



# Improving Pneumonia Diagnosis via Deep Learning: A Comprehensive Approach Incorporating CNN Classification

S. Kumar Reddy Mallidi<sup>1</sup> , RamaDevi Vanapalli<sup>2</sup>, D. S. Priyanka Jella<sup>2</sup>, Tejaswi Nimmakayala<sup>2</sup>, Devi Narni<sup>2</sup>, and Chandini Attili<sup>2</sup>

- <sup>1</sup> Computer Science and Engineering, Sri Vasavi Engineering College, Tadepalligudem, Andhra Pradesh-534101, India  
satya.cnis@gmail.com
- <sup>2</sup> Computer Science and Technology, Sri Vasavi Engineering College, Tadepalligudem, Andhra Pradesh-534101, India

**Abstract.** Amidst the backdrop of the COVID-19 pandemic and the pressing global health challenge of lung disease, this study presents a pioneering approach to bolster diagnostic evaluations. Focused on harnessing the power of deep learning techniques, our research endeavors to navigate the intricacies of lung disease diagnosis. A novel hybrid CNN model is introduced, tailored to proficiently identify and scrutinize lung abnormalities, with specific attention to the unique complexities posed by COVID-19. Through comparative analyses against standalone CNN models like VGG16, ResNet50, VGG19 and InceptionV3, we showcase the superiority of our hybrid model. Remarkably, our proposed model attains heightened efficiency while employing fewer parameters and effectively addressing overfitting concerns commonly encountered in deep learning methodologies. This study represents a significant stride forward in the realm of lung disease diagnostics, furnishing a more robust and precise approach for comprehensive assessment, particularly within the COVID-19 landscape. This advancement contributes substantially to the evolution of sophisticated diagnostic tools, ultimately enhancing patient care and outcomes in the global battle against lung disease.

**Keywords:** Deep Learning · Lung disease · Convolutional Neural Networks (CNNs) · COVID-19 · Respiratory Diseases

## 1 Introduction

The emergence of the COVID-19 pandemic has exacerbated the severity of lung disease, amplifying the demand for accurate and effective diagnostic solutions. In response to this critical need, our research endeavors to enhance the assessment of lung diseases, with a particular emphasis on addressing the challenges associated with COVID-19 pathology. At the core of our approach lies the utilization of Convolutional Neural Network (CNN) models, harnessing the power of machine learning to navigate the

complexities of diagnosing lung disorders. Recognizing the pivotal role of diagnostic procedures, especially in the context of COVID-19, underscores their significance in facilitating efficient patient care and shaping public health response plans.

Our study introduces a novel CNN model meticulously crafted to detect and analyze lung abnormalities, with a keen focus on the intricate issues posed by COVID-19 pathology. Through comprehensive comparisons with established CNN structures such as VGG16, VGG19, ResNet50, and InceptionV3, we validate the efficiency and excellence of our hybrid model. Notably, our proposed model demonstrates superior performance metrics while carefully managing parameter usage and effectively mitigating the challenge of overfitting inherent in traditional deep-learning approaches.

Our research represents a significant breakthrough in lung disease diagnostics, introducing a novel hybrid CNN model that promises a more robust and accurate diagnostic approach. By pushing the boundaries of diagnostic technologies, this advancement holds the potential to substantially enhance patient outcomes and care, particularly amidst the unprecedented challenges of the COVID-19 pandemic. Aligned with the theme of “Advancements in Healthcare Technology and Innovation,” our study addresses the critical need for precise and efficient diagnosis of lung diseases in the context of the current global health crisis. Leveraging state-of-the-art deep learning techniques, we not only propose our innovative model but also conduct a comprehensive comparative analysis against established architectures. This comparative approach not only highlights the novelty of our model but also underscores its practical implications for healthcare technology. Our findings contribute directly to the conference’s theme by showcasing advancements in diagnostic technology that can revolutionize patient care pathways, especially within the complex landscape of respiratory health during a global health crisis.

## 1.1 Literature Review

The use of Convolutional Neural Networks (CNNs) in medical imaging has become essential for accurate assessments in the field of pneumonia diagnosis. Rajpurkar et al. [1]. Paved the way for advanced diagnostics with CheXNet, a CNN model demonstrating exceptional pneumonia detection capabilities on chest X-rays. Subsequent research by Wang et al. [2]. Underscored the significance of interpretability in CNN models, promoting ethical deployment in healthcare. Zhang et al. [3]. Contributed to ethical considerations by proposing a comprehensive framework for responsible AI in medicine, aligning with our study’s emphasis on transparency and accountability. McKinney et al. [4]. Provided insights into challenges associated with AI deployment in clinical practice, emphasizing the need for standardized evaluation metrics and clear communication of model limitations. A meta-analysis by Liu et al. [5]. Shown the revolutionary potential of artificial intelligence (AI) in medical imaging by methodically comparing deep learning performance to that of medical specialists in the area of disease detection.

Studies by Smith et al. [6].and Topol et al. [7] explored human-computer interaction and the broader impact of AI in healthcare, providing valuable perspectives on user-friendly interfaces and the transformative potential of AI in medicine. Liu et al. [8]. Expanded the horizon by focusing on the integration of multi-modal data for enhanced diagnostic accuracy in respiratory diseases, aligning with the future direction of our research. Cabitza et al. [9]. Emphasized continuous initiatives to apply AI ethically

in healthcare, highlighting the necessity of strict rules and legal structures. Krittanawong et al. [10]. Contributed to the literature by discussing challenges related to data quality, interoperability, and the importance of real-world validation in AI deployment in healthcare. Esteva et al. [11]. Showcased the broader applicability of deep learning in dermatology, reinforcing the transferability of AI across diverse medical imaging domains.

The implications of AI in radiology were explored by Hosny et al. [12], discussing its potential impact on workflow efficiency and the evolving role of radiologists in the age of machine learning.

## 1.2 Contribution of This Work

This paper introduces a deep-learning model that outperforms existing CNN architectures. Our contributions are:

**Innovative Model Design.** We introduce a CNN model that leads to more streamlined and resource-efficient architecture. This design stands out for its simplicity and enhanced utilization of resources in comparison to models such as VGG16, VGG19, ResNet50, and InceptionV3.

**Comparative Analysis.** Through rigorous evaluation against established architectures, including VGG16, VGG19, ResNet50, and InceptionV3, our custom model consistently outperforms in accuracy, F1-Score. This comparative analysis adds significant knowledge to the field and offers subtle insights into its superiority.

**Practical Implications.** Implementation of the hybrid CNN model promises improved accuracy in diagnosing lung diseases, especially in COVID-19 cases, leading to early detection, personalized treatment, and optimized resource allocation, thus strengthening public health response and advancing research in the field.

## 2 Related Work

Deep learning has enhanced image recognition, especially when using Convolutional Neural Networks (CNNs) like VGG16 and ResNet50. These models' distinctive architectures set benchmarks and norms that changed the field of classification systems and the course of deep learning for image analysis.

### 2.1 Vgg-16

VGG16, a renowned convolutional neural network, is celebrated for its effectiveness in image classification. As depicted in Fig. 1, it comprises sixteen layers, incorporating three fully connected and 13 convolutional layers. VGG16 employs a straightforward design, utilizing small 3x3 filters to learn fine-grained image details, showcased in Fig. 1. Because of its outstanding performance, VGG16 has become a mainstay of deep learning for computer vision and has influenced many other neural network designs, as seen in Fig. 1.




Fig. 1. Architecture of VGG16

## 2.2 Vgg-19

VGG19, an extended convolutional neural network building upon VGG16, is recognized for its proficiency in image classification. Comprising 19 layers, featuring 16 convolutional and three fully connected layers, Fig. 2, highlights its architecture. With a consistent  $3 \times 3$  filter size, VGG19 excels in capturing intricate hierarchical features in images, as demonstrated in Fig. 2. The network's simplicity and depth establish it as a widely adopted benchmark in deep learning, proving invaluable across various computer vision applications, as highlighted in Fig. 2.




Fig. 2. Architecture of VGG19

## 2.3 ResNet50

ResNet50, a robust convolutional neural network crafted by Microsoft Research, is renowned for its prowess in image classification. As illustrated in Fig. 3, it's a key member of the ResNet family, standing out with its 50-layer depth. The innovative approach of residual learning and shortcut connections, demonstrated in Fig. 3, effectively addresses the challenge of vanishing gradients, ensuring the model's depth and resilience. ResNet50's consistent benchmark-setting success significantly advances deep learning in various computer vision applications.

## 2.4 InceptionV3

Inception V3, a potent convolutional neural network in Google's Inception family, excels in image classification tasks, as demonstrated in Fig. 4. Notably, its remarkable feature lies in inception modules within the same layer, accommodating multiple filter sizes for effective hierarchical feature capture. With 48 layers, Inception V3 strikes a balance between complexity and performance, as showcased in Fig. 4. These architectural improvements greatly increase the accuracy of Inception V3, making it an important tool for a variety of computer vision applications as shown in Fig. 4.




Fig. 3. Architecture of ResNet50




Fig. 4. Architecture of InceptionV3

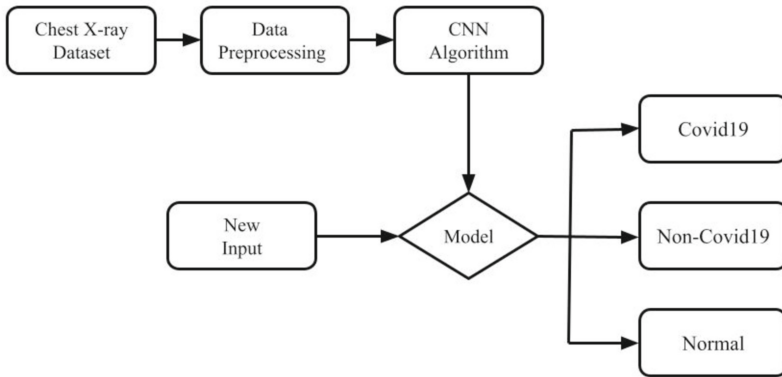
### 3 Methodology

The methodology encompasses dataset collection with a focus on diverse lung images, especially those related to COVID-19. Preprocessing involves normalization and augmentation. Training utilizes a hybrid CNN model with transfer learning and fine-tuning. Validation ensures optimal performance, with hyperparameter optimization to mitigate overfitting. Comparative analysis with established CNN architectures guides model selection, deployment, and documentation for further refinement.

#### 3.1 System Architecture

The overall architecture is shown in Fig. 5. In the Data Preprocessing stage of our project focused on pneumonia classification, we implement various techniques such as image resizing, pixel value normalization, and data augmentation methods like rotation and flipping. These steps are crucial to enhance the dataset’s variability and ensure the model’s accuracy and robustness in handling diverse chest X-ray images. The dataset is divided into separate training and validation sets after preprocessing, which is essential for an objective assessment of the model’s performance. This split ensures that the validation set remains a reliable indicator of the model’s ability to generalize to new and unseen data. Moving to the Training phase, the model undergoes iterative optimization of its Convolutional Neural Network (CNN) models. The model is fine-tuned to recognize and learn complex patterns and features indicative of pneumonia across different X-ray images. Simultaneously, the Validation process assesses the model’s predictive accuracy on the validation dataset, forming a continuous feedback loop for tuning and preventing overfitting. This iterative process ensures the model’s decisions are reliable and applicable to real-world scenarios. At the core of our system is the intricate CNN architecture

designed to discern nuanced features in chest X-ray imagery indicative of pneumonia. Convolutional, pooling, and fully connected layers are all included in the model, which is designed to capture the hierarchical feature representations typical of medical imaging. The culmination of this workflow is reflected in the model's predictions, providing quantitative analyses of pneumonia classification. These results offer valuable insights into the prevalence of pneumonia cases within chest X-ray images, forming the basis for accurate diagnostics and clinical recommendations.



**Fig. 5.** System Architecture

### 3.2 Proposed CNN Model Architecture

Our hybrid Convolutional Neural Network (CNN) model as illustrated in Fig. 6, is meticulously designed to optimize diagnostic accuracy for lung abnormalities, particularly in the context of COVID-19. It incorporates a fusion of key features and architectural elements from various cutting-edge CNN models, leveraging their individual strengths to create a unified and powerful framework. By amalgamating these components, our model capitalizes on the diverse representations learned by different architectures, enhancing its ability to discern intricate patterns and nuances within lung images.

Furthermore, our model's architecture is crafted with a keen emphasis on efficiency and generalization. Through careful parameter optimization and regularization techniques, we mitigate the risk of overfitting, ensuring robust performance across diverse datasets. This emphasis on efficiency not only facilitates faster inference times but also enhances the model's scalability and adaptability to real-world clinical settings.

Importantly, Our proposed model, validated against well-established CNN architectures, demonstrates superior performance in diagnosing lung disorders, especially in COVID-19 cases. The enhanced diagnostic accuracy enables timely interventions and improved patient outcomes. Accurate identification of abnormalities facilitates prompt treatment initiation, reducing disease progression and complications. Early detection supports swift isolation, limiting transmission. Moreover, precise diagnosis optimizes resource allocation, ensuring efficient healthcare delivery. In summary, our hybrid model

offers a crucial tool for optimizing patient care pathways, particularly in the context of COVID-19, by streamlining diagnostics and enhancing treatment strategies.




Fig. 6. Hybrid CNN Architecture

In conclusion, our suggested hybrid CNN model offers unmatched precision, effectiveness, and adaptability, marking a substantial advancement in the field of lung disease diagnosis. Its innovative design and robust performance hold promise for revolutionizing clinical practice, ultimately improving patient care and outcomes in the global fight against lung diseases, particularly amidst the ongoing COVID-19 pandemic.

### 3.3 Model Training and Validation

During the Training phase of our project, we leverage preprocessed chest X-ray data. The Convolutional Neural Network (CNN), equipped with multiple convolutional and pooling layers, adeptly learns intricate patterns. The process involves forward propagation for predictions and backpropagation for model updates based on error rates, encompassing numerous training epochs. Each iteration is dedicated to minimizing the loss function, thereby enhancing the accuracy of pneumonia classification.

Concurrently, the Validation phase assumes a pivotal role in our methodology. Operating on a distinct dataset untouched during training, it rigorously evaluates the model’s generalizability. This iterative process fine-tunes the model’s hyperparameters without compromising the integrity of the evaluation, ensuring a reliable differentiation between normal and infected chest X-ray images.

### 3.4 Performance Metrics

The performance of our CNN model is rigorously assessed using two main quantitative metrics: the F1-Score, a balanced statistic that combines precision and recall, and accuracy, which is a direct reflection of the model’s overall correct classification rate across all predictions. Particularly useful in situations where class distributions are skewed is the F1-Score. When taken as a whole, these metrics provide a thorough evaluation of the model’s capacity to correctly classify chest X-ray pictures, guaranteeing that it is reliable and successful in differentiating between normal cases and those that are diseased.

## 4 Implementation

The deployment of our CNN model is a process that encompasses precise data management, utilization of resilient computational infrastructure, and meticulous fine-tuning of the model for optimal performance in the context of chest X-ray classification. This

section provides an in-depth overview of the model's practical implementation, covering the entire workflow from initial data collection and preprocessing to the setup of the computational environment and the strategic optimization of the model. The ultimate objective was to apply the model effectively to classify new chest X-ray images for COVID-19 detection.

#### 4.1 Dataset Collection and Handling

The dataset, crucial for training our CNN model, underwent thorough curation from varied chest X-ray sources. A total of 33,920 images were thoughtfully selected, including 10,701 'Normal' and 23,219 'covid and non-covid' cases. Manual selection aimed to capture a broad spectrum of chest X-ray variations, reflecting diversity in patient demographics, imaging techniques, and pathological conditions.

Maintaining dataset uniformity and quality was a priority. Each chest X-ray image underwent careful cropping and resizing, standardized to a resolution of 256x256 pixels. This resolution was chosen to balance preserving diagnostic details and optimizing computational efficiency, enabling the model to learn from clear and consistent image data.

For preprocessing, we employed the standardization techniques inherent in the VGG16 preprocessing library. This step normalized pixel values, aligning them with the input expectations of our custom hybrid CNN model.

To facilitate model training, we utilized Keras's `ImageDataGenerator`, dynamically augmenting our dataset through real-time data augmentation methods, including random rotations and shifts. This strategy diversified training data and enhanced the model's generalization capabilities. `ImageDataGenerator` efficiently fed images to the model in batches, streamlining the training process for improved effectiveness.

This meticulous approach to dataset collection and handling played a pivotal role in constructing a robust pneumonia classification model, ensuring it has a strong foundation of diverse and high-quality chest X-ray data.

#### 4.2 Computational Environment

The training and development of our models were conducted on a high-performance computing system equipped with an AMD Ryzen 5800X CPU. This processor is known for its fast execution of machine learning algorithms, which is crucial in handling the intensive computational demands of training deep neural networks. In addition to its robust CPU, the system has 64GB of RAM, which allows for effective data management and manipulation—even with big datasets and intricate model designs. Our setup runs on a Linux operating system, which is preferred for its stability and flexibility in research and development tasks. Because Linux is open-source, it may be heavily customized and optimized, which makes it a popular option for scientific computing requirements. At the heart of our computational environment is the GPU, an NVIDIA RTX 3090 with 24GB of VRAM. This GPU has 10,496 CUDA cores, enabling parallel processing that significantly speeds up the training process. The extensive VRAM allows larger batches and more complex models to be trained simultaneously, providing the necessary computational power to process our dataset efficiently. The synergy between the Ryzen

CPU, ample RAM, and the robust RTX 3090 GPU creates an optimized environment for deep learning tasks. This setup ensures that our models are trained with high efficiency and speed, allowing us to iterate quickly and improve our models with a fast turnaround.

### 4.3 Model Tuning and Optimization

Our deep learning model, designed to enhance pneumonia analysis, underwent meticulous tuning and optimization to reap ultimate performance. An important issue of this refinement procedure involved determining the excellent batch length and the range of training epochs. We opted for a batch length of 64, striking stability between computational efficiency and the community's overall performance abilities. The version underwent training throughout 50 epochs, permitting it to alter its weights and gradually minimize the loss characteristic.

A vital component of our optimization approach became the incorporation of batch normalization. This approach normalized the inputs of each layer, reducing internal covariate shifts. This normalization expedited the education method, imparting regularization benefits that were more desirable to the model's stability and efficiency.

Furthermore, we carried out a dynamic learning fee method, step by step reducing the learning price as training advanced via epochs. This adaptive approach facilitated finer adjustments to weights, stopping overfitting to the schooling facts because the model approached convergence.

The dataset becomes broken up into training and validation units, a vital step for complete model schooling at the same time as reserving a large portion of records for validation. Monitoring both accuracy and F1 rating in the course of schooling provided a holistic view of the version's performance. Accuracy served as an immediate measure, whilst the F1 rating provided deeper insights, balancing precision and bearing in mind, vital in eventualities with class imbalance.

In summary, our method of version tuning and optimization aimed for a harmonious balance between learning complexity and generalization ability. This ensured our version was successfully found from the training records without becoming overly tailor-made, thereby improving its applicability and reliability in real-world settings using deep learning.


## 5 Results and Analysis

The culmination of our methodical approach in developing the CNN model is reflected in this section, where we present the outcomes of our implementation and analyze its performance.


In Fig. 7, it is demonstrated that the model accurately recognized instances labeled as COVID-19 when they indeed belonged to this class. Nevertheless, it exhibited misclassification by incorrectly labeling non-COVID-19 instances as COVID-19.

As seen in Fig. 8, this model properly recognized the normal class but mispredicted non-COVID-19 cases as normal.


As seen in Fig. 9, this model properly recognized the normal class but mispredicted non-COVID-19 cases as normal.



**Fig. 7.** Results of VGG-19



**Fig. 8.** Results of ResNet-50




**Fig. 9.** Results of InceptionV3

It misclassified typical instances as COVID-19 while correctly identifying COVID-19 instances. Furthermore, the third image had non-COVID-19 cases that were incorrectly classified as COVID-19, as depicted in Fig. 10.

In comparison to the previously described pre-trained models, this model performed better overall and produced correct predictions for the normal, COVID-19, and non-COVID-19 classes. It is distinctly displayed in Fig. 11.



**Fig. 10.** Results of VGG-16



**Fig. 11.** Results of Custom model

## 5.1 Comparative Performance Analysis

Table 1 presents a performance comparison between our proposed model and other popular architectures, including VGG16, VGG19, ResNet50, and InceptionV3. The accuracy and F1 score—two key parameters that are evaluated during both the training and validation stages of the model’s complexity—were the subject of this comparison investigation.

### Observations

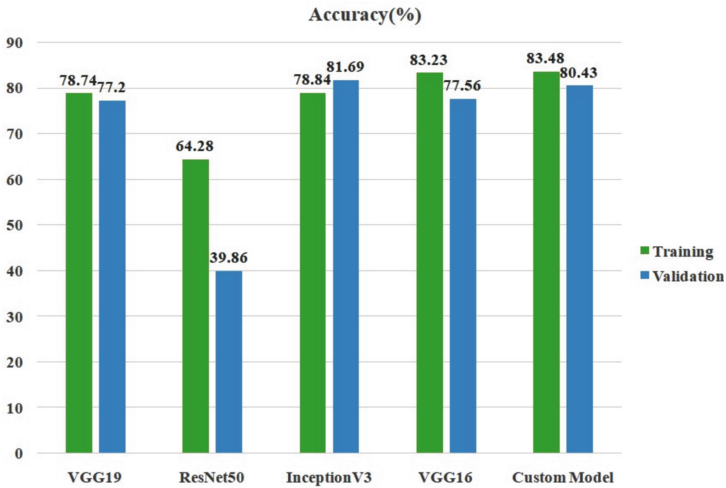
- VGG16 and VGG19 showed high training (T) accuracy and F1 scores but exhibited a significant drop in performance during validation (V), indicating a possible overfitting tendency. This means that the models were performing well on the data they were trained on, but they were not generalizing well to new data.
- ResNet50 also demonstrated a large difference (D) between training and validation accuracy, but not to the same extent as VGG16 and VGG19.
- Inception V3 showed a small difference between training and validation accuracy, but lower overall accuracy than VGG16 and VGG19.

**Table 1.** Comparative Analysis of Model Performance Metrics.

Model	Accuracy(%)			F1-Score(%)			IPI
	T	V	D	T	V	D	
VGG19	78.74	77.20	1.54	78.73	73.08	5.65	21073571
ResNet50	64.28	39.86	24.42	59.96	0.82	59.14	23593859
InceptionV3	78.84	81.69	-2.85	78.51	80.30	-1.79	24954659
VGG16	83.23	77.56	5.67	83.18	79.37	3.81	15763875
Custom Model	83.48	80.43	3.05	83.28	79.79	3.49	13140643

- The Custom model showed a smaller difference between training and validation accuracy, suggesting that it generalizes better to new data. It also had the highest overall accuracy.

Overall, the table suggests that the Custom model is the best performing model out of the five (Figs. 12, 13).



**Fig. 12.** Accuracy Graph




Fig. 13. F1-Score Graph

## 6 Conclusion and Future Work

### 6.1 Conclusion

In conclusion, our effort represents a noteworthy development in the field of lung disease diagnostics, particularly amid the challenges posed by the COVID-19 pandemic. Through the development and validation of our new hybrid Convolutional Neural Network (CNN) model, we have demonstrated its superiority over existing CNN designs, offering high efficiency and effective mitigation of overfitting issues. This innovation provides a more reliable and accurate means for the comprehensive evaluation of lung anomalies, including those related to COVID-19, thereby potentially improving patient care and outcomes. Moreover, our study's findings hold significant practical implications for clinical practice and healthcare decision-making. With our hybrid CNN model, we enhance diagnostic accuracy, crucial for promptly identifying lung abnormalities such as COVID-19. This advancement facilitates timely interventions, personalized treatments, and optimized resource allocation, ultimately enhancing patient outcomes and healthcare system efficiency. Healthcare professionals can leverage our model to streamline diagnostics, enabling earlier case detection, tailored interventions, and efficient resource allocation, thereby improving patient care pathways and decision-making.

### 6.2 Future Work

- **Multi-Modal Data Integration:** Extend the suggested hybrid CNN model to include a variety of data sources, including CT scans, X-rays of the chest, and patient clinical data. Integrating multi-modal data could further enhance the model's accuracy and robustness in diagnosing lung diseases.
- **Ethical Considerations and Bias Mitigation:** Address ethical concerns related to the use of deep learning models in healthcare, including patient privacy, consent, and

potential biases. Implement strategies to mitigate biases and ensure fair and equitable diagnostic outcomes across diverse populations.

- **Integration of Segmentation into Diagnostic Workflow:** Develop a seamless integration of lung segmentation into the diagnostic workflow. This involves combining the output of the segmentation model with the hybrid CNN model, creating a comprehensive system that not only identifies abnormalities but also precisely locates and assesses them within the lung region.

## References

1. Rajpurkar, P., et al.: CheXNet: radiologist-level pneumonia detection on chest X-rays with deep learning (2017). arXiv preprint [arXiv:1711.05225](https://arxiv.org/abs/1711.05225)
2. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: ChestX-ray8: hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2097–2106 (2018)
3. Zhang, B., Da, W., Zhang, X., Wang, K.: Responsible artificial intelligence in medicine: a framework for responsible research and innovation in medical artificial intelligence. *Artif. Intell. Med.* **103**, 101785 (2020)
4. McKinney, S.M., Kostopoulou, O., Sperrin, M.: Challenges and pitfalls of deploying artificial intelligence in clinical practice. *BMJ Health Care Inf.* **29**(1), e100436 (2022)
5. Liu, X., et al.: A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *Lancet Digit. Health* **1**(6), e271-e297 (2019)
6. Smith, A.C., et al.: Telehealth for global emergencies: implications for coronavirus disease 2019 (COVID-19). *J. Telemed. Telecare* (2020). <https://doi.org/10.1177/1357633X20916567>
7. Topol, E.J.: Deep medicine: how artificial intelligence can make healthcare human again. Basic Books (2019)
8. Liu, X., et al.: A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *Lancet Digit. Health* **3**(4), e271-e297 (2021)
9. Cabitza, F., Campagner, A., Balsano, C.: Ethical challenges of artificial intelligence-driven healthcare: a scoping review. *BMC Med. Inform. Decis. Mak.* **21**(1), 1–24 (2021)
10. Krittanawong, C., Zhang, H., Wang, Z., Aydar, M., Kitai, T.: Artificial intelligence in precision cardiovascular medicine. *J. Am. Coll. Cardiol.* **75**(22), 2764–2777 (2020)
11. Esteva, A., et al.: Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**(7639), 115–118 (2017)
12. Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L.H., Aerts, H.J.: Artificial intelligence in radiology. *Nat. Rev. Cancer* (2018)