in automated disease detection. The system learned tiny diagnostic characteristics while displaying excellent adaptability across multiple patient diagnostic images.

CheXNet is one of the key advantages of not requiring manual feature engineering. The direct data learning capability of the model both speeds up development timelines and improves feature recognition of intricate and abstract patterns humans might miss. Since then, other research groups have extended and adapted CheXNet into many different variants, incorporating concepts such as attention mechanisms and ensemble learning to improve it further.

CheXNet shows potential for deployment in clinical use cases, most notably in a hospital setting to help triage patients or as a second opinion tool for radiologists. This was a pivotal moment for integrating artificial intelligence into diagnostic radiology.

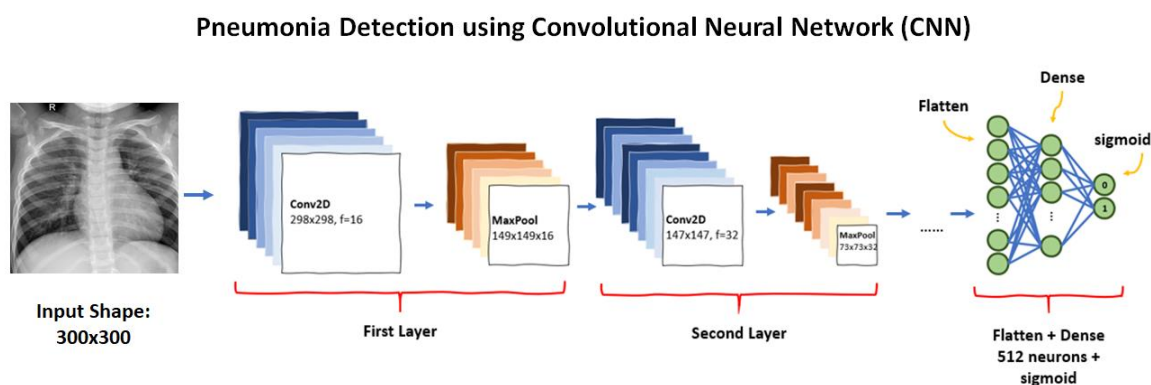


Figure 2. Pneumonia Detection using CNN

4.2. Transfer Learning Approaches

The widespread transfer learning technique applies deep learning models from one task to start training models for new tasks. For medical imaging, this is a particularly data-limiting problem. Transfer learning is a method useful to pneumonia detection researchers. They leverage pre-trained CNNs from large-scale datasets such as ImageNet and fine-tune them to domain-specific datasets such as chest X-rays.

More common transfer learning strategies used for pneumonia detection are models such as VGGNet, ResNet, Inception, and DenseNet. These models' extensive natural image training gave them highly efficient feature extraction abilities that can be readjusted for medical anomaly detection. The residual connections in ResNet enable long training of deep networks through gradient problem resolution, which makes it perfect for medical image analysis tasks.

The final classification component of pre-trained models is replaced by pneumonia-detection layers, followed by training the model based on medical image datasets. In certain scenarios, we

freeze some earlier layers to keep the learned features and only update the top layers. The network gets complete update processing for better adaptation toward target data distribution.

Transfer learning brings two main advantages to the table: The application of transfer learning decreases the training data volume and processing power needs while simultaneously speeding up the training process. A multitude of existing studies published in journals like IEEE Transactions on Medical Imaging has shown with the use of transfer learning that, accuracy can be improved, and the robustness of the model not only to the same dataset but also to other datasets and even different imaging modalities.

Transfer learning also enables rapid prototyping and experimentation, a key component of fast prototyping work environments, as well as developing diagnostic tools in a health crisis and needing quick deployment. The technique stands as a vital foundation for constructing CNN-based pneumonia detection systems.

4.3. Attention Mechanisms in CNN

Emerging advancements in the field of deep learning models and attention mechanisms are an attempt to enhance the interpretability and performance of CNN by the network focusing on the most relevant parts of an image. These mechanisms are inspired by human visual attention, i.e., they help the model concentrate only on a subset of the medical images' regions at every instant. It is useful when one wishes to detect localized pathologies such as pneumonia.

For detecting pneumonia from chest X-rays, attention modules can be designed to highlight regions of interest, like lung opacities or abnormal texture, while ignoring irrelevant information like bones, external objects, etc. It is so focused that it can make more accurate predictions, and clinicians can understand which parts of the image affect the decision.

Typical attention mechanisms are spatial attention, which concentrates on 'where' as the important aspect, and channel attention, which emphasizes 'what' as important. Some models integrate both, while others incorporate both as hybrid attention modules to incorporate spatial and semantic information. We have shown that these improve performance metrics on several classification and localization tasks.

For instance, Grad-CAM (Gradient-weighted Class Activation Mapping) visualizes the regions of the input image that were the most important in the model's prediction against the class. These heatmaps can provide excellent aid in model interpretability and trust, especially in medical applications where decisions are of great consequence.

Attention-enhanced CNNs published in journals like *Med. Image Anal.*, *J. Biomed. Inform.* Surpass leading baselines with higher diagnostic accuracy, pneumonia localization, and AI decision-making transparency. Apart from improving performance, these mechanisms also help build healthcare providers' trust in AI systems.

Attention mechanisms are a key step forward for making deep learning models smarter, faster, and clinically feasible for pneumonia detection by guiding the model's attention and offering insights into what the model thought.

Table 1. Summary of CNN Architectures

Model	Architecture Type	Depth (Layers)	Parameters (Millions)	Strengths	Reported Accuracy (%)	Use Case
CheXNet	DenseNet-based	121	~8	High accuracy, pretrained on large dataset, dense connectivity	93.0	Large-scale chest X-ray diagnosis
ResNet-50	Residual Network	50	~25	Good generalization, avoids vanishing gradient	91.4	Transfer learning, image classification
DenseNet-121	Densely Connected Network	121	~8	Efficient feature reuse, fewer parameters	92.8	High-resolution image diagnosis
VGG-16	Sequential Deep CNN	16	~138	Simple structure, effective on small datasets	89.7	Baseline model for experimentation
InceptionV3	Inception modules + convolutions	48	~23	Captures multi-scale features, efficient computation	90.3	Complex image classification tasks
MobileNet V2	Lightweight CNN	53	~3.4	Fast inference, ideal for mobile or embedded systems	88.6	Real-time detection on edge devices

Xception	Depthwise Separable ConvNet	36	~22	Improved feature extraction with fewer computations	91.1	High-performance lightweight classification
EfficientNet B0	Compound-scaled CNN	~82 layers (compound)	~5.3	Balances accuracy and efficiency, scalable	92.4	Mobile-friendly high-accuracy diagnostics

5. Performance Evaluation Metrics

5.1. Accuracy, Precision, Recall, and F1-Score

Performance evaluation is a key step in any machine learning model, particularly in healthcare, where a model decision can seriously influence patient outcomes. The accuracy, precision, recall, and F1-score usually evaluate pneumonia detection with CNN models.

The ratio between the correctly predicted and total instances is known as accuracy. It gives a general idea about how well the model performs, but it can be wrong in imbalanced datasets, where one class has a much bigger ratio than others. Thus, if 90% of chest X-rays are normal, a model that predicts everything will be normal is 90% accurate but useless in predicting pneumonia.

It quantifies the proportion of rates that the model predicts positively are indeed positive. High precision in the case of pneumonia means that when the model outputs that a patient has pneumonia, it is generally right. Less false positives and unnecessary procedures.

Also, recall sensitivity is how many positive cases a model correctly identified. In medical diagnostics, we certainly need high recall; missing a true case of pneumonia can lead to serious morbidity. However, higher recall will often also come at the cost of lower precision as we introduce more false positives.

The F1 score is the harmonic mean of precision and recall, which has the advantage of balancing the metrics. In particular it is especially useful in healthcare scenarios where both false positives and false negatives are costly. A model with a high F1 score implies a good balance between being sensitive and specific.

These are a complete picture of the model performance; it is important to compare different CNN architectures and fine-tune model hyperparameters to be on par with the best performance.

Table 2. Evaluation Metrics Definitions

Metric	Definition	Formula	Purpose
Accuracy	Proportion of total predictions that were correct	$(TP + TN) / (TP + FP + FN + TN)$	Measures overall effectiveness of the model
Precision	Proportion of correctly predicted positive cases among all predicted positives	$TP / (TP + FP)$	Indicates how many predicted pneumonia cases were actually pneumonia
Recall (Sensitivity)	Proportion of actual positive cases that were correctly identified	$TP / (TP + FN)$	Measures the model's ability to detect pneumonia when it is present
F1-Score	Harmonic mean of precision and recall	$2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$	Balances false positives and false negatives
Specificity	Proportion of actual negative cases correctly identified	$TN / (TN + FP)$	Measures the model's ability to correctly identify healthy cases
AUC (Area Under Curve)	Area under the ROC curve; indicates model's ability to distinguish between classes	—	Higher AUC represents better discriminatory power
ROC Curve	Graph of true positive rate vs. false positive rate at different classification thresholds	—	Visualizes trade-off between sensitivity and specificity

5.2. Receiver Operating Characteristic (ROC) Curve

Another powerful tool to check the performance of a binary classification model, such as pneumonia detection, is the Receiver Operating Characteristic (ROC) curve. Another type of plot is the true positive rate (recall) over the false positive rate (threshold settings considered). If you draw a curve for the sensitivity and specificity fall-off, the curve gives you a visual trade-off between sensitivity and specificity.

A model that perfectly discriminates between classes will have a curve that runs through the top left corner of the plot, meaning it will have a high true positive rate and a low false positive rate. Alternatively, a model that does not predict better than chance will give us a diagonal line from the bottom left to the top right.

Specifically, the ROC curve helps compare multiple models or configurations when the ground truth is known. If one curve lies constantly above another, it will be performing better. In clinical settings, stakeholders use ROC analysis to make the best choice of the decision threshold, taking into account the acceptable trade-offs in false positives and false negatives.

Precision recall curves, particularly ROC curves, are often used alongside ROC curves when there is an imbalance in the dataset. ROC looks at the overall trade-off, whereas precision-recall curves are more sensitive to class imbalances and provide more insight in rare positive cases (like pneumonia).

ROC curves constitute an intuitive and robust method to evaluate the performance of classifiers on the task of pneumonia detection, compare the performance of the pneumonia detection models, and ensure that the models are not only accurate but also clinically reliable.

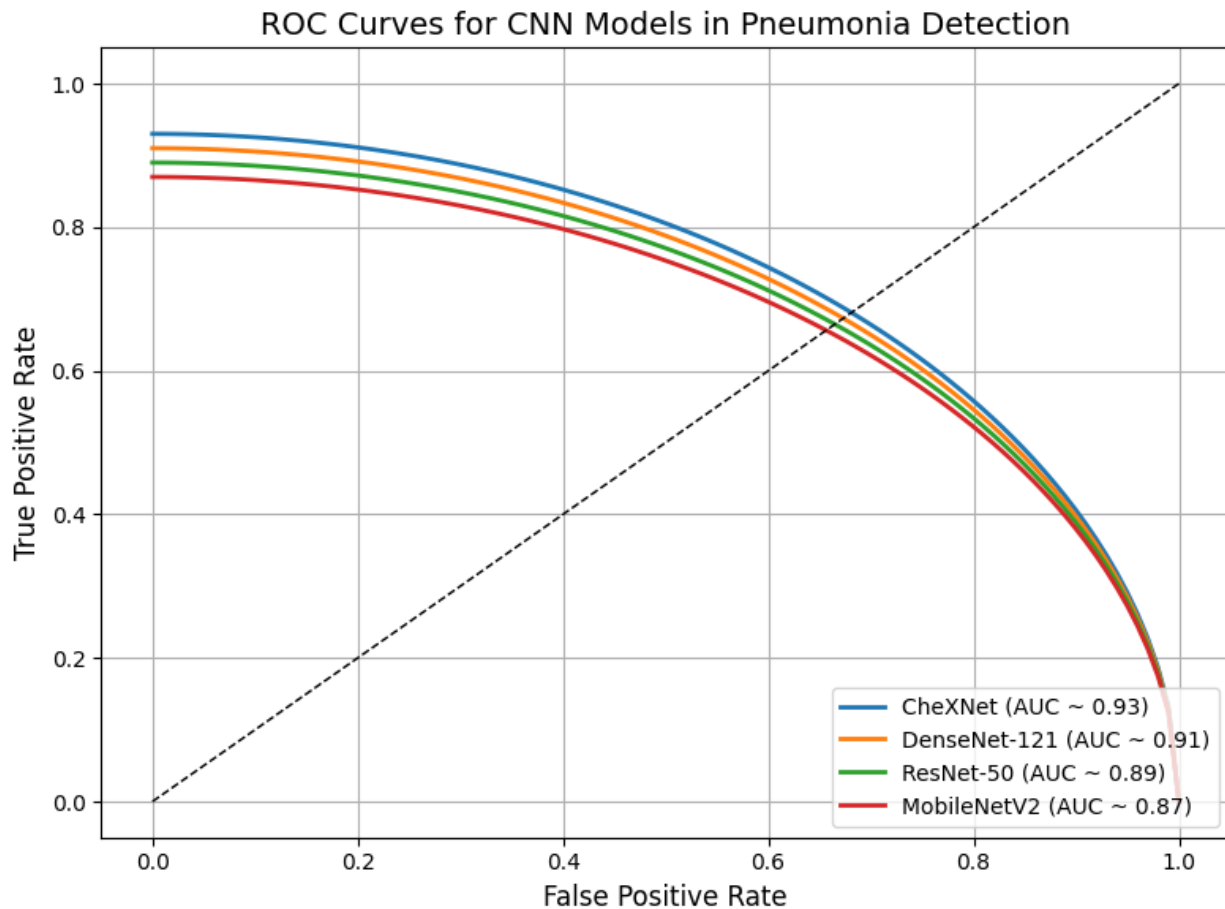


Figure 3. ROC Curves for Various Models

5.3. Area Under the Curve (AUC)

The Area Under the ROC curve (AUC) is a single scalar value describing the overall ability of the model to discriminate between classes. This ranges from 0 to 1, with perfect classification being 1, and where you aren't getting better than chance classified being 0.5. In the context of medical diagnostics, an AUC is greater than 0.90 (excellent), 0.80 to 0.90 (good), or 0.70 to 0.80 (fair).

AUC is a very useful metric since it is independent of the decision threshold, rendering accuracy alone to be significantly less useable. It gives us a single number to evaluate the models and establish their ranking in different tasks or datasets.

CNN-based models have also reported high AUC values in pneumonia detection. An example is that the CheXNet model reported an AUC of 0.93 for pneumonia classification, surpassing the performance of expert radiologists on the same test set. Showing this level of performance shows that deep learning models can help or even replace human expertise.

Additionally, AUC makes model refinement easier. Overall, a rising AUC during training epochs usually means that the model is learning well, and a stagnant AUC or even a dropping AUC can be a sign of overfitting or if the data is not being preprocessed or augmented well enough.

6. Comparative Analysis of CNN Models

6.1. ResNet vs. DenseNet

Two types of deep learning architecture that are quite popular in medical image analysis, especially in pneumonia detection from chest X-rays, are ResNet (Residual Network) and DenseNet (Densely Connected Convolutional Network). Each has unique structural advantages and how they perform, learn, and are applied to particular tasks.

Those introduced by He et al. are ResNet. It is introduced by ResNet, which is the residual learning that works by using shortcut connections to skip over one or more layers. First, these residual connections help to deal with the vanishing gradient problem, allowing networks to be trained with hundreds or even thousands of layers. Of its certainly low efficient structure in learning deep representations and degradation while the depth of the network increases, ResNet is very efficient. ResNet variants like ResNet-50 and ResNet 101 significantly perform in the pneumonia detection task, with good accuracy and relatively short training time.

On the other hand, DenseNet, as introduced by Huang et al., connects a feed forward to all the layers. Unlike ResNet, which sums the previous layer's output to the next, DenseNet concatenates to allow the model to use features more efficiently. It achieves this through dense connectivity, leading to a better flow of gradients and promoting feature propagation and reuse, often leading to better generalization with fewer parameters than equally deep ResNet models.

In pneumonia detection, studies have found DenseNet to often outperform in accuracy and AUC compared to ResNet due to its capability of richer feature extraction. While ResNet is still popular due to faster computation and easier optimization, bypassing initialization of weights is more responsible for deep CNNs having a larger context size. For example, CheXNet, which is built with DenseNet-121, has shown better performance than the ResNet-based model in detecting pneumonia.

In the end, which of ResNet or DenseNet should be used depends on the actual constraints of the application (how many computational resources they have access to, how fast they need the inference to be, what their dataset looks like, etc.)

6.2. MobileNet and Lightweight Models

However, deep learning models are often infeasible in real-world healthcare settings, such as remote or resource-constrained areas, where computational power, memory, and storage restrictions exist. Lightweight CNN architectures like MobileNet come into play here. In this post, I go through MobileNet, Google's model intended to be very efficient on mobile and embedded devices with depthwise separable convolutions to reduce the number of parameters and computations in the model.

MobileNet has real advantages regarding medical diagnostics on the edge: real-time analysis of chest X-rays on our portable devices without powerful GPUs. This provides possibilities to enable the implementation of detection tools for pneumonia in the context of field clinics, rural

hospitals, and even mobile screening units, which may not necessarily have connectivity to cloud-based systems.

Despite the simplicity of MobileNet, it is still competitive in terms of accuracy. MobileNet can achieve an accuracy comparable to heavier models, such as VGG or Resnet if fine-tuned and augmented appropriately (as shown by several studies). As the trade-off between latency and accuracy continued to improve, for example, the Variants MobileNetV2 and MobileNetV3 were more suited for clinical applications where rapid decision-making is important.

Some other lightweight models like SqueezeNet and ShuffleNet have also been attempted for pneumonia detection. These models employ clever architectural innovations that push down model size while maintaining performance. While not as powerful as more complex models for prediction, their ability to be deployed in practical scenarios makes them very valuable.

Overall, lightweight CNNs offer an effective solution that all but democratizes the ability to access AI-powered diagnostics in a manner that does not depend on heavy infrastructure demands in heterogeneous healthcare environments.

Table 3: Comparative Analysis of Lightweight Models

Model	Architecture Type	Model Size (MB)	Parameters (Millions)	Inference Speed	Accuracy (%)	Strengths	Use Case
MobileNetV1	Depthwise Separable CNN	~16	~4.2	Very fast on mobile CPUs	86.5	Low latency, optimized for mobile devices	Mobile apps, edge devices
MobileNetV2	Inverted Residuals + Linear Bottlenecks	~14	~3.4	Faster than V1	88.6	Improved accuracy with fewer operations	Real-time diagnostic tools
MobileNetV3	NAS-optimized Hybrid CNN	~10-16	~5.4	High FPS (frames/sec)	89.4	Balances speed and accuracy with squeeze-and-excite blocks	Wearables, embedded healthcare solutions
SqueezeNet	Fire Modules (1x1 + 3x3 filters)	~5	~1.2	Ultra-fast	83.1	Extremely compact, suitable for minimal storage footprint	IoT healthcare sensors, portable X-ray readers
ShuffleNet	Pointwise group convolution + channel	~5	~1.4	Very fast	84.8	Efficient feature learning with	Battery-powered clinical devices

	shuffle					reduced complexity	
EfficientNet-Lite0	Compound scaled architecture	~20	~5.3	Optimized for inference	90.1	High accuracy-to-size ratio	Cross-platform mobile and web apps

6.3. Ensemble Methods

Ensemble methods use the predictions made by multiple models to achieve better performance, lower variance, and higher robustness models in general. Ensemble learning is particularly helpful in the case of pneumonia detection as it helps offset the weaknesses of different CNN architectures and tends to generalize better on unseen data.

Ensembling CNN models can be done in several ways. For example, one commonplace method is model averaging, where the output from several trained CNNs (e.g., ResNet, DenseNet, and Inception) is averaged as the final output. This also helps to smooth out any of the individual model’s biases or errors. Another way is to do majority voting whereby each model decides on a class, and then the majority among the decisions is taken as the final prediction. Additionally, weighted voting can be used — where the more accurate models have a higher influence.

Some of the most advanced ensemble methods are stacking, where the outputs of base learners are fed into a meta-learner that makes the final prediction. Given the complex patterns and interdependencies, this approach can capture individual models that might fail to pick up.

In the medical image classification domain, it is well known that ensemble methods often perform better than a single model. For example, in this example with the pneumonia dataset, an ensemble of ResNet and DenseNet has been shown to yield better AUC and F1 scores than both models working alone. In addition, ensembles give us more stable predictions across various patient data and imaging conditions.

However, it comes at a trade-off in terms of increased computational cost and complexity. Multiple models need to be trained and maintained, requiring more resources than can be afforded for low-resource environments. Thus, centralized diagnostic systems need ensemble methods for better accuracy and reliability.

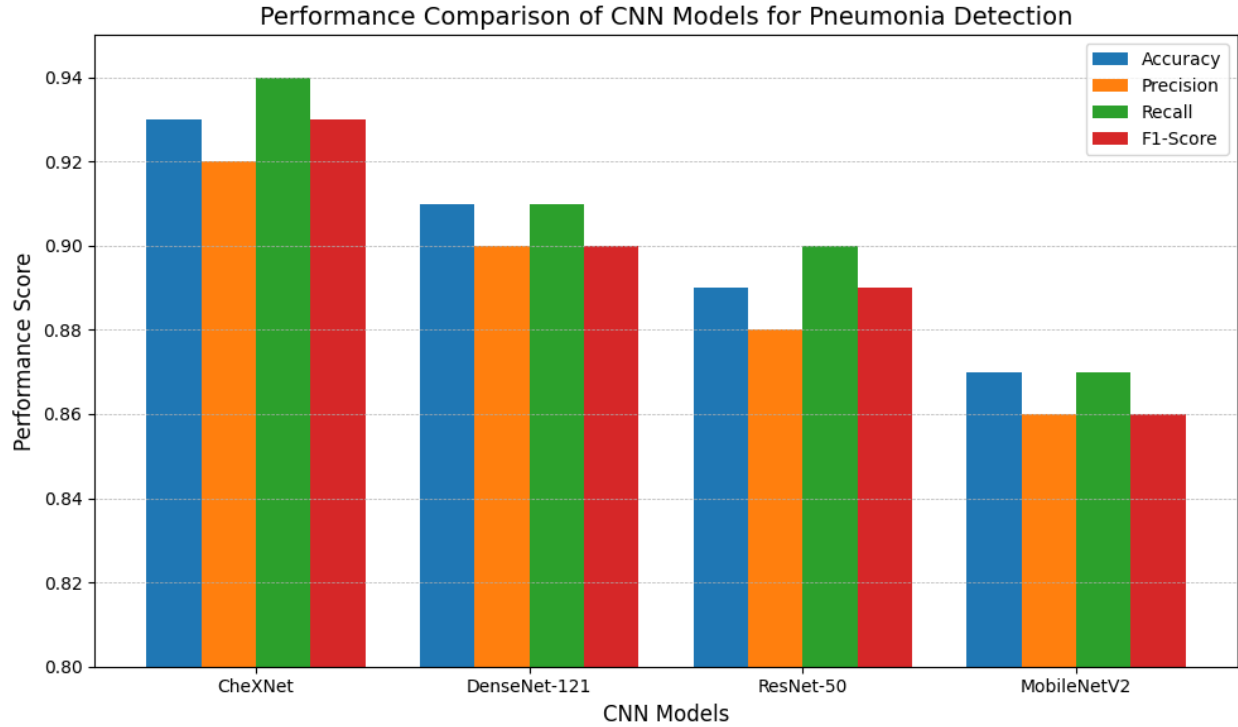


Figure 4. Performance Comparison of CNN Models

7. The Integration of CNN with Other Techniques

7.1. Localization with CNN and YOLO combined

CNNs are very good at classifying, e.g., detecting the presence of pneumonia (or not), but less good at precisely localizing the region of pathology on the chest X-ray. Researchers have started incorporating CNNs with object detection techniques such as YOLO (You Only Look Once), which can detect and localize more than one object within an image in real time to supplement this limitation.

YOLO is a state-of-the-art single-shot detection algorithm that divides the image into grids and detects bounding boxes and class probability for each grid cell at once. Combining YOLO and a CNN trained for pneumonia detection can assist in identifying the exact regions of infected lungs where signs of infection, such as infiltrates or consolidation, are present.

This hybrid approach promises promising results in engendering better interpretability and clinical value of AI models. The system can flag an image as pneumonia-positive and show specific regions of the image where the areas are affected, which helps radiologists verify the diagnosis with confidence. Localization is also used to assess the infection's severity and monitor changes with time, which is necessary for treatment planning.

In studies, CNN-YOLO architectures have reached high accuracy in localizing and detecting pneumonia lesions. In particular, YOLOv4, YOLOv5, etc., integrated with feature extraction backbones like ResNet or EfficientNet, show promising performance in the task of pneumonia detection in the RSNA Pneumonia Detection Challenge.

Using these combined sensing systems allows for more advanced diagnostic profiles that detect disease and can provide information useful in localizing disease through visual means, closing the loop in delivering machine learning outputs to clinical interpretation.

7.2. Hybrid Models and Multimodal Approaches

Hybrid models and multimodal approaches of various Neural Networks in conjunction with different information sources, enhancing performance in the cutting-edge direction of pneumonia detection. Typically, a hybrid model consists of CNNs combined with another deep learning framework such as RNN, LSTM networks or some Transformer based architecture.

For instance, while CNNs can easily learn spatial features of images across different layers of feature hierarchy, LSTMs are suitable to learn temporal dynamics of sequential data. Hybrid CNN LSTM models are ripe because they can handle time series data like sequence chest X-rays taken in a hospitalization setting of a patient for progression of disease tracking and prediction of outcomes.

An approach of such type, which integrates image data with non image data like patient's demographic, clinical notes, laboratory results and Electronic Health Records (EHR) into chest X-ray images, is being proposed. This allows to see a better picture of a patient's overall health status at a more accurate personal level for the diagnosis. For instance, a patient's age, white blood cell count, or oxygen saturation levels can affect the probability of getting pneumonia or the degree of pneumonia and thus help to improve the predictions of the model.

Attention mechanisms have also been proposed to learn attention over both the image features and clinical contextual information in the context of advanced hybrid models. Research published in artificial intelligence in medicine and IEEE Access shows that these models perform better.

Learning architectures are interconnected, each with the other, using multiple data modalities to produce more robust, interpretable, and clinically relevant AI systems. These models are especially useful in real-world problems where decisions are made based on other connected factors besides imaging.

Table 4. Integration Techniques of CNN with Other Methods

Technique	Integrated Method	Purpose	Advantages	Use Case/Application
CNN + YOLO	You Only Look Once (object detection)	Localization of pneumonia regions in X-rays	Real-time detection with bounding boxes	Region-specific diagnosis, visual explanations
CNN + LSTM	Long Short-Term Memory networks	Sequence modeling for progressive analysis over time	Captures temporal patterns in X-ray series	Monitoring disease progression in ICU patients
CNN + Grad-CAM	Gradient-weighted Class Activation Mapping	Interpretability and visual justification	Heatmaps show areas of model focus	Explainable AI for clinical trust
CNN + Transfer Learning	Pretrained models (e.g., ImageNet)	Faster training on small medical datasets	Boosts accuracy with less data	Domain adaptation and cross-institutional use

CNN Ensemble Methods	+	Averaging, majority voting, stacking	Improve robustness and accuracy	Reduces overfitting, enhances generalization	Clinical decision support systems
CNN + Feature Fusion		Combining CNN features with handcrafted or radiomics features	Richer feature representation	Leverages both learned and expert-driven insights	Comprehensive screening solutions
CNN + EHR Data		Electronic Health Record integration	Multimodal analysis combining image and clinical context	Enhances diagnosis with patient history	Personalized diagnostics and risk stratification
CNN Transformer Models	+	Vision Transformers	Global context modeling in images	Improved attention over spatial features	Advanced segmentation and classification

8. Challenges and Limitations

8.1. Data Quality and Variability

However, the quality and variability of the data are some of the most pressing challenges in developing CNN-based systems for pneumonia detection. The presenting datasets of Chest X-rays often suffer from problems, such as inconsistent labeling, different image resolution and artifacts, and different imaging techniques used in other institutions. This can severely degrade the CNN model’s performance and generalizability.

Labels are thus retrieved from automated natural language processing (NLP) systems, which are parsing from radiology reports in many public datasets, which in turn means that labels may be inaccurate. For instance, a label might be not unambiguous and/or not correct, such as in a case when a patient might be clinically diagnosed with pneumonia. At the same time, chest X-ray does not demonstrate clear radiographic evidence. However, this noise in the training data can hinder the learning process and thus end up with models that are not robust.

In addition, imaging protocols can be different; for example, patient positioning, exposure levels, and hardware are various settings, which could influence how you see pneumonia in the X-rays. These domain shifts mean that a dataset from one hospital may work poorly in training a model that will be used on images from another. The lack of cross-institutional generalizability limits the scaling of AI systems for medical imaging to wide-scale usage and their trustworthiness.

The second factor is the dataset is imbalanced. This may make pneumonia cases appear underrepresented compared to other images, leading a model to be biased toward the majority class as it has seen such an image more often. Techniques such as data augmentation and resampling help but can never entirely replace data distribution problems.

As such, curating good quality, diverse, and well-texted datasets remains an important task. To secure large and representative datasets that can improve the robustness and reliability of CNN models in the real world, radiologists, data scientists, and institutions must collaborate.

8.2. Overfitting and Generalization

Machine learning oversells frequently when a model performs well in training data; however, it cannot generalize in unseen data. This is even more critical when using small or nondiverse datasets in CNN-based pneumonia detection. However, excessively complex models, overfitting, and ‘memorizing’ training images rather than learning generalizable features severely limit such models' application in clinical applications.

Medical image analysis suffers from many reasons for overfitting. Second, one of the issues with medical datasets is that they tend to have a limited sample, particularly in rare conditions and pediatric cases. Since CNNs are big domains, they can quickly learn to remember those little examples without proper regularization. Second, the models are trained on data with images from a single source, thereby limiting these models to only performing well in one environment, with one imaging device or a particular demographic of patients.

To prevent overfitting, researchers use several techniques. Dropout, weight decay, and early stopping are regularization techniques that prevent the model from overfitting and depending too much on certain features. In addition, data augmentation, where training images are transformed using rotation, scaling, and translation, also helps teach the network to learn invariant features that enhance generalizability. Another way to avoid overly optimistic numbers with the performance metrics is to use cross-validation.

However, realizing generalization across datasets is a major hurdle. In some cases, models that achieve excellent results on public benchmarks fail when deployed in hospitals with other patient populations or imaging protocols. Because these models exhibit such a huge gap between experimental and real-world performance, there has been growing interest in domain adaptation and transfer learning to bring the performance gap closer.

There is a promise that CNNs might be good at detecting pneumonia. However, we still need to validate these models at scale, in the real world, and do a better job showing that they generalize in different clinical settings.

8.3. Interpretability of CNN Models

Understanding and trusting the decision-making process of a model, or interpretability, is a big issue in deploying CNNs for pneumonia detection in clinics because their internal workings and reasoning are not transparent enough to human users, especially clinicians who base their diagnosis on justifiable evidence.

However, these models lack interpretability, a barrier against clinical adoption. Moreover, physicians must understand ‘why’ a model has predicted pneumonia in a given image. However, simply providing a binary classification without any form of explanation or visual cue undermines confidence in the system, especially if it differs from the clinical judgment.

Researchers address this by developing saliency maps, Grad-CAM, and Layer-wise Relevance Propagation (LRP). These tools produce heatmaps, showing where the model was most affected within the input image. Visual explanations of such can aid radiologists in confirming the model’s output and perhaps find subtle abnormalities missed earlier.

Nevertheless, there are some limitations to these interpretability methods. It often comes with additional computational steps, resulting in inconsistent and misleading visualizations. They also cannot fully reason over the logic of feature extraction and combination in deep networks, which occur across deep (i.e., hundreds of) layers.

Recently, several efforts have been made to make CNNs more transparent. To this end, we propose approaches to developing inherently interpretable models, integrating clinical domain knowledge into model design, and fusing deep learning with symbolic reasoning in building

interpretability into hybrid systems. In addition, regulatory bodies are now starting to require interpretability to be used as a criterion for approving AI-based medical devices. Therefore, it needs to be more interpretable so that the people working in the field can have trust, and it fulfills the legal and ethical standards of the deployment of AI. More progress in this Area is vital to AI's responsible use in detecting pneumonia and many other diseases.

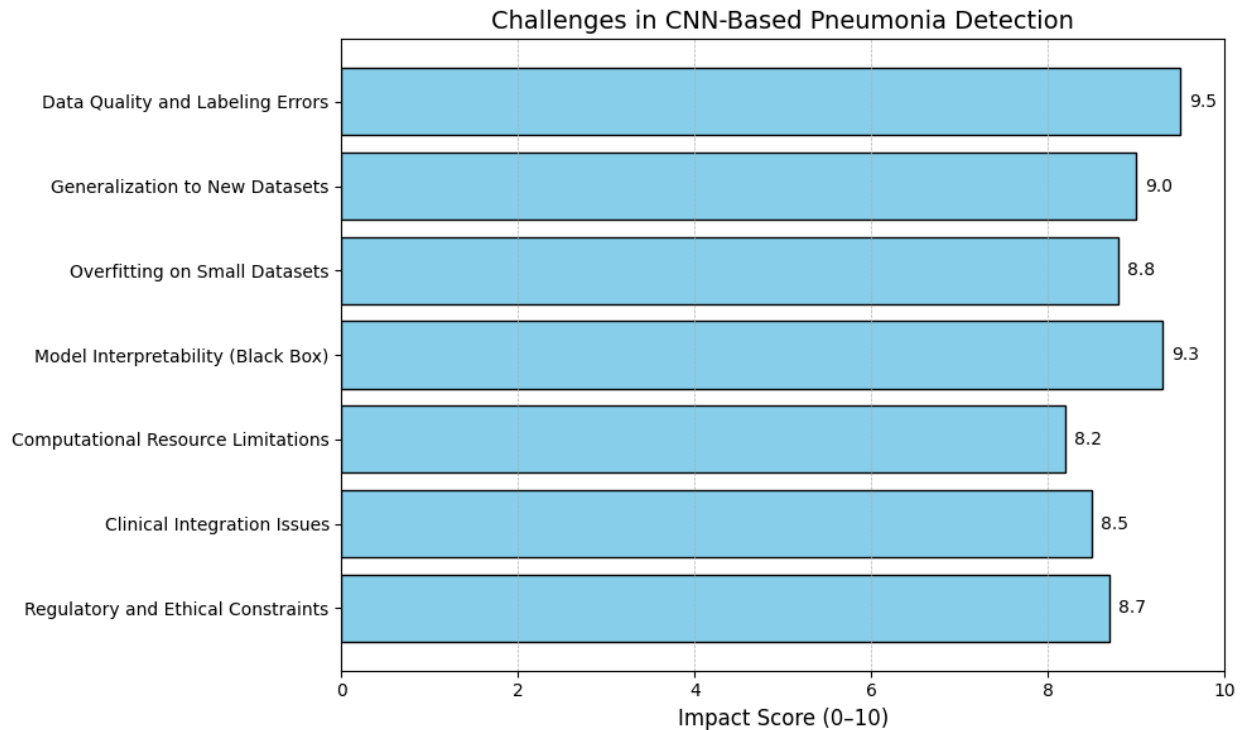


Figure 5. Challenges in CNN-Based Pneumonia Detection

9. Future Directions

9.1. Explainable AI in Medical Imaging

With explainable AI (XAI), medical imaging is set to take a transformational path and solve one of the most lingering problems in deep learning. Interpretability. CNN-based models are outperforming traditional diagnostic techniques. However, the demand for a more transparent and understandable AI system has grown, especially in high-stakes areas such as healthcare.

The objective of XAI, in this case, is to close the gap between high accuracy and clinical usability for pneumonia detection. Techniques such as Grad CAM, SHAP (Shapley Additive explanations), and LIME (Local Interpretable Model agnostic Explanations) are used to gain insight into how a CNN model arrives at its predictions. These tools produce intuitive visualizations and statistical explanations, enabling radiologists and clinicians to validate the AI recommendation more easily.

Besides visualization, research is moving towards inherently interpretable model architectures that combine domain knowledge with the learning process. For instance, models can be built to mimic the radiologists' diagnostic workflows and explicitly localize anatomical landmarks and pathological features before coming to a final decision. Not only does this type of structured

interpretability help explain the model, but it also helps improve their clinical workflow integration.

One promising avenue is if we can integrate clinical context into interpretability. Models can explain decisions beyond just image data by incorporating data from an electronic health record (EHR), patient history, and lab results. It provides this multimodal interpretability, improving both trust and diagnostic accuracy.

Yet another reason explainable AI will continue to expand in importance is that regulatory agencies, like the FDA, increasingly emphasize transparency in their AI tools. In future systems, high performance and the rationale of their outputs have to be demonstrated, making AI a trusted tool instead of an opaque black box in healthcare diagnostics.

9.2. Real-Time Detection Systems

The path towards realizing real-time systems to diagnose pneumonia in situ and immediately is the next frontier for detection using CNNs. Emergency departments, ICUs, and rural clinics are just some examples where the difference between life and death can come down to decisions that need to be made in real time.

Several challenges are needed to achieve real-time performance. To be computationally efficient, the first thing models need to be. Real optimizations are being done to deploy devices with limited processing power, especially for lightweight CNN architectures like MobileNet, EfficientNet-Lite, and the quantized models. With advances in edge computing and embedded AI hardware (e.g., NVIDIA Jetson, Google Coral), executing these models on the same machine as medical imaging devices or portable devices becomes possible.

Second, real-time systems must be accurate across various imaging environments in real-time. They require robust preprocessing pipelines and adaptive models that can cope with image quality, patient positioning, and scanner-type variability.

Furthermore, these systems can be integrated into existing clinical software such as PACS (Picture Archiving and Communication Systems) and EHRs (Electronic Health Records). AI can immediately augment an X-ray uploaded by a radiologist, and suggestions can be fed back to the radiologist to accelerate the triage and priority of critical cases dramatically.

Pilot studies in hospitals show that AI tools with real-time capabilities save diagnosis time and lead to better consistency and reduced diagnostic errors. These systems are also being adopted for telemedicine so that their diagnostic capability can travel to places where they are needed.

Moving forward, combining AI with real-time imaging technologies will be instrumental in future improvements in healthcare, making high-quality care available anytime and everywhere.

9.3. Integration into Clinical Workflows

Any CNN-based pneumonia detection system needs to be seamlessly integrated with clinical workflow. Model accuracy and speed are important, but an AI system's real value is supporting clinicians without taking away time from their day-to-day work.

One way is to embed AI tools into the radiology information systems (RISs) and picture archiving communication systems (PACSs) so that radiologists can view the AI predictions and heatmaps for traditional radiographs. The co-pilot model adds a degree of diagnostic confidence, especially for junior radiologists or when there is a high volume to handle the task that can decrease accuracy due to fatigue.

CDSS with CNN outputs can also suggest follow-up actions based on the AI findings, like ordering an additional test or initiating treatment protocols. To ensure the adoption of an AI-enhanced system, it must be user-friendly with minimal manual input (intuitive interfaces) and all the interfaces and available data presented to the end user.

Regulatory compliance and validation are also key factors. Rigorous clinical trials and approval from the FDA, EMA, or local health authorities (whoever governs the medical discipline AI is used for) are required for artificial intelligence systems. In addition, continuous monitoring and feedback loops are needed to ensure the effectiveness of models in the face of changes in clinical practice and the population of patients.

Their integration needs to be trained and collaborate. Clinicians require education on how the AI system works, what its capabilities are not, and how to interpret outputs. On the other hand, medical professionals require insight from developers to ensure the tools meet real-world needs. Ultimately, it is not only a technical but also a human problem. It's the vision where AI becomes a symbiotic relationship where it makes human expertise run faster, more accurately, and more compassionately run patient care.

10. Conclusion

Although pneumonia is still a major public health problem worldwide, especially in the vulnerable population (children, elderly, patients on immunosuppression) can be a good quality indicator of local sanitary status. Those improvements depend on timely and accurate diagnosis. Specifically in this context, the recent developments in deep learning, especially in Convolutional Neural Networks (CNNs), make them very promising tools to automatically detect signs of pneumonia in chest Xrays, in this case.

In this article, we discussed the structural basis of CNNs, the usages of CNNs in medical imaging, and several architectures that proved effective with pneumonia detection. DenseNet and ResNet have shown the best performance among the models, and MobileNet achieves their performance whilst minimizing the requirement of resources. Techniques like utilizing YOLO to integrate CNNs, as well as using ensemble models, help improve diagnostic accuracy and robustness.

However, there are still some issues. Overfitting is still an issue, particularly in the context of limited or not diverse datasets, which can lead to bad quality and variable data that heavily influence model performance. Finally, interpretability remains a major stumbling block, with clinicians insisting on transparency and explainability of AI's diagnosis. These problems, however, can only be solved by bringing together data scientists, healthcare professionals, and policymakers.

From here, the future of pneumonia detection with CNNs involves further progress on explainable AI, real-time detection systems, and graceful integration into clinical workflows. These advances promise both increased diagnostic accuracy and ease of access to higher levels of healthcare in all settings and populations.

It is paramount to guarantee ethical use of the technology, vet it robustly, and involve clinicians at every stage of technology development. Through this, we can capitalize on the untapped potential of AI to revolutionize the checking for pneumonia and, in fact, medical imaging in its entirety—to change the future of healthcare worldwide.

References

- [1] Rajpurkar, P., Irvin, J., Zhu, K., et al. (2017). CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning. arXiv preprint arXiv:1711.05225.
- [2] Jiang, Z. (2020). Chest X-ray Pneumonia Detection Based on Convolutional Neural Networks. 2020 International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE).
- [3] Patil, N., Doshi, A., & Dodiya, V. (2024). Pneumonia detection using deep neural networks based on chest x-ray images. In Algorithms for intelligent systems (pp. 395–406). https://doi.org/10.1007/978-981-97-0327-2_29
- [4] Kulkarni, R. V., Paldewar, A., Paliwal, S., Ajinath, P. A., & Panchal, S. (2024). Addressing class imbalance for improved pneumonia diagnosis: Comparative analysis of GANs and weighted loss function in classification of the chest X-ray images. 2024 5th International Conference for Emerging Technology (INCET), 1–7. <https://doi.org/10.1109/INCET61516.2024.10593502>
- [5] Gabruseva, T., Poplavskiy, D., & Kalinin, A. A. (2020). Deep Learning for Automatic Pneumonia Detection. arXiv preprint arXiv:2005.13899.
- [6] Baltruschat, I. M., Nickisch, H., Grass, M., Knopp, T., & Saalbach, A. (2019). Comparison of Deep Learning Approaches for Multi-Label Chest X-Ray Classification. Scientific Reports, 9, 6381.
- [7] Radočaj, P., Radočaj, D., & Martinović, G. (2025). Optimizing Convolutional Neural Network Architectures with Optimal Activation Functions for Pediatric Pneumonia Diagnosis Using Chest X-Rays. Big Data and Cognitive Computing, 9(2), 25.
- [8] Bairwa, H., & Jangid, R. (2024). Pneumonia Detection from Chest X-rays Using the CheXNet Deep Learning Algorithm. Preprints.org.
- [9] Ma, Y., & Gonzales, R. A. (2024). Convolutional Neural Network-Based Analysis of Pediatric Chest X-ray Images for Pneumonia Detection. Journal of Emerging Investigators, 7.
- [10] Radočaj, P., Radočaj, D., & Martinović, G. (2025). Optimizing Convolutional Neural Network Architectures with Optimal Activation Functions for Pediatric Pneumonia Diagnosis Using Chest X-Rays. Big Data and Cognitive Computing, 9(2), 25.
- [11] Natashia, C. (2020, June 6). Chest X-rays Pneumonia Detection using Convolutional Neural Network. Medium; TDS Archive. <https://medium.com/data-science/chest-x-rays-pneumonia-detection-using-convolutional-neural-network-63d6ec2d1dee>
- [12] Samat, N. A., Salleh, M., & Ali, H. (2020). The comparison of pooling functions in convolutional neural network for sentiment analysis task. In Proceedings of the 7th International Conference on Soft Computing & Pattern Recognition (SoCPaR 2019) (pp. [insert page numbers]). Springer. https://doi.org/10.1007/978-3-030-36056-6_20
- [13] Anbarasan, P., Kavyasree, N., & Sanjana, N. (2024). Comparative Analysis of Convolutional Neural Network Architectures for Pneumonia Diagnosis from Chest X-ray Images. 2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT).
- [14] Esfahani, Shirin Nasr, and Shahram Latifi. "A Survey of State-of-The-Art GAN-Based Approaches to Image Synthesis." 9th International Conference on Computer Science, Engineering and Applications (CCSEA 2019), 13 July 2019, csitcp.com/paper/9/99csit06.pdf, <https://doi.org/10.5121/csit.2019.90906>.
- [15] Nabati, R., & Qi, H. (2019). "RRPN: Radar Region Proposal Network for Object Detection in Autonomous Vehicles." 2019 IEEE International Conference on Image Processing (ICIP), Taipei, Taiwan, 2019, pp. 3093-3097, doi: 10.1109/ICIP.2019.8803392.
- [16] Rawat, W., & Wang, Z. (2017). "Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review." Neural Computation, 29(9), pp. 2352-2449, Sept. 2017, doi: 10.1162/neco_a_00990.
- [17] Wang, Weibin, et al. "Medical Image Classification Using Deep Learning." Intelligent Systems Reference Library, 19 Nov. 2019, pp. 33–51, https://doi.org/10.1007/978-3-030-32606-7_3.
- [18] Alom, Md Zahangir, et al. "The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches." ArXiv:1803.01164 [Cs], 12 Sept. 2018, arxiv.org/abs/1803.01164.
- [19] Frid-Adar, Maayan, et al. "GAN-Based Synthetic Medical Image Augmentation for Increased CNN Performance in Liver Lesion Classification." Neurocomputing, vol. 321, Dec. 2018, pp. 321–331, <https://doi.org/10.1016/j.neucom.2018.09.013>.
- [20] Karp, Rafal, and Zaneta Swiderska-Chadaj. Automatic Generation of Graphical Game Assets Using GAN. 13 July 2021, <https://doi.org/10.1145/3477911.3477913>.
- [21] L. Jiao and J. Zhao, "A Survey on the New Generation of Deep Learning in Image Processing," in IEEE Access, vol. 7, pp. 172231-172263, 2019, doi: 10.1109/ACCESS.2019.2956508.

- [22] L. Wang, W. Chen, W. Yang, F. Bi and F. R. Yu, "A State-of-the-Art Review on Image Synthesis With Generative Adversarial Networks," in *IEEE Access*, vol. 8, pp. 63514–63537, 2020, doi: 10.1109/ACCESS.2020.2982224.
- [23] Shorten, Connor, and Taghi M. Khoshgoftaar. "A Survey on Image Data Augmentation for Deep Learning." *Journal of Big Data*, vol. 6, no. 1, 6 July 2019, journalofbigdata.springeropen.com/articles/10.1186/s40537-019-0197-0, <https://doi.org/10.1186/s40537-019-0197-0>.
- [24] Kayalibay, Baris, et al. "CNN-Based Segmentation of Medical Imaging Data." *ArXiv:1701.03056 [Cs]*, 25 July 2017, arxiv.org/abs/1701.03056.
- [25] Jain, M., & None Arjun Srihari. (2023). House price prediction with Convolutional Neural Network (CNN). *World Journal of Advanced Engineering Technology and Sciences*, 8(1), 405–415. <https://doi.org/10.30574/wjaets.2023.8.1.0048>
- [26] Jain, M., & Arjun Srihari. (2024b). Comparison of Machine Learning Algorithm in Intrusion Detection Systems: A Review Using Binary Logistic Regression. *International Journal of Computer Science and Mobile Computing*, 13(10), 45–53. <https://doi.org/10.47760/ijcsmc.2024.v13i10.005>
- [27] Jain, M., & Shah, A. (2022). Machine Learning with Convolutional Neural Networks (CNNs) in Seismology for Earthquake Prediction. *Iconic Research and Engineering Journals*, 5(8), 389–398. <https://www.irejournals.com/paper-details/1707057>
- [28] Jain, M., & Arjun Srihari. (2024). Comparison of CAD Detection of Mammogram with SVM and CNN. *Iconic Research and Engineering Journals*, 8(6), 63–75. <https://www.irejournals.com/paper-details/1706647>
- [29] Kaushik, P., & Jain, M. A Low Power SRAM Cell for High Speed Applications Using 90nm Technology. *Csjournals. Com*, 10. <https://www.csjournals.com/IJEE/PDF10-2/66.%20Puneet.pdf>
- [30] Jain, M., & Arjun Srihari. (2024b). Comparison of Machine Learning Models for Stress Detection from Sensor Data Using Long Short-Term Memory (LSTM) Networks and Convolutional Neural Networks (CNNs). *International Journal of Scientific Research and Management (IJSRM)*, 12(12), 1775–1792. <https://doi.org/10.18535/ijrm/v12i12.ec02>
- [31] Kaushik, P., & Jain, M. (2018). Design of low power CMOS low pass filter for biomedical application. *International Journal of Electrical Engineering & Technology (IJEET)*, 9(5).
- [32] Jain, M., & Shah, A. (2024). Anomaly Detection Using Convolutional Neural Networks (CNN). *ESP International Journal of Advancements in Computational Technology (ESP-IJACT)*, 2(3), 12–22. <https://www.espjournals.org/IJACT/ijact-v2i3p102>
- [33] Kumar, Y., Saini, S., & Payal, R. (2020). Comparative Analysis for Fraud Detection Using Logistic Regression, Random Forest and Support Vector Machine. *SSRN Electronic Journal*.
- [34] Höppner, S., Baesens, B., Verbeke, W., & Verdonck, T. (2020). Instance-Dependent Cost-Sensitive Learning for Detecting Transfer Fraud. *arXiv preprint arXiv:2005.02488*.
- [35] Niu, X., Wang, L., & Yang, X. (2019). A Comparison Study of Credit Card Fraud Detection: Supervised versus Unsupervised. *arXiv preprint arXiv:1904.10604*.
- [36] Bhat, N. (2019). Fraud detection: Feature selection-over sampling. *Kaggle*. Retrieved from <https://www.kaggle.com/code/nareshbhat/fraud-detection-feature-selection-over-sampling>
- [37] InsiderFinance Wire. (2021). Logistic regression: A simple powerhouse in fraud detection. *Medium*. Retrieved from <https://wire.insiderfinance.io/logistic-regression-a-simple-powerhouse-in-fraud-detection-15ab984b2102>
- [38] Olaitan, V. O. (2020). Feature-based selection technique for credit card fraud detection. Master's Thesis, National College of Ireland. Retrieved from <https://norma.ncirl.ie/5122/1/olaitanvictoriaolanlokun.pdf>
- [39] Raymaekers, J., Verbeke, W., & Verdonck, T. (2021). Weight-of-evidence 2.0 with shrinkage and spline-binning. *arXiv preprint arXiv:2101.01494*. Retrieved from <https://arxiv.org/abs/2101.01494>
- [40] Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2017). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784–3797. <https://doi.org/10.1109/TNNLS.2017.2736643>
- [41] Carcillo, F., Dal Pozzolo, A., Le Borgne, Y. A., Caelen, O., Mazzer, Y., & Bontempi, G. (2019). Scarff: A scalable framework for streaming credit card fraud detection with spark. *Information Fusion*, 41, 182–194. <https://doi.org/10.1016/j.inffus.2017.09.005>
- [42] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66. <https://doi.org/10.1016/j.cose.2015.09.005>
- [43] Mohit Jain, Arjun Srihari (2024). Comparison of Machine Learning Models for Stress Detection from Sensor Data Using Long Short-Term Memory (LSTM) Networks and Convolutional Neural Networks (CNNs). <https://ijrm.net/index.php/ijrm/article/view/5912/3680> <https://doi.org/10.18535/ijrm/v12i12.ec02>
- [44] Zareapoor, M., & Shamsolmoali, P. (2015). Application of credit card fraud detection: Based on bagging ensemble classifier. *Procedia Computer Science*, 48, 679–685. <https://doi.org/10.1016/j.procs.2015.04.201>

- [45] Mohit Jain, Adit Shah (2024). Anomaly Detection Using Convolutional Neural Networks (CNN). *ESP International Journal of Advancements in Computational Technology*. <https://www.espjournals.org/IJACT/2024/Volume2-Issue3/IJACT-V2I3P102.pdf>
- [46] Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602–613. <https://doi.org/10.1016/j.dss.2010.08.008>
- [47] Mohit Jain and Arjun Srihari (2023). House price prediction with Convolutional Neural Network (CNN). <https://wjaets.com/sites/default/files/WJAETS-2023-0048.pdf>
- [48] Patel, H., & Zaveri, M. (2011). Credit card fraud detection using neural network. *International Journal of Innovative Research in Computer and Communication Engineering*, 1(2), 1–6. https://www.ijrccce.com/upload/2011/october/1_Credit.pdf
- [49] Puneet Kaushik, Mohit Jain, Gayatri Patidar, Paradayil Rhea Eapen, Chandra Prabha Sharma (2018). Smart Floor Cleaning Robot Using Android. *International Journal of Electronics Engineering*. <https://www.csjournals.com/IJEE/PDF10-2/64.%20Puneet.pdf>
- [50] Duman, E., & Ozcelik, M. H. (2011). Detecting credit card fraud by genetic algorithm and scatter search. *Expert Systems with Applications*, 38(10), 13057–13063. <https://doi.org/10.1016/j.eswa.2011.04.102>
- [51] Puneet Kaushik, Mohit Jain. “A Low Power SRAM Cell for High Speed Applications Using 90nm Technology.” *Csjournals.Com* 10, no. 2 (December 2018): 6. <https://www.csjournals.com/IJEE/PDF10-2/66.%20Puneet.pdf>
- [52] Jain, M., & Srihari, A. (2021). Comparison of CAD detection of mammogram with SVM and CNN. *IRE Journals*, 8(6), 63-75. <https://www.irejournals.com/formatedpaper/1706647.pdf>
- [53] Kaushik, P., Jain, M., & Jain, A. (2018). A pixel-based digital medical images protection using genetic algorithm. *International Journal of Electronics and Communication Engineering*, 31-37. http://www.irphouse.com/ijece18/ijecev11n1_05.pdf
- [54] Kaushik, P., Jain, M., & Shah, A. (2018). A Low Power Low Voltage CMOS Based Operational Transconductance Amplifier for Biomedical Application. <https://ijsetr.com/uploads/136245IJSETR17012-283.pdf>
- [55] Jain, M., & Shah, A. (2022). Machine Learning with Convolutional Neural Networks (CNNs) in Seismology for Earthquake Prediction. *Iconic Research and Engineering Journals*, 5(8), 389–398. <https://www.irejournals.com/paper-details/1707057>
- [56] Kaushik, P., & Jain, M. (2018). Design of low power CMOS low pass filter for biomedical application. *International Journal of Electrical Engineering & Technology (IJEET)*, 9(5).
- [57] Nettur, S. B., Karpurapu, S., Nettur, U., et al. (2025). Lightweight Weighted Average Ensemble Model for Pneumonia Detection in Chest X-Ray Images. *arXiv preprint arXiv:2501.16249*.
- [58] Ahmed Shuvo, S., Islam, M. A., Hoque, M. M., & Sulaiman, R. B. (2023). Comparative Study of Deep Learning Models for Binary Classification on Combined Pulmonary Chest X-ray Dataset. *arXiv preprint arXiv:2309.10829*.
- [59] El Asnaoui, K., Chawki, Y., & Idri, A. (2020). Automated Methods for Detection and Classification Pneumonia based on X-Ray Images Using Deep Learning. *arXiv preprint arXiv:2003.14363*.
- [60] Li, D. (2023). Attention-enhanced architecture for improved pneumonia detection in chest X-ray images. *BMC Medical Imaging*, 23, 1177.
- [61] Kermany, D. S., Zhang, K., & Goldbaum, M. (2018). Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. *Cell*, 172(5), 1122-1131.e9.
- [62] Wang, X., Peng, Y., Lu, L., et al. (2017). ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2097-2106.
- [63] Stephen, O., Sain, M., Maduh, U. J., & Jeong, D. U. (2019). An Efficient Deep Learning Approach to Pneumonia Classification in Healthcare. *Journal of Healthcare Engineering*, 2019, 4180949.
- [64] Chouhan, V., Singh, S. K., Khamparia, A., et al. (2020). A Novel Transfer Learning Based Approach for Pneumonia Detection in Chest X-ray Images. *Applied Sciences*, 10(2), 559.
- [65] Togacar, M., Ergen, B., & Comert, Z. (2020). COVID-19 Detection Using Deep Learning Models to Exploit Social Mimic Optimization and Structured Chest X-ray Images Using Fuzzy Color and Segmentation Features. *Computers in Biology and Medicine*, 121, 103805.
- [66] Lakhani, P., & Sundaram, B. (2017). Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks. *Radiology*, 284(2), 574-582.
- [67] Yao, L., Prosky, J., Poblens, E., Covington, B., & Lyman, K. (2018). Weakly Supervised Medical Diagnosis and Localization from Multiple Resolutions. *arXiv preprint arXiv:1803.07703*.
- [68] Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). Variable Generalization Performance of a Deep Learning Model to Detect Pneumonia in Chest Radiographs: A Cross-Sectional Study. *PLoS Medicine*, 15(11), e1002683.

- [69] Islam, M. T., Aowal, M. A., Minhaz, A. T., & Ashraf, K. (2017). Abnormality Detection and Localization in Chest X-Rays Using Deep Convolutional Neural Networks. arXiv preprint arXiv:1705.09850.
- [70] Rajpurkar, P., Irvin, J., Ball, R. L., et al. (2018). Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of the CheXNeXt Algorithm to Practicing Radiologists. *PLoS Medicine*, 15(11), e1002686.
- [71] Rubin, J., Sanghavi, D., Zhao, C., et al. (2018). Large Scale Automated Reading of Frontal and Lateral Chest X-Rays Using Dual Convolutional Neural Networks. arXiv preprint arXiv:1804.07839.
- [72] Guan, Q., Huang, Y., Zhong, Z., et al. (2018). Diagnose Like a Radiologist: Attention Guided Convolutional Neural Network for Thorax Disease Classification. arXiv preprint arXiv:1801.09927.
- [73] Khan, A. I., Shah, J. L., & Bhat, M. M. (2020). CoroNet: A Deep Neural Network for Detection and Diagnosis of COVID-19 from Chest X-ray Images. *Computer Methods and Programs in Biomedicine*, 196, 105581.
- [74] Apostolopoulos, I. D., & Mpesiana, T. A. (2020). COVID-19: Automatic Detection from X-Ray Images Utilizing Transfer Learning with Convolutional Neural Networks. *Physical and Engineering Sciences in Medicine*, 43(2), 635-640.
- [75] Ozturk, T., Talo, M., Yildirim, E. A., et al. (2020). Automated Detection of COVID-19 Cases Using Deep Neural Networks with X-ray Images. *Computers in Biology and Medicine*, 121, 103792.
- [76] Chowdhury, M. E. H., Rahman, T., Khandakar, A., et al. (2020). Can AI Help in Screening Viral and COVID-19 Pneumonia? *IEEE Access*, 8, 132665-132676.
- [77] He, X., Yang, X., Zhang, S., et al. (2020). Sample-Efficient Deep Learning for COVID-19 Diagnosis Based on CT Scans. medRxiv.
- [78] Oh, Y., Park, S., & Ye, J. C. (2020). Deep Learning COVID-19 Features on CXR Using Limited Training Data Sets. *IEEE Transactions on Medical Imaging*, 39(8), 2688-2700.
- [79] Singh, D., Kumar, V., & Kaur, M. (2020). Classification of COVID-19 Patients from Chest CT Images Using Multi-Objective Differential Evolution-Based Convolutional Neural Networks. *European Journal of Clinical Microbiology & Infectious Diseases*, 39(7), 1379-1389.
- [80] Horry, M. J., Chakraborty, S., Paul, M., et al. (2020). COVID-19 Detection through Transfer Learning Using Multimodal Imaging Data. *IEEE Access*, 8, 149808-149824.
- [81] Asif, S., Wenhui, Y., Jin, H., & Tao, Y. (2020). Classification of COVID-19 from Chest X-ray Images Using Deep Convolutional Neural Networks. *Interdisciplinary Sciences: Computational Life Sciences*, 13(2), 1-10.
- [82] Loey, M., Smarandache, F., & M Khalifa, N. E. (2020). Within the Lack of Chest COVID-19 X-ray Dataset: A Novel Detection Model Based on GAN and Deep Transfer Learning. *Symmetry*, 12(4), 651.
- [83] Narin, A., Kaya, C., & Pamuk, Z. (2020). Automatic Detection of Coronavirus Disease (COVID-19) Using X-ray Images and Deep Convolutional Neural Networks. *Pattern Analysis and Applications*, 24(3), 1207-1220.
- [84] Khan, S. H., Sohail, A., Zahoor, U., & Qureshi, A. S. (2020). A Survey of the Recent Architectures of Deep Convolutional Neural Networks. *Artificial Intelligence Review*, 53(8), 5455-5516.