

Article

Not peer-reviewed version

---

# Vision-Based Detection and Robotic Intervention System for Early Identification of Pneumonia from Chest X-Ray Images

---

[Jenisha Raut](#)\*

Posted Date: 14 April 2026

doi: 10.20944/preprints202604.0878.v1

Keywords: pneumonia detection; chest X-ray; deep learning; CNN; computer vision; robotic intervention; healthcare AI



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Vision-Based Detection and Robotic Intervention System for Early Identification of Pneumonia from Chest X-Ray Images

Jenisha Raut

Department of Nursing and Health Professions 1, North Hennepin Community College, Minneapolis, USA;  
jenisharahut@gmail.com

## Abstract

Pneumonia is one of the major causes of death, especially in paediatric groups, with late diagnosis being a major risk factor that complicates the clinical situation and treatment. Traditional diagnostic methods based on the use of manual decoding of chest X-rays are limited by inter-observer error and the lack of specialists, especially in high-volume or resource-intensive environments. This paper suggests a vision-based detection system based on a convolutional neural network that was trained on 5,863 chest X-ray images to classify pneumonia (binary). The framework is also not limited to diagnosis, it incorporates a simulated robotic intervention module that can invoke automated alerts, clinical notification, and preliminary response activities. The model uses standardised preprocessing and has an efficient feature extraction with about 11.17 million parameters. The system has a test accuracy of 89.74% with a recall of about 96.7%, which means that the system is highly sensitive in detecting pneumonia cases. Combining AI-based detection with robotic action shows that a scalable solution can be used in real-time clinical assistance, especially in intelligent and remote healthcare settings.

**Keywords:** pneumonia detection; chest X-ray; deep learning; CNN; computer vision; robotic intervention; healthcare AI

---

## 1. Introduction

Pneumonia is a worldwide health issue of concern, and it causes about 14-15 percent of all deaths among children below the age of five and about 700,000 deaths every year [1]. Low- and middle-income areas are especially harsh since not every area has adequate access to timely diagnostic facilities. They should be identified early because later diagnosis leads to high morbidity, hospitalisation and mortality. Even with the development of the medical imaging system, the failure to detect the disease efficiently still exists, mainly related to the shortage of resources and the lack of radiologists with high competence. As a result, scalable, accurate, and fast diagnostic solutions are urgently needed that can be used both in clinical and resource-restricted settings.

Chest X-ray imaging is the most used diagnostic modality in the detection of pneumonia because it is cost-effective and readily available. Yet, optical reading of radiographs is inherently subjective and is likely to be affected by inter-observer errors. Research has shown that radiologists might reach a 20% variation in diagnostic agreement, especially about the differentiation of viral and bacterial pneumonia patterns. Moreover, the growing amount of imaging data places cognitive and time-related burdens on clinicians and increases the chances of misdiagnosis. Current models are generally dependent upon detailed imaging or genomics data and may not be feasible to adopt in resource-limited environments [2]. These shortcomings underscore the incompetence of depending on the conventional method of diagnostic procedures, especially in healthcare institutions where demand is high.

The latest advances in artificial intelligence, especially deep learning, have proven to have great potential in the analysis of medical images. Convolutional neural networks (CNNs) can automatically learn hierarchical features of raw images, which allows them to achieve strong classification results without using manual feature engineering. It has been empirically indicated that CNN-based models can deliver diagnostic accuracies of over 90% in pneumonia detection tasks [3]. Nevertheless, these models do enhance the accuracy of diagnostics, but their use is very restricted to decision support with little incorporation into the overall clinical response system.

The most critical research gap is that the current studies do not include comprehensive frameworks that go beyond detection to intervention. Majority of the studies only pay attention to classification accuracy and not the operational aspect of healthcare delivery [4]. Particularly, the possibility of integrating AI-based detection with robots to carry out automated response is under-researched. This is a restricting factor of such models to be applied practically to the real world where such models do not have the time to accurately diagnose but rather act on the situation as swiftly as possible.

This paper fills this gap by suggesting a vision-based pneumonia detecting system that is combined with a simulated robotic intervention system. Combining the results of the CNN model trained on 5,863 chest X-ray images, the system provides high diagnostic accuracy and allows automated clinical response, such as alert generation and preliminary intervention assistance.

## 2. Literature Review

### 2.1. Deep Learning in Medical Imaging

Deep learning has radically changed the field of medical imaging by allowing the extraction of features and high-dimensional pattern recognition of raw data without any human involvement. The convolutional neural networks (CNNs) especially are the most popular architecture because they can apprehend spatial hierarchies by convolutional and pooling operations [5]. CNNs have been highly effective at various diagnostic imaging tasks, such as tumour detection, retinal disease classification and pulmonary pathology detection [6]. The empirical results show that in a controlled environment, CNN-based systems can attain classification accuracy of over 90%, which can drastically decrease the need to rely on manual interpretation [7].

The efficacy of CNNs in the clinical setting, however, depends on their resiliency to noise, differences in image capture, and patient heterogeneity. Although high accuracy is commonly reported, such models can be found to have lower generalisability than unseen data distributions. Also, deep learning is black box, which makes it questionable in terms of interpretability and clinical trust [8]. Nevertheless, CNNs are not completely useless as a pneumonia detector because they are scalable and can easily handle large imaging datasets, so they are especially relevant to high-throughput screening tasks [9].

### 2.2. Pneumonia Detection Studies

Machine Learning is popular in education, multilingual interaction, industries, manufacturing, healthcare and disease detection [10]. Pneumonia detection Recent works in this area have heavily leveraged CNN-based architectures, both trained as custom-designed models and transfer learning models (ResNet and VGGNet). Average CNN models are expected to reach accuracy of 85 to 92 percent with more advanced models, especially ResNet50, having reported performance over 94 percent in optimised cases. Such advancements can be greatly explained by better network structures and residual learning processes, which improve feature representation and reduce the problem of vanishing gradients [11].

Although these steps have been taken, there is also the issue of performance variability. Most researchers achieve high accuracy on validation datasets, but do not achieve consistent accuracy on independent test datasets, which suggests overfitting. Also, sensitivity (recall) is frequently emphasized more than specificity, as it is clinically important to reduce the number of false negatives

[12]. Sensitivity in several cases should be greater than 95% and the specificity could be less than 85% which results in higher rates of false positives and unnecessary clinical intervention.

Importantly, most of the research conducted so far is concerned with the performance of classification in a vacuum and does not consider how detection systems can be integrated into clinical workflows. Such short-sightedness restricts the practical value of such models because correct diagnosis is not a guarantee of better patient outcomes unless there is immediate and coordinated action.

### 2.3. Dataset and Challenges

These models are highly sensitive to the quality and composition of the training data, which is crucial to the performance of deep learning models in identifying pneumonia. Accessible datasets, including chest X-ray collections of about 5,000 to 6,000 images, are typically used because of their accessibility [13]. Nonetheless, these datasets are usually not balanced, and the number of pneumonia cases is a lot higher than the number of normal cases or the other way round. This imbalance may skew the model in favour of the powerful group, making it less efficient at generalisation.

The other major challenge is the low variety of medical datasets. A lot of data comes out of individual institutions and demographics, including paediatric patients between one and five years old [14]. This limits the usefulness of the model in larger populations. Also, there is the issue of annotation reliability whereby even expert-labeled data can be inconsistent because of inter-observer variability. Multi-expert validation decreases the rate of error but fails to remove ambiguity in borderline cases. The restrictions indicate the necessity of effective preprocessing and validation procedures to guarantee credible model performance.

### 2.4. Research Gap

The main weakness of the existing literature is the absence of unification of diagnostic models with real-time clinical response systems. Although CNN-based methods have shown good results in image classification, they are still mostly used to do offline analysis or support systems. This lack of integration between detection and intervention negates the possible role of artificial intelligence in healthcare environments where the quick reaction is critical to enhance patient outcomes [15].

Furthermore, the topic of robotics introduction into the medical diagnosis process is still under-researched, especially in the context of respiratory illnesses management. The current systems do not often have automated systems of alert production, patient monitoring, or initial intervention after the diagnosis. This is a major loophole since, with AI-based detection combined with robotics, it is possible to have round-the-clock surveillance and instant reaction, particularly in resource-limited situations.

This gap is filled by current research, as it proposes an integrated framework that combines CNN-based pneumonia detection with simulated robotic intervention systems. This method is not only about classification but also involves practical outcomes that would be taken, which would make the proposed solution more practical and clinically relevant.

## 3. Materials and Methods

### 3.1. Dataset Description

The paper uses a publicly available chest X-ray dataset that consists of 5,863 X-ray images that are classified into two categories: Pneumonia and Normal. The data is based on the paediatric patients with the age range of one to five years, which makes them clinically relevant to the respiratory diagnosis in the early stage. Images were obtained within the regular anterior post radiographic conditions and were subjected to a multiphase quality control procedure. Poor quality and illegible scans were eliminated to improve the reliability of the data sets.

Two expert clinicians annotated each image, and the assessment subset was validated by a third expert to reduce diagnostic ambiguity. This multi-expert validation system improves the credibility of labels, but fails to completely remove inter-observer inconsistencies, especially in borderline cases where the border between viral and bacterial pneumonia lies between visual features.

The dataset is divided into three subsets: training (5,216 images), validation (16 images), and testing (624 images). Although this split is useful in the development and evaluation of the model, the proportion of the validation set is too small (about 0.27 of the total data), which may not be effective in hyperparameter tuning and may lead to unstable validation results. Irrespective of this shortcoming, the dataset can still be used in binary classification tasks and offer a sound basis when assessing deep learning-based diagnostic systems.

### 3.2. Data Preprocessing

To maintain uniformity and maximise the efficiency of the model, the images were downscaled to a set size size of  $224 \times 224$  pixels, which fits the common input demands of convolutional neural networks. The pixel intensity values were scaled to 0-1 range which stabilises gradient updates during training and speeds up convergence.

To overcome the threats of overfitting and improve generalisability, data augmentation methods were only used on the training set. These were random rotations (up to 15 degrees), horizontal inversion and zooming. These methods artificially increase the sample size and cause variability that mimics the conditions of real-world imaging.

Nevertheless, augmentation strategies are to be skilfully managed in medical imaging scenarios. Over-changes can alter features of clinical interest, which can be misleading to the model [16]. In the present work, moderate augmentation parameters were used to preserve anatomical integrity in addition to enhancing model robustness. This equilibrium is essential in preserving diagnostic validity.

### 3.3. Model Architecture

The suggested system uses a specially designed convolutional neural network (CNN), which is designed to binary classify the image of chest X-rays. The network has three convolutional layers with filter depths of 32, 64 and 128 respectively. A max-pooling operation (reduction of spatial dimensionality and computation) follows every convolutional layer but retains salient features.

The convolutional layers perform the task of extracting hierarchical image features, starting with low-level edges and textures and the more complicated pathological patterns that are related to pneumonia. After extracting features, a flattening layer is used to translate multidimensional feature maps into one-dimensional feature, and this is fed through a fully connected dense layer containing 128 neurons.

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 222, 222, 32)	896
max_pooling2d (MaxPooling2D)	(None, 111, 111, 32)	0
conv2d_1 (Conv2D)	(None, 109, 109, 64)	18,496
max_pooling2d_1 (MaxPooling2D)	(None, 54, 54, 64)	0
conv2d_2 (Conv2D)	(None, 52, 52, 128)	73,856
max_pooling2d_2 (MaxPooling2D)	(None, 26, 26, 128)	0
flatten (Flatten)	(None, 86528)	0
dense (Dense)	(None, 128)	11,075,712
dropout (Dropout)	(None, 128)	0
dense_1 (Dense)	(None, 1)	129

Total params: 11,169,089 (42.61 MB)  
 Trainable params: 11,169,089 (42.61 MB)  
 Non-trainable params: 0 (0.00 B)

Figure 1. Proposed CNN Architecture for Pneumonia Detection.

Figure 1 shows the sequential CNN model architecture employed to classify the images of the chest X-ray images as either normal or pneumonia based on hierarchical feature extraction.

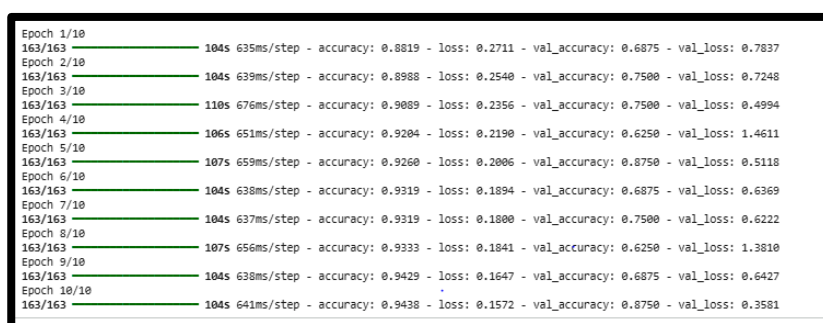
To prevent overfitting, a dropout layer with 0.5 rate is added to deactivate half of neurons randomly each time training is performed. The last output layer employs a sigmoid activation function that gives a probability value of 0 to 1 in binary classification.

The model has a total of about 11.17 million trainable parameters which represents a trade-off between representational and computational efficiency. Although more complex architectures like ResNet have better performance, the chosen model has more interpretability and can be trained in less time, which is why it can be used in a real-time deployment scenario. This design option is in line with the fact that the study is aimed at practical use and not just maximisation of accuracy.

### 3.4. Training Setup

Adam optimiser was used to train the model consisting of adaptive learning rates and momentum-based updates to increase the efficiency of convergence. The learning rate was chosen as 0.001 to achieve a balance between the speed of training and stability. A binary cross-entropy loss function was used because it is an appropriate loss in binary classification and the loss is given a probabilistic interpretation of the error in prediction.

The training was performed using 10 epochs and the batch size was 32, so the total number of steps per epoch was 163 steps depending on the size of the training dataset. The validation was done after every epoch to check the performance of the model and identify possible overfitting.



Epoch 1/10	163/163	104s	635ms/step	accuracy: 0.8819	loss: 0.2711	val_accuracy: 0.6875	val_loss: 0.7837
Epoch 2/10	163/163	104s	639ms/step	accuracy: 0.8988	loss: 0.2540	val_accuracy: 0.7500	val_loss: 0.7248
Epoch 3/10	163/163	110s	676ms/step	accuracy: 0.9089	loss: 0.2356	val_accuracy: 0.7500	val_loss: 0.4994
Epoch 4/10	163/163	106s	651ms/step	accuracy: 0.9204	loss: 0.2190	val_accuracy: 0.6250	val_loss: 1.4611
Epoch 5/10	163/163	107s	659ms/step	accuracy: 0.9260	loss: 0.2006	val_accuracy: 0.8750	val_loss: 0.5118
Epoch 6/10	163/163	104s	638ms/step	accuracy: 0.9319	loss: 0.1894	val_accuracy: 0.6875	val_loss: 0.6369
Epoch 7/10	163/163	104s	637ms/step	accuracy: 0.9319	loss: 0.1800	val_accuracy: 0.7500	val_loss: 0.6222
Epoch 8/10	163/163	107s	656ms/step	accuracy: 0.9333	loss: 0.1841	val_accuracy: 0.6250	val_loss: 1.3810
Epoch 9/10	163/163	104s	638ms/step	accuracy: 0.9429	loss: 0.1647	val_accuracy: 0.6875	val_loss: 0.6427
Epoch 10/10	163/163	104s	641ms/step	accuracy: 0.9438	loss: 0.1572	val_accuracy: 0.8750	val_loss: 0.3581

**Figure 2.** Training and Validation Performance of the Proposed CNN Model over Epochs.

This Figure 2 shows the training and validation accuracy and the loss trend of every epoch in this model, and it shows the model learning behavior and convergence to performance.

Even though the accuracy of training can be improved by adding additional epochs, the chances of overfitting also increase, especially with the small validation set. Hence, a conservative period was selected to preserve generalisability. A final training accuracy of about 94% was eventually found in the training process, which meant that the training was successful and the computational efficiency of the training was appropriate to use in applications that were deployment oriented.

### 3.5. Evaluation Metrics

There are four major metrics that were used to assess model performance; accuracy, precision, recall and F1-score. Accuracy is a measure of the general proportion of correctly classified instances, which gives some broad measure of model performance. But in medical diagnostics, recall (sensitivity) is of special concern, since it is a measure of how well the model detects pneumonia cases and minimises false negatives.

Precision measures the percentage of correct predictions of all positive predictions, which indicates how many positive predictions a model makes correctly. The harmonic meaning of the recall and precision is known as the F1-score; this is used to give a balanced analysis of the classification performance. Combined, these measures provide a full measure of diagnostic efficiency.

### 3.6. Robotic Intervention System

To go beyond passive diagnosis, the research involves a simulated robotic intervention framework operationalisation of CNN model output. When pneumonia is detected, the system automatically activates a series of preprogrammed behaviors, such as the creation of alerts, automatic notification of medical staff, and the organization of oxygen assistance.

This integration symbolizes the transition between the decision support and action-based healthcare systems. The robotic module in a simulated environment interprets the output of the model and triggers the relevant response in milliseconds, which means that the delay between diagnosis and intervention is minimized. This responsiveness is especially useful in high-risk situations where early treatment is essential [17].

The suggested framework is consistent with the notion of smart hospitals, where AI-based systems will be combined with automated infrastructure to maximize efficiency and patient care. Moreover, this method can address the absence of medical capabilities in distant healthcare facilities, so preliminary intervention is possible. Although the existing implementation is simulated, it can be used as a scaled basis in real-world robotic integration.

## 4. Results

### 4.1. Training Performance

It was observed that during the training process, there was a gradual rise in the model performance, as the training accuracy has grown to about 88 in the first epoch and 94.38 by the 10th epoch. At the same time, the loss of training steadily reduced, which is evidence of successful optimisation and model convergence. Nonetheless, the validation accuracy had a significant variance, varying between 62.5 and 87.5 with equal instability in the validation loss.

This discrepancy between training and validation performance indicates moderate overfitting whereby the model fits pattern-specific to the dataset and not to the generalisable features. This is further compounded by the fact that the validation set used is very small (16 images) which makes it less reliable as a performance measure. Irrespective of such fluctuations, the general positive trend in validation accuracy to the last epoch shows that the model still has some generalisation capacity. Application wise, the model is stable enough to be used in initial screening systems but to reach clinical grade reliability, additional validation on larger datasets would be needed.

### 4.2. Test Performance

On a separate test set of 624 images, the model was found to have a test accuracy of 89.74%, which is a high generalisation performance given the instability of validation exhibited. This degree of precision is in the higher end of the standard CNN-based pneumonia detection systems which usually report accuracy between 85% and 92%.

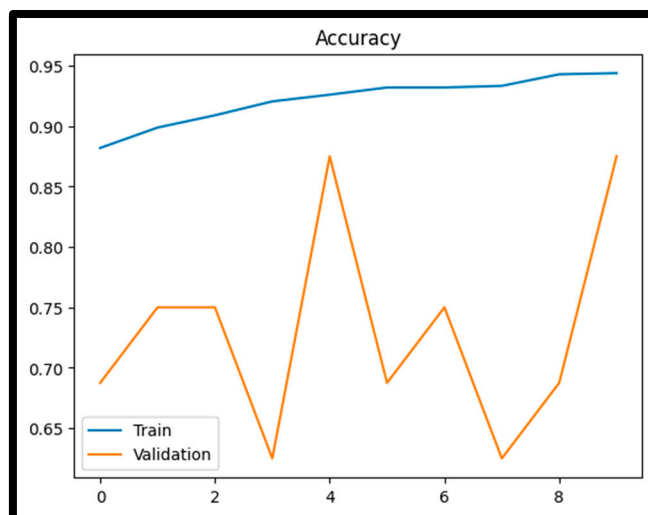


Figure 3. Training and Validation Accuracy Curve.

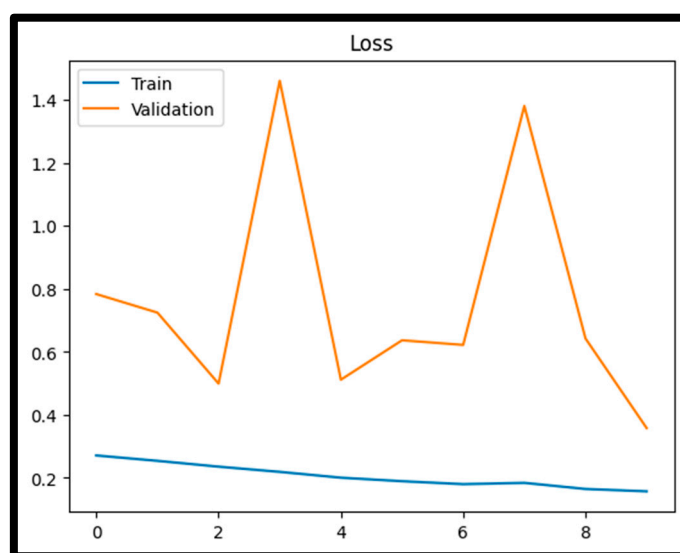


Figure 4. Training and Validation Loss Curve.

Further analysis of the classification metrics depicts a precision of 0.88 and a recall of 0.97 upon pneumonia detection which produces a F1-score of around 0.92. This high recall value is of special importance in a clinical setting; the evidence of the high recall value means that the model is effective in the detection of pneumonia cases and reduction of false negatives [18]. This feature is essential to early intervention systems, where failure to diagnose can be disastrous to health.

Nevertheless, the reduced precision is a hint that there are false positives that will cause false alarms to be raised in the real-life implementation. Although this can add workload, it is usually better than false negatives in high stakes medical cases.

#### 4.3. Confusion Matrix Analysis

The confusion matrix gives a thorough analysis of the classification performance, as illustrated in Table 1.

Table 1. Confusion Matrix Analysis.

Metric	Value
True Positive (TP)	377
True Negative (TN)	183

False Positive (FP)	51
False Negative (FN)	13

This model accurately identified 377 cases of pneumonia and 183 cases of normal with very good classification performance. False negative rate is also very low (13 cases) and this leads to a sensitivity (recall) of about 96.7 which is vital to early detection systems. This sensitivity is high enough to make sure that most of the infected patients are identified rightfully and subjected to timely intervention.




Figure 5. Confusion Matrix Analysis.

On the other hand, with 51 false positives, the specificity of 78.2 is seen, which shows that some healthy people can be mistakenly diagnosed with pneumonia [19]. This might result in unnecessary follow-up processes, but it is consistent with a risk-averse diagnostic approach that puts the patient at the centre. The trade-off is acceptable in the context of robotic intervention because precautionary responses are not as harmful as missed diagnoses.

#### 4.4. Prediction Example

One of the sample test images was properly detected as PNEUMONIA DETECTED with a confidence score of 0.999 which is a high degree of model confidence. This result shows that the system is able to produce credible predictions, which can be directly used by the robotic intervention module to respond to a clinical situation in real-time.




Figure 6. Prediction Example.

## 5. Discussion

The findings show that the vision-based system proposed has a test accuracy of 89.74%, and the recall rate is significantly high at about 96.7% on pneumonia detection. This means that the model is very useful in diagnosing positive cases which is a very important requirement in early-stage clinical intervention [20]. Application-wise, the system is more sensitive than specific, meaning that there is no risk that potentially infected patients are left unattended. This design option is consistent with risk-averse healthcare designs, in which false negatives impose much greater costs than false positives. Nevertheless, the noticed fluctuation in validation performance points out to the inherent instability during the training process, implying that the model can be vulnerable to data distribution and sample size [21]. Nevertheless, the high-test performance indicates that the model is still useful in diagnostic terms when used on unseen data.

One of the strengths of the suggested system is its balance between diagnostic performance and computational efficiency. The model has about 11.17 million parameters and yields close to 90% accuracy even without the complexity of more complicated architecture, which is suitable in real-time. The system can process and classify images in milliseconds in a simulated setting, allowing downstream responses in the robotic intervention module in real-time. The model becomes more of an active decision-support system than a passive diagnostic tool due to this integration, which allows it to issue alerts and begin initial responses [22]. This responsiveness would be especially useful in clinical environments with high demand and in remote care environments, when quick decision-making can result in a considerable impact on patient outcomes. Capability of automating initial intervention steps also lessens the workload of clinicians and improves efficiency of operations.

Despite these, a few weaknesses should be addressed with utmost importance. The first limitation is the very small validation set (16 images) which compromises the validity of validation metrics and adds to variation in performance. This is a limitation that limits the ability of the model to tune hyperparameters and amplifies the chances of overfitting. Also, the sample is restricted to paediatric cases in one source, which casts doubts on the generalisability of the results to other age groups and clinical conditions. One more significant weakness is that there is no physically realized robotic system. Although the simulation proves that automated intervention is a viable idea, it does not consider real-life limitations like hardware latency, sensor integration, and environmental variability. Therefore, the existing architecture is still a prototype and not a deployable solution.

The work in the future is to aim at improving the diagnostic and operational aspects of the system. Modelling wise, the feature extraction and the classification accuracy of more than 9295 percent might be enhanced by adopting transfer learning methods like ResNet50 [23]. Also, generalisability and bias would be minimised by increasing the range of samples to more diverse and equalised samples. Integration On the robotics side, end-to-end automation, e.g., patient monitoring and physical assistance, would be possible by integrating it with real-world platforms with frameworks like Robot Operating System (ROS) [24]. To be applied in clinical settings, the system should be subjected to stringent validation in hospital settings, with special attention to reliability, interpretability, and regulatory compliance. These developments would enable the shift of a simulated prototype to a scalable and intelligent healthcare solution.

## 6. Conclusions

This study has designed a vision-based diagnostic system that can be used to detect pneumonia early in patients using chest X-ray images, combined with a simulated robotic intervention framework. The model was trained on 5,863 paediatric radiographic images, and the test accuracy was 89.74, and the recall was high with a 96.7-rate, which means that it is good at detecting pneumonia cases. This performance is an indication that deep learning proves to be effective in assisting with analyzing medical images in a rapid and scalable manner.

In addition to detection, the main contribution of this study is the connection of diagnostic output to automated robotic response systems. The system allows timely responses like the

generation of alerts, notification to clinicians, and supportive care preparation, thus shortening the time gap between diagnosis and intervention. This integration is one of the changes in passive AI models to action-oriented healthcare systems, which is essential in any time-sensitive health case like pneumonia.

## Appendix

### Chest X-Ray Images data set:

<https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia/data>

## References

1. Application wise, the proposed framework can have tremendous potential in smart hospital setups and remote healthcare systems in which specialist expertise is not readily available. The system, though presently in the form of a simulation, forms a scalable base to deploy in the real-world, and it will promote an increase in diagnostic efficiency, decrease in clinical workload, and better patient outcomes. References Nough, M. I., & Nouran, M. Perspective of Pneumonia in the Health-Care Setting.
2. Uddin, A. Z. M. J., Begum, M. R., Akib, A. A. S., Islam, K., Hasib, A., Giri, A., & Shahi, A. (2025, December). LungNet: An Interpretable Machine Learning Framework for Early Lung Cancer Detection Using Structured Clinical Data. In *2025 IEEE 13th Conference on Systems, Process & Control (ICSPC)* (pp. 181-186). IEEE.
3. Rabbah, J., Ridouani, M., & Hassouni, L. (2025). Improving pneumonia diagnosis with high-accuracy CNN-Based chest X-ray image classification and integrated gradient. *Biomedical Signal Processing and Control*, *101*, 107239.
4. An, Q., Rahman, S., Zhou, J., & Kang, J. J. (2023). A comprehensive review on machine learning in healthcare industry: classification, restrictions, opportunities and challenges. *Sensors*, *23*(9), 4178.
5. Abad, M., Casas-Roma, J., & Prados, F. (2024). Generalizable disease detection using model ensemble on chest X-ray images. *Scientific Reports*, *14*(1), 5890.
6. Giri, A., Das, S. R., Joy, A. Z. M. J. U., Akib, A. S. M., Misat, M. M. H., Khadgi, M., ... & Shahi, B. (2025). Smart IoT Egg Incubator System with Machine Learning for Damaged Egg Detection. In *International conference on WorldS4* (pp. 236-245). Springer, Cham.
7. Mabrouk, A., Diaz Redondo, R. P., Dahou, A., Abd Elaziz, M., & Kayed, M. (2022). Pneumonia detection on chest X-ray images using ensemble of deep convolutional neural networks. *Applied Sciences*, *12*(13), 6448.
8. Arun Prakash, J., Asswin, C. R., Ravi, V., Sowmya, V., & Soman, K. P. (2023). Pediatric pneumonia diagnosis using stacked ensemble learning on multi-model deep CNN architectures. *Multimedia tools and applications*, *82*(14), 21311-21351.
9. Zhang, D., Ren, F., Li, Y., Na, L., & Ma, Y. (2021). Pneumonia detection from chest X-ray images based on convolutional neural network. *Electronics*, *10*(13), 1512.
10. Giri, A., Akib, A. A. S., Hasib, A., Acharya, A., Prithibi, M. A., Rahman, R. H., ... & Taha, H. I. C. (2025, April). Design and development of a cost effective and modular cnc plotter for educational and prototyping applications. In *2025 IEEE 4th International Conference on Computing and Machine Intelligence (ICMI)* (pp. 1-6). IEEE.
11. Kundu, R., Das, R., Geem, Z. W., Han, G. T., & Sarkar, R. (2021). Pneumonia detection in chest X-ray images using an ensemble of deep learning models. *PloS one*, *16*(9), e0256630.
12. Ayan, E., Karabulut, B., & Ünver, H. M. (2022). Diagnosis of pediatric pneumonia with ensemble of deep convolutional neural networks in chest X-ray images. *Arabian Journal for Science and Engineering*, *47*(2), 2123-2139.
13. Al-qaness, M. A., Zhu, J., AL-Alimi, D., Dahou, A., Alsamhi, S. H., Abd Elaziz, M., & Ewees, A. A. (2024). Chest X-ray Images for Lung Disease Detection Using Deep Learning Techniques: A Comprehensive Survey: MAA Al-qaness et al. *Archives of Computational Methods in Engineering*, *31*(6), 3267-3301.
14. Arora, A., Alderman, J. E., Palmer, J., Ganapathi, S., Laws, E., Mccradden, M. D., ... & Liu, X. (2023). The value of standards for health datasets in artificial intelligence-based applications. *Nature medicine*, *29*(11), 2929-2938.

15. Giri, A., Hasib, A., Islam, M., Tazim, M. F., Rahman, M. S., Khadgi, M., & Akib, A. A. S. (2025, June). Real-time human fall detection using yolov5 on raspberry pi: An edge ai solution for smart healthcare and safety monitoring. In *International Conference on Data Analytics & Management* (pp. 493-507). Cham: Springer Nature Switzerland.
16. Prakash, J. A., Ravi, V., Sowmya, V., & Soman, K. P. (2023). Stacked ensemble learning based on deep convolutional neural networks for pediatric pneumonia diagnosis using chest X-ray images. *Neural Computing and Applications*, 35(11), 8259-8279.
17. An, Q., Chen, W., & Shao, W. (2024). A deep convolutional neural network for pneumonia detection in x-ray images with attention ensemble. *Diagnostics*, 14(4), 390.
18. Rao, S., Zeng, Z., & Zhang, J. (2025). Robust Multiclass Pneumonia Classification via Multi-Head Attention and Transfer Learning Ensemble. *Applied Sciences*, 15(21), 11426.
19. Račić, L., Popović, T., & Šandi, S. (2021, February). Pneumonia detection using deep learning based on convolutional neural network. In *2021 25th International Conference on Information Technology (IT)* (pp. 1-4). IEEE.
20. Hussain, A., Amin, S. U., Lee, H., Khan, A., Khan, N. F., & Seo, S. (2023). An automated chest X-ray image analysis for covid-19 and pneumonia diagnosis using deep ensemble strategy. *IEEE Access*, 11, 97207-97220.
21. Masud, M., Bairagi, A. K., Nahid, A. A., Sikder, N., Rubaiee, S., Ahmed, A., & Anand, D. (2021). A pneumonia diagnosis scheme based on hybrid features extracted from chest radiographs using an ensemble learning algorithm. *Journal of Healthcare Engineering*, 2021(1), 8862089.
22. Kheirdoust, A., Barzanouni, F., Rasoulia, A., Behrouzi, F., Esmailzadeh, A., Ghaddaripouri, K., & Mazaheri Habibi, M. R. (2025). Evaluation of Machine Learning Methods Developed for Prediction and Diagnosis of Pneumonia: A Systematic Review. *Health Science Reports*, 8(12), e71446.
23. Sharma, S., & Guleria, K. (2023). A deep learning based model for the detection of pneumonia from chest X-ray images using VGG-16 and neural networks. *Procedia Computer Science*, 218, 357-366.
24. Akib, A. A. S., Giri, A., Islam, M., Sifa, F. J., Elahi, T. A., Aktia, A. N., ... & Khanna, A. (2024, October). Design and simulation of a quadruped robot. In *International Conference on Data-Processing and Networking* (pp. 373-385). Singapore: Springer Nature Singapore.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.