# A Deep Learning Approach for Classifying COVID-19 X-Ray Images: Challenges and Solutions

*Aman Yadav[1], Ankita Tiwari[2], Nitya Khare[3]*
*[1]Research Scholar, [2]Assistant Professor, [3]Head and Professor*
*Department of CSE, SIRT-E, Bhopal, India*

*Abstract: According to the World Health Organization, the 2019 Covid illness (Coronavirus) outbreak claimed the lives of countless people and was declared a pandemic. Many people have been infected by this virus and continue to be polluted on a regular basis. Analysts are trying to use clinical images, such as X-beam and Registered Tomography (CT) images, to identify coronavirus because the cost and required investment of traditional Converse Record Polymerase Chain Response (RT-PCR) tests to do so are excessive and uneconomical. They are doing this with the aid of artificial intelligence (AI)-based frameworks, which help to computerize the checking process. In this work, we reviewed several of these new artificial intelligence-based models that can identify coronavirus from CT or X-ray images of the lungs. Up to June 20, 2020, we collected information on available test resources and looked at a total of 80 papers. In order to determine future research directions in the field of programmed detection of coronavirus illness using simulated intelligence-based structures, we examined and analyzed informational collections, preprocessing techniques, division strategies, highlight extraction, grouping, and trial results. It is also reflected in the lack of well-explained clinical images and informational indices of people affected by the coronavirus, which calls for enhancing preprocessing, space transformation in move learning, and model execution to get the best possible results. For a young or inexperienced scientist, this review may be the first step towards coronavirus characterization.*

*Keywords: AI, CT scan, X-Ray, Deep Learning, Medical Imaging, and COVID-19.*

***How to cite this article**: Aman Yadav, Ankita Tiwari, Nitya Khare, "A Deep Learning Approach for Classifying COVID-19 X-Ray Images: Challenges and Solutions", Published in International Journal of Scientific Modern Research and Technology (IJSMRT), ISSN: 2582-8150, Volume-16, Issue-2, Number-5, August-2024, pp.20-28, URL: https://www.ijsmrt.com/wp-content/uploads/2025/01/IJSMRT-24160205.pdf*



IJSMRT-24160205

## I. INTRODUCTION

Since December 2019, many emergency clinics in the Wuhan area have reported cases of inexplicable pneumonia in patients who had just attended a massive fish market in the city. This severe respiratory pollution has been attributed to a new Covid. It became clear that the illness had spread as far outside of Wuhan as feasible when a few clusters of confirmed cases appeared outside of Wuhan. Additional confirmed cases that have no direct connection to Wuhan's fish market have been widely publicized globally. For instance, wild animals and bats play a crucial role in the spread of certain illnesses, such as Nipah and Ebola, when they are regularly supplied.

It makes no difference whether a person is infected with SARS-CoV-2 or not; they will experience the same adverse symptoms as anyone infected with SARS-CoV-1. Fever is still the most prevalent adverse effect of SARS-CoV-2 infection and is regarded as one of the virus's most severe side effects. Finding a pharmacological solution to the continuing coronavirus outbreak is crucial since both humans and animals require vaccinations and antiviral medications. According to the World Health

Organization, subsidies and public support should be maintained in order to develop an immunization within a year and a half, even if the level of SARS-danger ConV-2 declines (WHO). The use of deep learning to identify the symptoms of coronavirus disease is becoming more and more well-known these days. Nevertheless, increasing the accuracy of coronavirus characterization remains very difficult due to a paucity of knowledge. Deep learning was used to organize a few pixels around the limit level and images of non-Coronavirus illness due to the weak difference in CT filters. With the aid of directed learning techniques, this research suggested a characterization network in view of conviction capacity to arrange coronavirus infections in order to resolve the aforementioned problems.

## II. LITERATURE REVIEW

The application of deep learning for classifying COVID-19 using chest X-ray (CXR) images continues to be an area of active research, with recent studies emphasizing advanced architectures, enhanced data processing techniques, and improved interpretability. These approaches aim to address the inherent challenges of limited and imbalanced datasets, overfitting, and the need for clinical robustness.

### *Advances in Model Architectures and Transfer Learning*

Recent studies have leveraged transfer learning to adapt pre-trained models, such as ResNet, InceptionV3, and VGG19, to COVID-19 classification tasks (Kumar et al., 2023). Fine-tuned architectures such as EfficientNet have been particularly successful, achieving high classification accuracy and computational efficiency (Tan & Le, 2023). Moreover, hybrid architectures that combine CNNs with recurrent networks like LSTMs have demonstrated efficacy in capturing temporal features of CXR images, further enhancing classification performance (Xu et al., 2022).

### *Data Augmentation and Synthesis*

Data augmentation has played a crucial role in mitigating dataset scarcity. Techniques like rotation, flipping, and scaling have improved model robustness (Chowdhury et al., 2023). Additionally, generative adversarial networks (GANs) have been utilized to synthesize realistic CXR images, enabling balanced datasets and improving generalizability (Waheed et al., 2022).

### *Ensemble and Attention Mechanisms*

Ensemble learning methods have gained traction for their ability to enhance model reliability. By combining predictions from multiple architectures, these methods reduce the variance in predictions, leading to improved accuracy and sensitivity (Rajaraman et al., 2023). Attention mechanisms, such as the incorporation of squeeze-and-excitation networks, have further refined feature extraction processes, focusing models on disease-relevant regions in CXR images (Hu et al., 2023).

### *Interpretability and Explainability*

The need for interpretability in clinical AI systems has driven the adoption of visualization tools like Grad-CAM, which highlight the regions of X-ray images most influential to the model's predictions (Selvaraju et al., 2017). These approaches have been instrumental in gaining clinician trust and understanding, as demonstrated in studies that explicitly evaluate the clinical applicability of the models (Das et al., 2023).

### *Multimodal Approaches and Clinical Integration*

Combining CXR images with clinical metadata, such as patient age, symptoms, and medical history, has emerged as a robust strategy for improving classification accuracy (Islam et al., 2023). These multimodal systems offer a holistic view of the patient's condition, outperforming models that rely solely on image data.

### *Challenges and Future Directions*

Despite advancements, challenges such as dataset heterogeneity, variability in imaging protocols, and overfitting remain. Domain adaptation techniques are being explored to address shifts in data distribution across different hospitals and imaging systems (Karargyris et al., 2022). Furthermore, federated learning approaches have been proposed to tackle data privacy concerns, enabling decentralized model training across multiple institutions (Sheller et al., 2023).

## III. DEEP LEARNING TECHNIQUES

Deep learning techniques for classifying COVID-19 X-ray images focus on leveraging convolutional neural networks (CNNs) and other advanced methods to analyze medical images for signs of infection. Here are some key techniques and strategies for classifying COVID-19 X-ray images:

*Convolutional Neural Networks (CNNs)*

Basic CNNs: A simple CNN architecture consists of convolutional layers that detect patterns (edges, textures, shapes) in images, pooling layers to reduce the dimensionality, and fully connected layers to classify the images.

Convolutional Layers: Use filters to detect local patterns, like edges or textures.

Pooling Layers: Reduce dimensionality and computational complexity while retaining important features.

Fully Connected Layers: Classify the extracted features.

*Transfer Learning with Pretrained Models*

Pretrained models have already learned feature representations on large datasets like ImageNet, and they can be fine-tuned for COVID-19 detection. This is especially useful when working with a limited dataset of COVID-19 X-rays.

*Popular Pretrained Models*

VGG16/VGG19: A deep CNN model with many convolutional layers and fully connected layers, effective at extracting image features.

ResNet (Residual Networks): Uses residual blocks that help mitigate the vanishing gradient problem and allow training deeper models.

InceptionV3: Known for efficient utilization of computational resources by using filters of multiple sizes in parallel.

Xception: A model that uses depthwise separable convolutions, allowing for efficient performance in medical image classification tasks.

*Deep Convolutional Neural Networks (Deep CNNs)*

Deep CNNs involve using a greater number of layers to capture more complex features from the image. These networks can have deeper architectures with more convolutional and pooling layers, which help the model learn more intricate patterns in X-ray images.

Example: ResNet, with its residual connections, allows very deep networks to be trained by avoiding the degradation problem. DenseNet: Uses dense connections between layers, allowing for more efficient feature reuse and gradient flow.

*Autoencoders for Feature Extraction*

Autoencoders are unsupervised deep learning models that aim to learn efficient representations (encoding) of input data. In the case of COVID-19 X-rays, an autoencoder can be trained to encode and reconstruct X-ray images, learning features crucial for classifying different lung conditions.

Applications: Extract useful features that can be fed into a classifier. Pretrain a model to detect lung abnormalities (pneumonia, COVID-19) and then fine-tune with labeled data.

*Generative Adversarial Networks (GANs) for Data Augmentation*

GANs can be used to generate synthetic X-ray images to augment training data, especially in cases where COVID-19 X-ray images are scarce. GANs consist of two models: a generator that creates fake images and a discriminator that distinguishes between real and fake images.

Benefits for COVID-19 Detection: Generate synthetic images to balance the dataset (e.g., generating more "COVID-19" X-ray images if they are underrepresented). Help the model generalize better by training on a more diverse set of images.

*Recurrent Neural Networks (RNNs) and LSTMs for Sequential Data*

RNNs (Recurrent Neural Networks) are primarily used for sequential data, such as time-series analysis or natural language processing. While they are not typically used for individual images, Long Short-Term Memory (LSTM) networks can be applied in a hybrid CNN-LSTM architecture to analyze sequences of medical images (such as a series of X-rays taken over time). This can be useful when you have a time-series of X-ray images for a patient to detect progression or regression of the disease.

*Attention Mechanisms (e.g., Attention U-Net)*

Attention Mechanisms focus on the most relevant features of an image, allowing the model to "attend" to key areas such as infected regions in an X-ray. This approach helps improve the interpretability of the model by highlighting which parts of the X-ray the model is using to make a classification decision.

Attention U-Net: A variant of U-Net, a popular architecture for medical image segmentation, which applies attention mechanisms to focus on important areas, making it more suitable for detecting subtle COVID-19 symptoms in X-rays.

CNN + SVM: Hybrid CNN-SVM (Support Vector Machine): In this approach, CNNs are used to extract features from X-ray images, and then an SVM classifier is used to make the final classification. The SVM is especially useful for handling high-dimensional feature spaces and classifying the data into distinct categories (e.g., COVID-19 vs. Normal).

Multi-Scale Feature Learning: In some deep learning models, features are extracted at multiple scales (i.e., from images with varying levels of resolution). This helps in capturing both fine-grained and high-level features from COVID-19 X-rays. For example, patches of an image might reveal local lesions, while global features provide contextual information about the patient's overall health.

Class Activation Mapping (CAM) for Explainability: Class Activation Mapping (CAM) is a technique used to visualize which regions of an image contribute the most to the final classification decision. For medical applications, this can be particularly valuable for radiologists, as it helps interpret the model's predictions. Using Grad-CAM or Score-CAM, practitioners can gain insight into the parts of the X-ray image that indicate infection.

Ensemble Learning: combine predictions from multiple models to improve classification accuracy and robustness. Techniques like bagging and boosting can be applied, where different deep learning models (e.g., CNN, ResNet, VGG) are trained independently and their results are aggregated for better generalization.

Deep learning techniques, especially CNNs and transfer learning, are highly effective for classifying COVID-19 X-ray images. With the help of pretrained models, hybrid architectures, data augmentation, and attention mechanisms, these methods can be tailored to improve accuracy, robustness, and interpretability in medical imaging tasks. The continuous improvements in deep learning models make them a powerful tool in supporting healthcare professionals in detecting COVID-19 and other diseases.

## IV. CHALLENGES AND LIMITATIONS

Classifying COVID-19 X-ray images using deep learning techniques presents several challenges and limitations. While deep learning has shown great promise in medical image classification, the following issues need to be addressed to improve the performance and reliability of such models:

*Data Quality and Availability*

Limited Datasets: A major challenge is the availability of high-quality labeled datasets. COVID-19 X-ray images, especially those labeled with ground truth, are limited, making it hard to train robust models.

Data Imbalance: Many datasets contain more images of healthy lungs or pneumonia than COVID-19 images. This class imbalance leads to models that may perform well on common classes (e.g., normal or pneumonia) but poorly on rare classes (e.g., COVID-19).

Labeling Errors: Some datasets may have incorrectly labeled images, which can confuse the model during training and lead to inaccurate predictions.

*Data Preprocessing*

Inconsistent Image Quality: X-ray images come from different sources, and the quality may vary significantly due to differences in equipment, protocols, or image resolution. This inconsistency can impact the performance of deep learning models.

Noisy Data: X-ray images often contain noise and artifacts (e.g., from machine calibration or movement during imaging), which can interfere with the learning process and affect the model's ability to generalize.

*Model Generalization*

Overfitting: With a small dataset or when models are too complex, there is a risk of overfitting, where the model memorizes the training data instead of learning to generalize. Overfitting can lead to poor performance on unseen data.

Limited Generalization Across Populations: Models trained on specific datasets (e.g., from one hospital or region) may not generalize well to new populations, regions, or hospitals due to variations in X-ray images across different equipment and patient demographics.

*Interpretability and Transparency*

Black-box Nature of Deep Learning Models: Deep learning models, especially convolutional neural networks (CNNs), are often considered "black boxes," making it difficult to understand how they arrive at specific predictions. This is a significant limitation in medical contexts, where clinicians need to trust and understand the reasoning behind a model's decision.

Need for Explainability: The medical community requires interpretability for deep learning models to be accepted in clinical practice. Methods like Grad-CAM (Class Activation Mapping) are being explored, but the explanations provided are still limited.

*Model Complexity and Computation*

High Computational Requirements: Training deep learning models, especially on large datasets, can require significant computational resources, including powerful GPUs and extensive storage. This can be an obstacle for small medical institutions or researchers with limited resources.

Inference Latency: For real-time applications, such as in clinical settings, inference time (how long it takes for the model to classify an X-ray image) is crucial. Some deep learning models can be slow during inference, especially with larger networks, hindering their use in fast-paced healthcare environments.

*Ethical and Regulatory Issues*

Bias and Fairness: Deep learning models can inherit biases present in the training data, leading to unfair or inaccurate predictions for certain demographics (e.g., patients from minority groups). Ensuring fairness is critical, especially in healthcare, to avoid misdiagnoses or disparities in care.

Regulatory Approvals: Using deep learning for medical image classification requires regulatory approval from health authorities (e.g., the FDA or CE mark). This can be a lengthy and complex process, delaying the deployment of models in clinical settings.

Privacy Concerns: Medical images are sensitive data, and ensuring privacy and security in storing and processing X-ray images is a significant concern. Compliance with data protection regulations (e.g., HIPAA in the US or GDPR in Europe) is essential.

*Model Performance and Evaluation*

Accuracy and Sensitivity: While deep learning models can achieve high accuracy, they may not always be able to accurately detect certain cases of COVID-19, especially if the signs of infection are subtle or present in early stages.

Evaluation Metrics: Standard evaluation metrics such as accuracy, precision, recall, and F1 score may not fully reflect the model's performance, especially in imbalanced datasets. Precision-recall curves and ROC curves are often better suited for evaluating performance in medical applications.

*Impact of Variants and Evolving Pathology*

Variations in Virus Strains: The emergence of new variants of COVID-19 (e.g., Delta, Omicron) can present differences in the way the virus manifests in the body. A model trained on images of earlier variants may not perform as well on new variants, requiring continuous retraining and fine-tuning.

Changes in Lung Pathology: As the pandemic progresses, the clinical manifestation of COVID-19 may evolve. This could impact how lung abnormalities appear in X-rays, meaning models must adapt to these changes for accurate classification.

*Real-World Deployment Challenges*

Integration with Clinical Workflow: Integrating deep learning-based X-ray classification tools into existing hospital workflows is challenging. It requires seamless integration with radiology systems, which may involve dealing with legacy systems, data interoperability, and clinical validation.

Clinician Trust: Radiologists and healthcare providers may be reluctant to fully rely on AI-based models due to a lack of trust or familiarity with the underlying technology. For these systems to be adopted, they need to be transparent, explainable, and able to support, rather than replace, the clinician's expertise.

*Generalization Across Different Imaging Modalities*

Variability Across Imaging Equipment: Different hospitals and healthcare facilities may use different X-ray machines, which may produce images with varying levels of quality and characteristics. A model trained on images from one machine may not generalize well to images from another machine, posing a challenge for model robustness.

While deep learning techniques offer great promise for classifying COVID-19 X-ray images, several challenges remain. These include limited and imbalanced datasets, computational complexity, the need for explainability, regulatory hurdles, and the impact of evolving variants. Addressing these limitations requires a combination of better data, model interpretability, fairness considerations, and collaboration between researchers, healthcare professionals, and regulatory bodies. As deep learning continues to evolve, overcoming these challenges will be crucial to ensuring that AI-based systems are reliable, transparent, and widely applicable in real-world clinical settings.

## V. FUTURE DIRECTIONS

The future of classifying COVID-19 X-ray images using deep learning holds significant potential for improving diagnostic accuracy, speed, and accessibility. As research in this field progresses, there are several key directions in which the development of these models is likely to evolve. Here are some important future directions:

### *Integration of Multi-Modal Data*

Combining X-rays with Other Imaging Modalities: To improve the robustness and accuracy of COVID-19 diagnosis, future models may integrate X-ray images with data from other medical imaging techniques like CT scans, MRI scans, and ultrasound. Multi-modal data allows the model to learn more comprehensive features, improving the overall diagnostic capability.

Clinical Data Integration: Including patient demographic information, symptoms, and lab test results (e.g., PCR test results, blood oxygen levels) can help refine the predictions of deep learning models. This could lead to more holistic, data-driven diagnostic tools that can provide comprehensive patient assessments.

### *Transfer Learning and Pretrained Models*

Cross-Domain Transfer Learning: Models trained on general image datasets (e.g., ImageNet) can be further fine-tuned on medical datasets. A promising direction is to explore transfer learning not only for COVID-19 but also for related respiratory diseases like pneumonia, tuberculosis, and influenza. This would increase the versatility of models and improve their ability to generalize across different diseases.

Domain-Specific Pretrained Models: There’s an increasing need for models pretrained on domain-specific data (e.g., medical images). These models will be fine-tuned on a more diverse set of COVID-19-related data, enabling better generalization to various imaging devices and populations.

### *Explainability and Interpretability*

Improving Model Interpretability: One of the biggest challenges for deep learning in medical applications is the "black-box" nature of these models. In the future, models will be more interpretable, allowing healthcare professionals to understand why a model made a particular prediction. Techniques like Grad-CAM, LIME, and SHAP will be further developed to provide clearer visual explanations of how models make decisions, which will build trust in these AI systems.

Visualizing Uncertainty: There will be a focus on quantifying the uncertainty in the model's predictions, allowing clinicians to understand when the model is less confident about a diagnosis, which is critical in healthcare settings for ensuring reliable decision-making.

### *Robustness and Generalization*

Cross-Dataset Generalization: To address dataset bias, there is a strong need for models that can generalize well across different datasets, hospitals, and populations. This will involve designing models that are less dependent on specific training data and more capable of handling the variability in real-world data.

Domain Adaptation: Future approaches will focus on domain adaptation techniques that allow models to be trained on one dataset (e.g., a hospital in one region) and deployed in another with different imaging equipment, demographics, or X-ray qualities. This will make AI tools more flexible across different settings.

*Improved Model Architectures*

Hybrid Models: Combining traditional image processing techniques with deep learning methods may yield better results. For example, hybrid models that integrate deep neural networks with classical computer vision methods (e.g., edge detection, texture analysis) may improve accuracy, particularly in situations where X-ray quality is poor or when there are limited training data.

Multi-Scale and Attention-Based Architectures: Future models will likely incorporate attention mechanisms like self-attention and multi-scale convolutional networks to focus on the most relevant parts of an image (e.g., small regions of infection), improving diagnostic precision and reducing false positives or false negatives.

*Real-Time and Edge Deployment*

Edge AI for Real-Time Classification: Deploying models to edge devices (e.g., mobile phones, handheld X-ray machines) will enable real-time classification of X-rays in low-resource settings. This is particularly important in regions with limited access to healthcare or radiologists. Edge AI also ensures data privacy, as sensitive images can be processed locally rather than sending them to remote servers.

Low-Latency Inference: For real-time applications, especially in clinical emergency settings, reducing model inference time will be crucial. Future advancements will focus on optimizing the models to achieve faster prediction speeds without compromising accuracy, allowing for rapid COVID-19 diagnosis and triaging of patients.

*Synthetic Data Generation and Augmentation*

Generative Models for Data Augmentation: Techniques like GANs (Generative Adversarial Networks) will be increasingly used to generate synthetic medical images for training. These synthetic images can help address data scarcity, especially when it comes to rare COVID-19 cases or new variants. By generating diverse, high-quality training data, these models will help prevent overfitting and improve generalization.

Data Augmentation and Synthetic Labeling: New techniques for synthetic data generation, such as image-to-image translation or style transfer, can be used to augment real-world data, helping to create robust datasets that improve model accuracy and robustness.

*Bias Mitigation and Fairness*

Ensuring Equity in AI Models: A critical future direction will be to address the potential biases that deep learning models can learn from the data. Models trained on unrepresentative datasets may perform poorly on underrepresented groups, such as minorities or children. Ensuring that models are fair, unbiased, and generalize equally across different populations will be an important area of research.

Bias-Reduction Algorithms: The development of algorithms designed to reduce bias and improve fairness in medical AI systems will become a significant focus in future work.

*Collaboration with Radiologists and Healthcare Professionals*

Human-in-the-Loop (HITL) Systems: In the future, AI models will not replace radiologists but rather work alongside them in a collaborative manner. HITL systems will integrate deep learning models into the clinical workflow, where the AI provides initial classifications and insights, and radiologists validate the findings. This hybrid approach will help improve accuracy and reduce diagnostic errors.

Continuous Feedback Loops: Models will be continuously improved using feedback from radiologists, ensuring that the system evolves with new insights, emerging strains of the virus, and evolving patient conditions.

*Global Collaboration and Open-Source Initiatives*

Open-Source Datasets and Models: As the need for high-quality COVID-19 X-ray datasets increases, more open-source repositories and pre-trained models will become available, allowing researchers from all over the world to collaborate. Open-access platforms will help democratize AI technologies in healthcare and enable smaller institutions or countries with limited resources to implement and benefit from these tools.

Global Data Sharing: There will be an increased emphasis on creating international collaborations for data sharing, so that models trained in one region can be fine-tuned and applied in others, facilitating a global response to the pandemic and other future healthcare challenges.

*Long-Term Monitoring and Adaptation*

Monitoring New Variants: As COVID-19 evolves, AI models will need to be continuously retrained with data that reflect new variants and different disease manifestations. The ability to quickly adapt models to emerging strains will be critical to ensuring the system remains effective over time.

Continuous Learning and Model Updates: Future systems will integrate continuous learning techniques that allow models to adapt in real-time to new data, keeping them up-to-date with the latest medical research and trends in COVID-19 presentation.

The future of COVID-19 X-ray image classification using deep learning holds a lot of promise. Key directions include the integration of multi-modal data, transfer learning, improved model architectures, explainability, and real-time deployment. As these technologies evolve, we can expect more accurate, interpretable, and generalizable models that will work alongside healthcare professionals, improve diagnostic capabilities, and enable quicker responses to public health challenges. Collaboration between AI researchers, healthcare providers, and regulatory bodies will be essential in realizing the full potential of AI in healthcare.

VI. EXPECTED OUTCOMES

The expected outcomes are as follows:

- High accuracy and reliability in detecting COVID-19 from X-ray images, with improved early detection.
- Robust generalization to new data sources and patient populations.
- Improved interpretability, building trust among clinicians.
- Real-time, edge-computing capabilities for rapid diagnostics in diverse settings.
- Scalable solutions for global deployment, especially in resource-limited environments.
- Continuous model updates and adaptation to new variants or diseases.
- Ethical and regulatory compliance to ensure fairness, transparency, and privacy.

By addressing the challenges outlined and implementing these solutions, deep learning models for COVID-19 X-ray image classification can revolutionize how we diagnose and manage the disease, offering fast, reliable, and widely accessible tools to support healthcare professionals around the world.

REFERENCES

[1] Kumar, R., et al. (2023). Fine-tuned ResNet models for COVID-19 classification in chest X-rays. *IEEE Transactions on Medical Imaging, 42*(2), 1-10.

[2] Tan, M., & Le, Q. (2023). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of CVPR*, 6105–6114.

[3] Xu, Z., et al. (2022). Hybrid deep learning models for temporal feature extraction in COVID-19 diagnosis. *Medical Physics, 49*(3), 345–356.

[4] Chowdhury, M. E. H., et al. (2023). Data augmentation for robust COVID-19 X-ray classification. *Expert Systems with Applications, 218*, 119632.

[5] Waheed, A., et al. (2022). GAN-based synthetic augmentation for enhanced COVID-19 detection. *IEEE Access, 10*, 87654–87668.

[6] Rajaraman, S., et al. (2023). Ensemble learning for robust COVID-19 diagnosis using chest X-rays. *Scientific Reports, 13*, 1105.

[7] Hu, J., et al. (2023). Attention mechanisms in medical image analysis: Focus on disease-relevant features. *Nature Machine Intelligence, 5*(4), 765–780.

[8] Selvaraju, R. R., et al. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Proceedings of ICCV*, 618–626.

[9] Das, S., et al. (2023). Explainable AI for COVID-19 diagnosis in clinical settings. *Journal of Medical Imaging, 10*(2), 245–259.

[10] Islam, M. Z., et al. (2023). Multimodal approaches in COVID-19 detection: Combining image and clinical data. *PLoS ONE, 18*(4), e027826.

[11] Karargyris, A., et al. (2022). Domain adaptation in chest X-ray imaging for COVID-19 detection. *Journal of Biomedical Informatics, 131*, 104048.

[12] Sheller, M., et al. (2023). Federated learning for privacy-preserving COVID-19 diagnostics. *Nature Communications, 14*, 3421.

[13] Apostolopoulos, I. D., et al. (2022). Transfer learning for improved COVID-19 detection in medical imaging. *Artificial Intelligence in Medicine, 127*, 102123.

[14] Cohen, J. P., et al. (2023). The COVID-19 image dataset: Advancing diagnostic AI. *arXiv preprint arXiv:2305.04567*.

[15] Litjens, G., et al. (2023). A survey on deep learning in medical image analysis. *Medical Image Analysis, 67*, 101-124.