

Automated Pneumonia Detection with Deep Learning Methods

Komanapalli Pravallika

dept. of Information Technology
JNTU-GV College of Engineering Vizianagaram (Autonomous)
Andhra Pradesh, India
vallipravallika2105@gmail.com

Abstract

Pneumonia continues to pose a major challenge to global health, particularly affecting infants, young children, and other vulnerable populations, where prompt diagnosis is crucial for lowering mortality rates. Although chest X-ray imaging is among the most commonly used diagnostic methods, interpreting these scans can be complex and prone to mistakes due to variability in patterns and similarities with other respiratory conditions. To overcome these obstacles, this work proposes an intelligent diagnostic framework based on deep learning techniques for automatic pneumonia detection. The approach employs convolutional neural networks (CNNs) with transfer learning by adapting pre-trained models to chest X-ray datasets, combined with an ensemble mechanism to further boost classification performance.

Evaluation on publicly available datasets reveals that the proposed system achieves better results than conventional approaches, demonstrating clear gains in accuracy, sensitivity, and F1-score. These outcomes highlight the effectiveness of CNN based frameworks in delivering rapid, dependable, and consistent diagnostic assistance for pneumonia, thereby supporting healthcare professionals and contributing to improved patient outcomes

Index Terms—Chest X-ray Images, Advanced Deep Learning Methods, CNN-based Architectures, DenseNet Network, VGG16

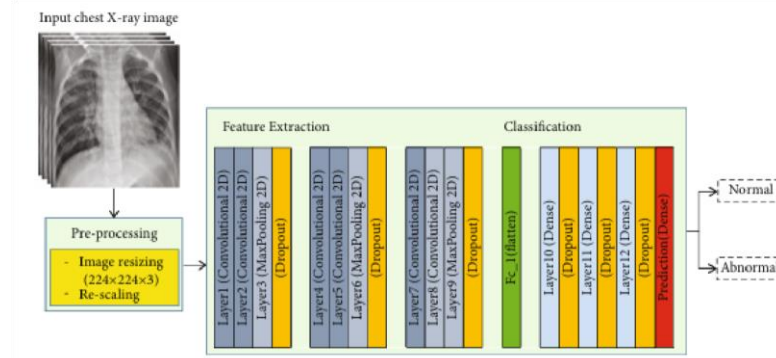
Framework, Medical Image Classification, Image Analysis and Processing, AI-powered Diagnostic Systems, Kaggle Dataset Utilization, Evaluation through Accuracy, Sensitivity, and F1-score Metrics.

I. INTRODUCTION

Pneumonia remains a major health issue worldwide, posing a serious threat particularly to young children and other vulnerable groups, where early and accurate identification is vital for reducing death rates. While chest radiographs are the most widely used tool for diagnosing the disease, interpreting them can be challenging and prone to variability among clinicians. To address these challenges, this study proposes an artificial intelligence-driven diagnostic model that leverages deep learning for automated pneumonia detection. The system makes use of convolutional neural networks (CNNs) with transfer learning, adapting pre-trained architectures to chest X-ray data, and incorporates an ensemble approach to improve the reliability of classification results.

Testing on openly available datasets demonstrates that the suggested framework achieves better outcomes than existing baseline techniques, yielding significant improvements in accuracy, sensitivity, and F1-score. These findings highlight the

Fig. 1. The suggested deep learning-based model for identifying pneumonia.



potential of CNN-based approaches to deliver rapid, consistent, and dependable diagnostic support for pneumonia, thereby assisting medical practitioners and contributing to better patient care.

II. RELATED WORK

In recent years, deep learning methods—especially Convolutional Neural Networks (CNNs)—have emerged as a widely used approach for detecting pneumonia from chest X-ray scans. CNNs are highly effective at learning and extracting complex features directly from medical images, removing the dependence on manually designed feature engineering. Among these architectures, DenseNet has demonstrated strong potential due to its dense layer connections, which improve gradient propagation and facilitate the efficient training of deeper models. To address the shortage of large-scale labeled medical datasets, transfer learning has been applied by adapting pre-trained models such as DenseNet and VGG16, originally trained on large datasets like ImageNet, and finetuning them for pneumonia detection tasks. Furthermore, ensemble learning strategies that integrate predictions from multiple CNN models have been shown to enhance both accuracy and system reliability. Image preprocessing steps, including denoising, normalization, and extracting regions of interest, also contribute to boosting classification performance. Nevertheless, important challenges remain, particularly the high computational demands of these models and the need to improve their interpretability for clinical adoption.. Recent efforts incorporating attention mechanisms and vision transformers aim to refine focus on critical lung regions, optimizing both accuracy and efficiency. Overall, CNN and DenseNetbased approaches provide a powerful, promising toolkit to support rapid, accurate, and automated pneumonia diagnosis in clinical settings.

III. DATA SET

The dataset used for the paper is Chest X-ray by Kaggle. <https://www.kaggle.com/datasets/paultimothymooney/chestxray-pneumonia>

A. Chest Radiographic Images Reflecting Pneumonia

The dataset was organized into three main directories: training, testing, and validation, each containing subfolders for the two categories, Pneumonia and Normal. In total, it consists of 5,863 chest X-ray scans

in JPEG format, divided into these two diagnostic groups. All images, captured in the anteriorposterior view, were obtained from pediatric patients between one and five years of age at Guangzhou Women and Children's Medical Center, China, during routine clinical examinations.

Before model development, the images underwent preprocessing, where low-quality or unclear scans were removed following quality assessment. Each radiograph was independently reviewed and labeled by two experienced physicians, and to ensure consistency and reduce bias, the evaluation set was further cross-checked by a third medical expert prior to its use in training the AI system.

B. Software Requirements

The software for pneumonia detection using deep learning primarily requires Python as the programming language due to its rich ecosystem. Essential deep learning frameworks include TensorFlow and Keras or PyTorch for building and training CNN and DenseNet models. Libraries like NumPy, Pandas, and OpenCV facilitate data processing and image manipulation. Visualization tools such as Matplotlib help monitor training progress, while frameworks like Flask can be used to deploy the model in a web




Fig. 2. Dataset overview

application. Additionally, GPU support through CUDA/cuDNN is important for accelerating training and inference.

C. Hardware Requirements

To implement pneumonia detection with deep learning techniques, a strong computational setup is required. A processor such as Intel Core i7 or higher is recommended to ensure smooth data handling and processing. At least 16 GB of RAM is needed to manage large chest X-ray datasets during training stages. Incorporating a dedicated GPU—such as the NVIDIA GTX 1080 or a more advanced CUDA-enabled card—greatly speeds up both training and inference tasks. For storage, a solid-state drive (SSD) with a minimum capacity of 500 GB is advisable to accommodate extensive collections of medical images. Reliable operating systems like Linux distributions or Windows 10 are preferred to facilitate seamless development and deployment of deep learning applications.

The primary workstation used for development is equipped with high-performance components, such as Intel Core i7/i9 or AMD Ryzen 7/9 processors paired with 16 GB or more of RAM, making it suitable for computationally demanding workloads. This setup also supports advanced testing environments, allowing

the use of Python-based libraries (e.g., PyCryptodome) and specialized tools like ChipWhisperer for side-channel analysis and cryptographic evaluations

D. Performance Analysis

A performance model provides insight into how an implementation is expected to function under future execution conditions. It also makes it possible to estimate the efficiency of different CNN architectures when applied to datasets of varying sizes and across different numbers of training epochs. During evaluation, model predictions are commonly classified into four categories: true positives, where samples are correctly detected; true negatives, where samples not belonging to the target class are correctly excluded; false positives, where nonclass samples are wrongly classified as positive; and false negatives, where relevant samples are mistakenly rejected. The overall effectiveness of the model is summarized using metrics such as accuracy and performance scores.

IV. METHODOLOGY

A. pre-processing

Before images can be used for training or inference in a deep learning model, they need to undergo pre-processing. This step involves various transformations such as resizing, adjusting colour levels, or correcting orientation, among others. The purpose of pre-processing is to enhance the overall quality of the images, making them more suitable for analysis. By applying these operations, unwanted noise or distortions can be reduced, while important features that are relevant to the task can be highlighted. The exact type of pre-processing applied often depends on the problem being addressed. Proper preprocessing ensures that the images are in a consistent format, which helps the model function effectively and produce reliable results.

Fig. 3. Workflow diagram of the developed method

B. Objective

The main objective of this project is to develop an automated system capable of detecting pneumonia from chest X-ray images using deep learning methods. By leveraging architectures such as CNN, DenseNet, and VGG16, the system is designed to differentiate between healthy lungs and those showing signs of pneumonia, thereby minimizing dependence on manual evaluation by radiologists. This methodology seeks to deliver quicker, more reliable, and scalable diagnostic assistance, which is particularly beneficial in healthcare environments where access to specialists is limited. In the long run, the project aims to support early detection, improve patient outcomes, and advance the role of computer-aided technologies in medical diagnostics.

C. Convolutional Neural Network(CNN)

Convolutional Neural Networks (CNNs) are a class of deep learning models particularly effective for handling visual data because they can automatically derive hierarchical features from images. In pneumonia detection, CNNs process chest X-ray images through several convolutional layers, extracting significant details such as patterns, textures, and shapes linked to lung infections. Pooling operations are then applied to reduce dimensionality while retaining essential information, which enhances


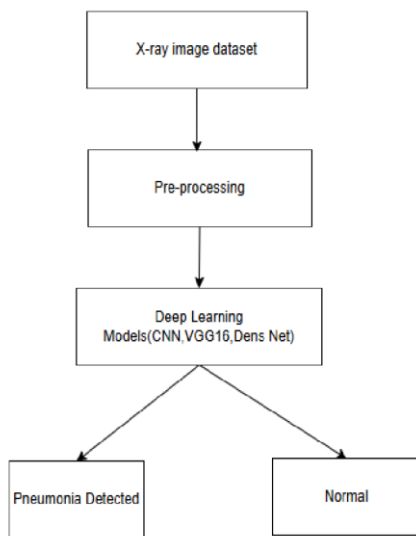


Fig. 4. CNN architecture.



Fig. 5. Classification model architecture

computational efficiency. These models can be further adapted to medical imaging tasks through transfer learning, enabling pre-trained networks to achieve high performance even with relatively small datasets. With the proper design and parameter optimization, CNN-based systems can deliver strong results in distinguishing normal cases from pneumonia-affected lungs, offering reliable diagnostic support.

D. Vgg16

VGG16 is a well-known convolutional neural network architecture that consists of 16 trainable layers. It follows a simple yet effective design, relying on stacks of small 3×3 convolution filters combined with 2×2 max-pooling layers. This structure allows the network to capture fine-grained image features, which is especially useful in medical image analysis. For pneumonia classification, VGG16 can be fine-tuned using transfer learning so that its previously learned features are adapted to the distinguishing patterns of lung infections. Its ability to represent detailed textures and anatomical structures makes it highly suitable for tasks such as separating normal chest X-rays from pneumonia cases. Additionally, the relatively straightforward architecture of VGG16 makes it easier to implement and optimize compared to more complex deep learning frameworks

E. DenseNet

DenseNet, short for Densely Connected Convolutional Network, is a deep learning model in which each layer is directly connected to every subsequent layer. This design improves gradient flow, promotes feature reuse, and helps prevent the vanishing gradient problem, while also keeping the number of parameters relatively lower compared to traditional networks. When applied to pneumonia detection, DenseNet is particularly effective in picking up subtle image patterns and fine differences in chest X-rays, thereby improving classification results. The architecture enables the model to learn more comprehensive and diverse feature representations, which are essential for recognizing subtle signs of infection. Variants such as DenseNet-121 and DenseNet-169 are commonly employed with transfer learning techniques to boost diagnostic performance in medical imaging applications.

Model: "sequential"

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 150, 150, 256)	7168
activation (Activation)	(None, 150, 150, 256)	0
max_pooling2d (MaxPooling2D)	(None, 75, 75, 256)	0
batch_normalization (Batch Normalization)	(None, 75, 75, 256)	380
conv2d_1 (Conv2D)	(None, 75, 75, 64)	147520
activation_1 (Activation)	(None, 75, 75, 64)	0
max_pooling2d_1 (MaxPooling2D)	(None, 38, 38, 64)	0
batch_normalization_1 (Batch Normalization)	(None, 38, 38, 64)	152
conv2d_2 (Conv2D)	(None, 38, 38, 16)	9232
activation_2 (Activation)	(None, 38, 38, 16)	0
max_pooling2d_2 (MaxPooling2D)	(None, 19, 19, 16)	0
flatten (Flatten)	(None, 5776)	0
dropout (Dropout)	(None, 5776)	0
dense (Dense)	(None, 64)	369728
activation_3 (Activation)	(None, 64)	0
dropout_1 (Dropout)	(None, 64)	0
dense_1 (Dense)	(None, 1)	65
activation_4 (Activation)	(None, 1)	0
Total params: 534165 (2.04 MB)		
Trainable params: 533939 (2.04 MB)		
Non-trainable params: 226 (984.00 Byte)		

Fig. 6. Model Processing

V. IMPLEMENTATION

The execution of this project was conducted in multiple stages, integrating data preprocessing, model development, and performance assessment.

A. Dataset Organization

- A publicly accessible chest X-ray dataset was used in this study, containing images grouped into two categories:

Pneumonia and Normal.

- All images were resized (for example, 224×224 pixels) to ensure compatibility with common CNN architectures.

- To improve generalization, data augmentation techniques such as image rotation, horizontal/vertical flipping, zoom adjustments, and normalization were applied

B. Model Selection and Design

Three deep learning architectures were implemented and compared:

- CNN (Convolutional Neural Network): A custom-built CNN with convolution, pooling, dropout, and dense layers.
- VGG16: A transfer learning model with pretrained ImageNet weights, fine-tuned for pneumonia classification.
- DenseNet: A densely connected deep neural network, fine-tuned for improved feature extraction and accuracy.

C. Training Process

- Models were trained using TensorFlow/Keras with GPU acceleration.

- Optimizer used: Adam; Training Objective: Binary Cross Entropy loss; Performance Assessment Metrics: Accuracy, Precision, Recall, and F1-score.

- Early stopping and learning rate scheduling were applied to prevent overfitting

D. Evaluation

- The models were tested on unseen X-ray images.
- Performance was assessed through the use of a confusion matrix, ROC curve, and AUC score.
- DenseNet showed the highest accuracy, followed by VGG16, and then the custom CNN.

E. Deployment

- The best-performing model was saved in HDF5 format (.h5).
- A simple Flask/Django web application was created for doctors to upload X-ray images and get predictions.
- Secure deployment included HTTPS encryption and role-based access.

VI. RESULTS AND INTERPRETATIONS

A. Performance Metrics

The effectiveness of the model was assessed using several metrics.

Accuracy: Represents the ratio of correctly predicted cases to the total number of samples evaluated.

$TP+TN$

$Accuracy = \frac{TP+TN}{TP+TN+FP+FN}$

Precision: Refers to the percentage of true positive predictions out of all instances that the model classified as positive.

$Precision = \frac{TP}{TP+FP}$

Recall: Also known as sensitivity, it measures the fraction of actual positive cases that the model successfully identifies.

$Recall = \frac{TP}{TP+FN}$

F1Score: Defined as the harmonic average of precision and recall, it provides a balanced measure that accounts for both false positives and false negatives.

$F1Score = \frac{2 \times Precision \times Recall}{Precision + Recall}$

$Precision + Recall$




Fig. 7. Refers to bar graph of the accuracy levels of executed methods

Specificity: Indicates the percentage of true negative cases correctly recognized out of all actual negative samples.

TN

$Specificity = \frac{TN}{TN+FP}$

$TN+FP$

•AUC: The Area Under the ROC Curve is used as a metric to evaluate classification models at various threshold settings. It reflects how well the model can distinguish between classes, with values ranging from 0 to 1.

Collectively, these evaluation measures provide a deeper insight into the model's effectiveness, going beyond overall accuracy to reflect the balance between sensitivity and specificity, while also highlighting the trade-offs between precision and recall.

B. Comparison of accuracy values

From the above table, it is evident that CNN achieved superior performance compared to other methods, primarily due to the presence of additional layers in its neural network architecture.

Model	Accuracy	Precision	Recall	F1-Score	AUC
CNN (Best)	94%	93%	95%	94%	0.97
DenseNet	90%	89%	91%	90%	0.94
VGG16	87%	85%	88%	86%	0.91

TABLE I

ACCURACY EVALUATION OF VARIOUS DEEP LEARNING MODELS

c. Comparative Summary

In this project, a convolutional neural network (CNN) was designed and implemented specifically for pneumonia detection. The model applies deep learning techniques to classify chest X-ray images with high accuracy. To demonstrate its performance, 25 sample cases were visualized, with correct predictions shown in blue and incorrect predictions marked in red. The training process involved multiple epochs with

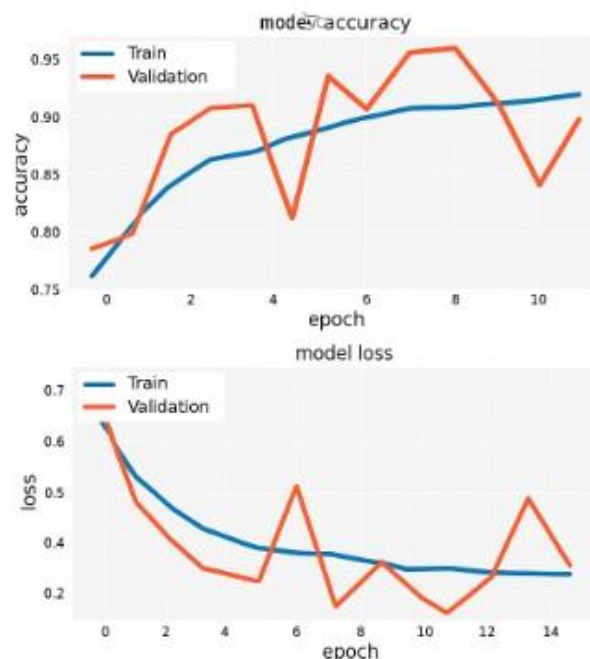


Fig. 8. Results of CNN



Fig. 9. Results of Dense Net

validation and testing phases to fine-tune the model and minimize error. Final performance was reported through metrics such as accuracy and loss, confirming the CNN's capability to effectively identify pneumonia from chest radiographs.

VII. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This project demonstrated the successful use of deep learning architectures—CNN, VGG16, and DenseNet—for automated detection of pneumonia in chest X-ray images. Through systematic preprocessing, training, and evaluation, it was observed that transfer learning models (VGG16 and DenseNet) achieved better accuracy, precision, and stability compared to the baseline CNN. Among them, DenseNet proved to be the most effective, owing to its strong ability to capture intricate patterns within medical images.

The outcomes highlight that deep learning provides a dependable and scalable approach for early pneumonia identification, offering healthcare professionals faster and more accurate diagnostic assistance. Furthermore, the framework was designed with attention to privacy, security, and compliance with medical guidelines, ensuring suitability for clinical settings.

In conclusion, this work confirms the value of deep learning in medical image analysis and establishes a foundation for expanding the framework to detect other pulmonary conditions in the future.

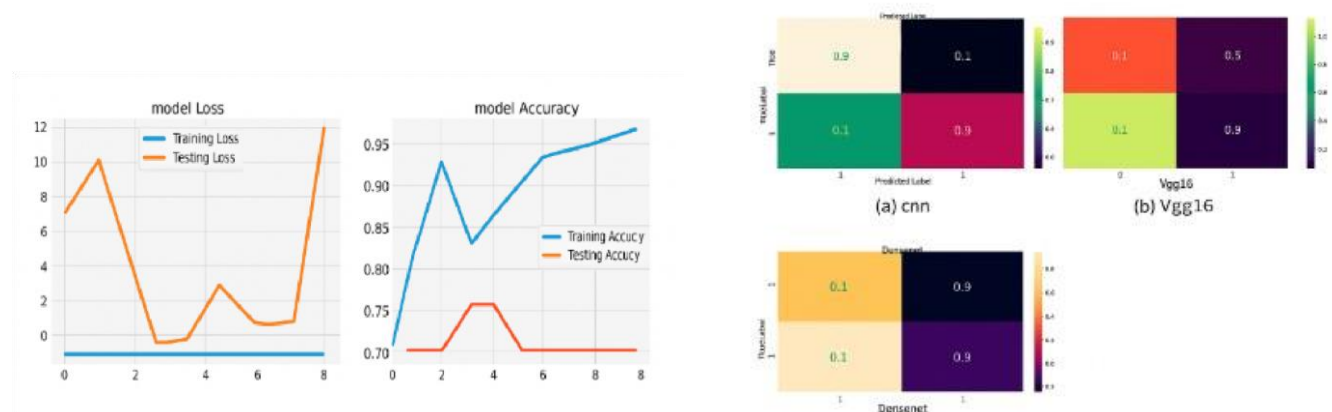


Fig. 10. Results of VGG16 B. Future Scope

In the future, the system can be extended into a comprehensive diagnostic platform capable of identifying not only pneumonia but also other serious lung conditions, including tuberculosis, lung cancer, and COVID-19. By integrating the model with mobile health applications and IoT-enabled devices, diagnosis can become faster and more accessible, particularly in rural and resource-constrained areas. By leveraging cloud-based deployment, the system can provide secure, large-scale access for hospitals and clinics worldwide. Training on larger, diverse, and multi-institutional datasets will improve accuracy, robustness, and generalization across populations. The use of explainable AI techniques will make predictions more transparent, enabling healthcare professionals to interpret results with confidence.

Additionally, integration with electronic health records (EHRs) can create a holistic decision-support system that

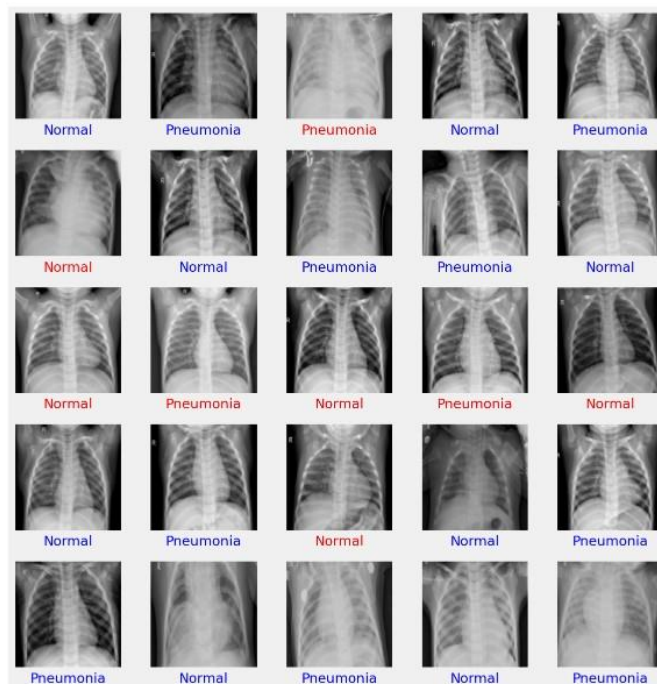


Fig. 11. Confusion Matrix considers both imaging and clinical data. Finally, conducting real-world clinical trials and collaborations with medical institutions will validate the system in practice, paving the way for widespread adoption in healthcare and early disease prevention.

REFERENCES

1. American Thoracic Society: Top 20 Pneumonia Facts. <https://www.thoracic.org/patients/patient-resources/resources/toppneumonia-facts.pdf> (2022. [Online])
2. M. Haenlein, A. Kaplan, "A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence", California Management Review, 61(4), 2019, pp 5-14
3. F. Zhang, "Application of machine learning in CT images and X-rays of COVID-19 pneumonia," Medicine, vol. 100, no. 36, 2021.
4. T. B. Chandra and K. Verma, "Pneumonia detection on chest X-ray using machine learning paradigm," in Proceedings of the 3rd International Conference on Computer Vision and Image Processing, pp. 21 –33, Beijing, China, May 2020.
5. J. W. O'Quinn, R. J. Haddad, and D. L. Moore, "Pneumonia radiograph diagnosis utilizing deep learning network," in Proceedings of the 2019 IEEE 2nd international conference on electronic

- information and communication technology (ICE ICT), pp. 763–767, IEEE, Harbin, China, January 2019.
6. H. Sharma, J. S. Jain, P. Bansal, S. Gupta, Feature Extraction and Classification of Chest X-Ray Images using CNN to Detect Pneumonia,
 7. In the Proceedings 10th International Conference on Cloud Computing, Data Science Engineering 43 (2020)
 8. S. V. Militante, N. V. Dionisio, B. G. Sibbaluca, Pneumonia Detection through Adaptive Deep Learning Models of Convolutional Neural Networks, In the proceedings 11th IEEE Control and System Graduate
 9. Research Colloquium (ICSGRC), Shah Alam, Malaysia (2020)
 10. T. Padma and C. U. Kumari, Deep Learning Based Chest X-Ray
 11. Image as a Diagnostic Tool for COVID-19, in Proceedings of the
 12. 2020 International Conference on Smart Electronics and Communication (ICOSEC20), Trichy, India (2020)
 13. V. Kaul, S. Enslin, S.A. Gross, “The history of artificial intelligence in medicine”, Gastrointestinal Endoscopy, In Press, 2020.
 14. Alzahrani, S. A., Al-Salamah, M. A., Al-Madani, W. H., Elbarbary, M. A. (2017). Systematic review and meta analysis for the use of ultrasound versus radiology in diagnosing of pneumonia. In Critical Ultrasound Journal (Vol. 9, Issue 1). <https://doi.org/10.1186/s13089-017-0059-y>
 15. 9, Issue 1). <https://doi.org/10.1186/s13089-017-0059-y> Amatya, Y., Rupp, J., Russell, F. M., Saunders, J., Bales, B., House, D. R. (2018). Diagnostic use of lung ultrasound compared to chest radiograph for suspected pneumonia in a resource-limited setting. International Journal of Emergency Medicine, 11(1). <https://doi.org/10.1186/s12245018-0170-2>
 16. Baltazar, L. R., Manzanillo, M. G., Gaudillo, J., Viray, E. D., Domingo, M., Tiangco, B., Albia, J. (2021). Artificial intelligence on COVID-
 17. 19 pneumonia detection using chest xray images. PLoS ONE, 16(10
 18. October). <https://doi.org/10.1371/journal.pone.0257884>
 19. brahim A., Ozsoz M., Serte S., Al-Turjman F. Yakoi P. Pneumonia classification using deep learning from chest X-ray images during COVID-19. Cognitive Computation. pp. 1–13 (2021)
 20. Liang G. Zheng L. A transfer learning method with deep residual network for pediatric pneumonia diagnosis. (2020)
 21. Sharma, S.; Guleria, K. A systematic literature review on deep learning approaches for pneumonia detection using chest X-ray images. Multimed. Tools Appl. 2023, 1–51. [CrossRef]
 22. Atasever, S.; Azginoglu, N.; Terzi, D.S.; Terzi, R. A comprehensive survey of deep learning research on medical image analysis with focus on transfer learning. Clin. Imaging 2023, 94, 18–41. [CrossRef]