

Pneumonia Detection Enhanced by Conditional Generative Adversarial Networks cGAN Addressing Class Imbalance with High Quality Synthetic Data

Vijayalakshmi S.¹, Vijayasekaran G.², Vinoth Kumar P.³, Ayub Khan A.⁴

¹Associate Professor, ⁴Assistant Professor, Department of Electronics and Communication Engineering, Sona College of Technology, Salem, India.

²Associate Professor, Computer Science and Engineering, New Horizon College of Engineering, Bengaluru, Karnataka, India.

³Assistant Professor, Department of Electronics and Communication Engineering, Nandha College of Technology, Erode, India.

E-mail: ¹vijisaumiya@gmail.com, ²gunavijay90@gmail.com, ³vinothkumarp.research@gmail.com, ⁴ayubkhan.slm@gmail.com

Abstract

Pneumonia has recently reported the highest number of deaths in the entire world. Consequently, diagnostic procedures ought to be precise. Conversely, the problem of class imbalance is a universal issue in medical image classification that may result in biased models that do not perform well in underrepresented classes. The issue of class disproportion is successfully addressed in the work under consideration, as a cGAN generates high-quality synthetic images of the minor classes, allowing for the creation of a balanced dataset that, in turn, leads to high sensitivity of the model and consequently improves its overall performance. This paper presents a new Conditional Generative Adversarial Network (cGAN) architecture to improve the detection of pneumonia in the provided chest X-rays. To mitigate the issue of data imbalance in the dataset, this work suggests a conditional GAN-based augmentation process for synthetic X-ray images by producing clinically viable and label-coherent synthetic X-ray images. The framework also includes validation achieved through the application of either SSIM/FID or balanced training, which improves the accuracy of pneumonia detection and leads to diagnostic conclusions unlike those of existing methods. Current approaches are less successful compared to this study, as the system achieves 96.5% accuracy, 95.2% sensitivity, and 96.1% specificity, with a high F1 score. The suggested framework has a good performance scale, making it applicable in medical applications. This paper demonstrates the capacity of cGAN to develop pneumonia diagnosis machines that are feasible and user-friendly in healthcare institutions.

Keywords: Pneumonia Detection; Chest X-ray Imaging; Conditional Generative Adversarial Networks (cGAN); Synthetic Image Validation; Attention Mechanisms.

1. Introduction

Pneumonia is a health issue that has alarming consequences across the globe and is a cause of death among a large number of individuals, particularly in developing countries. The World Health Organization reports that each year, 800,000 children below the age of five die of pneumonia, and millions of children worldwide are experiencing its health impacts. The diagnosis should be made early and accurately to reduce these numbers; however, there are issues with the current diagnostic tools that need to be addressed. The healthcare system needs precise diagnostic equipment that does not require human specialists and is able to handle all patients irrespective of their location. Figure 1 demonstrates recent trends among reported cases of pneumonia and deaths from this disease during that period.




Figure 1. Pneumonia Cases Reported vs. Deaths in Recent Decade

The recent artificial intelligence systems that are guided by deep learning technologies can do much to change medical imaging. CNNs are good at analyzing chest X-ray images to find pneumonia; however, they also perform when dealing with substantial medical scans. Models have several issues in the process of pneumonia identification, such as imbalanced data resulting in decreased accuracy of classification and scarce training data that leads to overfitting, which cannot handle different factors of patient populations. The solution to these problems is the development of strong medical image applications. Recent developments in artificial intelligence, particularly deep learning, offer great potential for medical imaging. Optimal findings can be achieved through CNN technology analysis of chest X-rays using deep learning to detect pneumonia. Various issues may arise in deep learning models pertaining to patient health care data, such as unequal training data, which essentially demands a more diverse group of patients. This necessity for addressing this issue would lead to the establishment of reliable health solutions that would indeed be more helpful to patients.

1.1 Research Gap

Although current studies have made significant advancements in the field of deep learning-based pneumonia detection, there are several critical limitations that form a significant gap which the present research will fill. Most past research uses disproportional samples of chest X-rays, where normal cases are exceedingly higher than those of pneumonia. This

imbalance leads to biased training and models with less sensitivity, showing uneven performance when taken to novel clinical environments. Unlike several studies, which have tried using GANs for data augmentation, the quality of the resultant diagnostic images is rarely checked with the aid of powerful and clinically significant diagnostic indicators such as SSIM or FID. Synthetic samples in past research are therefore usually structurally unfaithful and thus of less value in training medical diagnostic systems. Another gap identified in recent literature is that most of these methods test their models on just one dataset available in the public domain, raising the question of whether their models can be generalized. Deep learning systems also do not train their mechanisms for diagnostic consistency when trained on minority-class synthetic samples. Additionally, more structured models like attention networks and CNN-LSTM models require significant computing power and cannot be accommodated in systems with low computing capability or constrained healthcare infrastructures. Furthermore, most existing GAN-based augmentation models do not condition labels; thus, the synthetic images generated do not contain the correct disease-specific radiological features that could allow for successful classification of pneumonia. In order to overcome such constraints, we must address the challenge of synthesizing clinically significant synthetic images, verifying such images through objective quality metrics, balancing the minority class samples without creating artifacts, and ensuring that the entire framework is computationally efficient so that it can apply to real-world cases. The proposed cGAN-based framework directly addresses these gaps in this study regarding high-fidelity and condition-specific synthetic image generation, authenticated by both SSIM and FID, dataset balancing, and enhancing classification accuracy, which is also practical in clinical practice.

The key objectives of the research are:

1. Present a framework of cGAN to generate synthetic high-quality chest X-ray images with specific conditions.
2. Assess the diagnostic value of synthetic images through the use of sound quantitative metrics, such as SSIM and FID.
3. 3.Improve the accuracy of pneumonia detection by including synthetic images in the training dataset.
4. Evaluate the proposed framework in terms of scalability and computational efficiency for practical applications.

The proposed cGAN system comprises Generator and Discriminator components optimized for medical imaging tasks. While the Generator creates fake images from input class labels, the Discriminator checks if those images are suitable for medical use. This model uses real and synthetic images from the enhanced dataset to refine its classification performance. The proposed framework demonstrates its performance through numerical evaluation and in comparison, with existing leading methods.

This paper proposes a new conditional GAN-assisted augmentation model, specifically designed to detect pneumonia in children using chest X-ray images. The approach differs from current research as it generates class-conditioned synthetic radiographs that do not alter the patterns of disease-relevant structures, and includes a quantitative validation step to remove low-quality samples according to SSIM and FID. In contrast to previous literature, which utilizes classical oversampling or unverified GAN images only, the proposed framework

generates a balanced dataset composed of clinically sound synthetic and real images, enabling a downstream classifier to learn more about pneumonia features. This contribution opens a new trend where adversarial synthesis and structural similarity validation are combined to mitigate class imbalance and enhance diagnostic robustness.

In the rest of the paper, we describe how we build and validate the cGAN model, report the results of our comparative analysis, and discuss the relevance of these results for the field of AI in medical diagnostics in general.

2. Related Work

Deep learning technology has become the leading medical data analysis solution for diagnosing illnesses through chest X-rays. Convolutional Neural Networks are the top choice in machine learning because they effectively read spatial and organizational patterns in image data. These deep learning architectures show exceptional results in classifying pneumonia through their complex design. These models need substantial amounts of balanced data to function at their best, according to research [11][12].

Medical professionals find Generative Adversarial Networks (GANs) highly effective for creating enhancements in medical datasets. GANs consist of two parts—a Generator and a Discriminator—which create synthetic images that align with the statistics of real datasets. Many research findings prove GAN technology's capacity to reduce class imbalance through the generation of synthetic minority class samples. One study used a GAN system to enhance breast cancer diagnosis by creating realistic synthetic mammograms for the small minority class [15]. By adding condition vectors, cGANs outperform basic GANs in their ability to create specific types of data output. According to [16], cGANs created realistic retinal images based on disease severity, which helped train robust models. Most GAN-based methods fail to demonstrate the actual quality of their synthetic data, making them unsuitable for healthcare use [17]. Training these models becomes unstable and requires specific adjustments of control parameters. Deep learning systems benefit from attention mechanisms that help them learn valuable medical image features by looking at specific areas of interest in the images. In medical imaging applications, Transformer networks demonstrate their utility through self-attention, which proves effective in segmenting and classifying brain MRI measurements [18]. Researchers have built hybrid CNN designs with attention systems, including CNN-LSTMs, to process spatial and temporal medical image data. Research demonstrates that combining different neural network types helps better identify pneumonia along with other respiratory conditions [19]. Models that use attention require significant computing power and many examples to prevent the problem of overfitting to specific data. Their fixed attention systems tend to miss important information in datasets that contain fewer common examples [20], but recent techniques use transfer learning to address this issue. Transfer learning approaches achieve good results but struggle to perform well across different imaging modalities and clinical settings because of domain-specific challenges [13][14].

Recent research shows that cGANs can create diverse, realistic images that help solve class imbalance problems [23]. Clinical use of these models remains limited due to a lack of trustworthy validation standards, including SSIM and FID [24]. Researchers have tested various state-of-the-art methods to find new ways of identifying pneumonia. A 2022 study suggested using attention-based CNNs to analyze the NIH Chest X-ray dataset and achieved high detection accuracy [25]. A hybrid CNN-LSTM system outperformed basic CNN models

by recognizing correlations between chest X-ray image areas and their sequential development [26].

Despite new techniques being adopted for pneumonia detection, further research is needed to ensure the system works well for every patient and X-ray environment. Most research models work with public databases, but these datasets often lack adequate patient diversity [27]. Current research shows limited use of synthetic data to improve pneumonia detection results in studies [28].

Table 1. Comparative Analysis of Existing Recent Methods

Technique	Merits	Limitations	References
Semi-Supervised Learning	Reduces dependency on labeled data, making it suitable for limited datasets.	Performance depends heavily on the quality of the small, labeled subset.	[29]
Transfer Learning with Fine-Tuning	Achieves high performance with small datasets by leveraging pre-trained networks.	Domain shifts may cause degradation in accuracy and generalizability.	[30]
DenseNet Architectures	Provides high feature extraction efficiency, especially in medical imaging tasks.	Computationally expensive with memory limitations for large datasets.	[22]
Hybrid Transformer-CNN Models	Combines CNN's feature extraction with Transformer's global context learning, boosting accuracy.	Requires large computational resources and training time.	[23]
Augmented Dataset Validation with FID and SSIM	Validates synthetic data quality for clinical use, ensuring diagnostic relevance.	Limited adoption in real-world scenarios; validation often lacks expert evaluation.	[24]
LSTM-Based Time-Series Analysis in Imaging	Captures temporal dependencies in medical image sequences, improving detection rates.	Limited applicability for static chest X-rays; prone to overfitting.	[26]
DCNNs	High accuracy; effective in identifying complex patterns.	Requires extensive data; issues with interpretability.	[31], [32]
Attention Mechanisms	Enhances interpretability; focuses analysis on key areas.	Computationally intensive; integration complexity.	[33]

GANs	Useful in data augmentation; enhances model robustness.	Training instability; potential for generating artifacts.	[34]
Transfer Learning	Reduces training resources; leverages pre-trained networks.	Domain mismatch can affect performance.	[35]
Ensemble Methods	Improves overall model accuracy and reliability.	Increases computational cost; complex management needed.	[36], [38]

From the reviewed literature, the following gaps are identified:

1. While generative models such as GANs and cGANs are widely used, the quality and diagnostic relevance of synthetic images are rarely validated using robust metrics like SSIM or FID.
2. Existing methods for handling class imbalance often fail to generate diverse and clinically relevant samples, limiting their impact on diagnostic accuracy.
3. Most studies focus on single datasets, raising concerns about the generalizability of models across diverse clinical settings.
4. Advanced architectures, such as attention mechanisms and hybrid models, often require significant computational resources, making them impractical for resource-constrained environments.

The proposed research develops a new cGAN model to bridge these knowledge gaps and diagnose pneumonia from chest X-ray images. Our model generates synthetic training images that solve the data shortage problem and make high-quality model training possible in limited data conditions. The approach uses enhanced validation tests to check if the synthetic generated images look realistic enough and match the requirements of medical diagnosis.

3. Proposed Work

The proposed cGAN model consists of two main components: the generator and discriminator. The generator develops X-ray images that look like actual pneumonia scans to enlarge the training dataset available. Our additional dataset helps solve the typical lack of medical examples that appear rarely. The discriminator tests image authenticity and checks whether synthetic images work well enough in medical diagnoses. When the generator creates new images, it takes instructions from the training needs to make underrepresented class expert images. The technique produces synthetic images that exactly meet the gaps in our dataset. The system uses attention mechanisms to detect and highlight the most important image areas needed by doctors to make their diagnostic decisions. The system shows doctors exactly where and how AI makes decisions, allowing them to trust the AI medical diagnosis platform more easily. The proposed research design meets medical imaging AI needs while creating better ways to apply AI tools directly in clinical practice for better pneumonia diagnosis. This work will help medical teams spot pneumonia sooner with better accuracy, which directly benefits patients in disadvantaged populations.

Figure 2 illustrates the architecture of the proposed Conditional Generative Adversarial Network (cGAN) framework, highlighting its two main components: the Generator and the Discriminator. The Generator generates synthetic chest X-rays based on selected class identifiers or clinical points. This model takes both random noise and condition inputs, then merges them into a defined dimensional space before processing. The dense block transforms the representation into a deconvolution process, which enlarges the feature map into an image space. Batch normalization is added to smooth training, and the output layer uses Sigmoid activation to convert the generated image into values between 0 and 1. By examining synthetic images, the Discriminator tests the output made by the Generator. An input that includes both real and synthetic images plus condition data goes into the system. The Discriminator uses convolutional layers to extract image features, which it combines with the condition vector before processing through fully connected layers to produce a real or synthetic image classification score. The Generator and Discriminator compete against each other through training updates, where the Generator reads input to generate convincing outputs while the Discriminator determines their reality level.




Figure 2. Flow Process of the Proposed cGAN Model

Batch normalization plays a crucial role in stabilizing the adversarial training process of the proposed cGAN. By normalizing intermediate feature distributions within the Generator, batch normalization reduces internal covariate shift, improves gradient flow, and prevents activation saturation. This leads to faster convergence and enables the Generator to model fine-grained radiological structures without mode collapse. Within the Discriminator, batch normalization ensures consistent feature scaling when evaluating real and synthetic X-ray samples, which is essential in adversarial learning, where fluctuations in intermediate

representations can destabilize training. Overall, batch normalization enhances training stability, improves synthetic image fidelity, and contributes directly to the reliability of the balanced dataset used for pneumonia classification.

Our framework starts by changing chest X-ray images to match uniform dimension requirements and normalizing pixel values for the dataset. The cGAN takes pre-processed images together with their medical condition values during learning. The Generator learns to generate specific condition images while the Discriminator analyzes image realism and condition alignment during this training period. During training the competing algorithms adjust performance based on the set loss parameters. The synthetic images require confirmation of their medical accuracy; hence their quality is measured using the Structural Similarity Index (SSIM) and Fréchet Inception Distance (FID) tests. After validating the synthetic images, they are blended with the original data to create a new dataset that handles class unfairness better. The model training uses the extended dataset to identify pneumonia cases while reporting performance through multiple evaluation measures. Tests determine the model's performance speed to validate its use in real healthcare applications. Our approach improves medical outcome testing without increasing computer workload to support medical use.

A cGAN consists of two main components: This design incorporates two networks - a generator G and a discriminator D to handle additional information y . When paired with conditioners the system trains models to produce more accurate images that reflect target feature sets for improved data balance. The generator G and discriminator D in a cGAN play a min-max game defined by the following objective function:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log(D(x|y))] + \mathbb{E}_{z \sim p_{data}(z)} [\log(1 - G(z|y))] \quad (1)$$

In Eq. (1), x is a real image from the dataset, z is a noise vector, sampled from a noise $p_z(z)$, y is the condition or label associated with the image, $G(z|y)$ is the output of the generator given noise z and condition y , producing synthetic images and $D(x|y)$ is the discriminator's estimate of the probability that x is a real image given the condition y , as opposed to a synthetic image produced by G .

The min-max goal in Equation (1) is a mathematical expression of the adversarial character of the cGAN with the Generator having to minimize the capability of the Discriminator to identify real versus synthetic samples, and the Discriminator having to maximize the difference. The saddle-point optimization is what pulls the Generator toward the generation of synthetic images with a conditional distribution identical to that of natural pneumonia images. Consequently, the balance between the min and max game provides better sample fidelity, less discrepancy between real and generated distributions, and convergence of the adversarial training procedure.

The generator is tasked with creating images that the discriminator will classify as real. It is a function $G(z|y)$ that maps the input noise vector z , conditioned on y , to the data space. The generator's architecture can typically be a deep convolutional neural network. The conditioning on y is usually implemented by concatenating y with z at the input layer, allowing the generator to produce data specific to the given condition.

The discriminator acts as a binary classifier that tries to distinguish between real data samples and fake data samples produced by the generator. The function $D(x|y)$ outputs a scalar representing the probability that x is real. Like the generator, the discriminator is also

conditioned on y , which can be incorporated into the discriminator architecture by concatenating y with the input x .

To train the cGAN, we use the binary cross-entropy loss for both the generator and the discriminator. For the discriminator, the loss is defined as:

$$\mathcal{L}_D = -\frac{1}{m} \sum_{i=1}^m [y_i \log D(x_i|y_i) + (1 - y_i) \log(1 - D(x_i|y_i))] \quad (2)$$

In Eq. (2), m represents batch size, y_i is the label indicating whether x_i is real or synthetic.

For the generator, the loss is:

$$\mathcal{L}_G = -\frac{1}{m} \sum_{i=1}^m \log D(G(z_i|y_i)) \quad (3)$$

The proposed cGAN model successfully creates realistic medical images for AI medical diagnostics training. By reducing class imbalance problems in the training data, the model helps advance AI healthcare systems that detect pneumonia in X-ray images. The design uses advanced mathematical rules to display clear operations that scientists can examine and modify for many medical imaging fields.

The cGAN model takes random noise combined with conditional data to create synthetic medical images that look just like real ones. The first stage involves conditioning the input noise vector and the additional information. Within the suggested cGAN architecture, the attention mechanism is used following the early convolutional stages of the generator. The initial convolutional blocks identify low-level radiographic features after which the attention module is deployed to enhance and emphasize pneumonia-related areas and the image is then gradually upsampled into the final synthetic chest X-ray. Such positioning will make the Generator focus on synthesis of clinically meaningful structures.

Let z be a latent space noise vector sampled from a standard normal distribution $p_z(z)$, and let y be a label or any other auxiliary information related to the specific characteristics of the image to be generated (e.g., "pneumonia" vs. "normal"). The conditioning is performed by concatenating z and y , creating an augmented input z_y :

$$z_y = \text{concat}(z, y) \quad (4)$$

The augmented input z_y is then fed into the generator G , which is a deep convolutional neural network. The generator function G maps z_y to the data space to produce a synthetic image $\hat{x}$

$$\hat{x} = G(z_y) \quad (5)$$

Through stacked transposed convolutions, batch normalization and ReLU blocking layers the generator scales conditioned inputs into detailed X-ray image outputs that match actual radiological results. The network D receives real images x from the dataset combined with synthetic input $\hat{x}$ along with condition values y . The discriminator D analyses both real X-ray scans and generated X-ray results based on y records. The discriminator network D accepts the complete x and synthetic $\hat{x}$ images, conditioned on y combination when handling real images yet receives the joint input of $\hat{x}, y$ for synthetic images. The discriminator outputs a probability indicating the likelihood that the given image is real:

$$D(x|y) = \sigma(f(\text{concat}(x, y))) \quad (6)$$

$$D(\hat{x}|y) = \sigma(f(\text{concat}(\hat{x}, y))) \quad (7)$$

In Eq. (6), σ is the sigmoid activation function, and f represents the internal mappings of the discriminator which include convolutional layers, batch normalization, and LeakyReLU activations. Each neural network model in a GAN follows separate loss functions to modify its internal parameters through the training data flow. The generator aims to maximize the probability that its generated images are classified as real by the discriminator:

$$\mathcal{L}_G = -\log D(\hat{x}|y) \quad (8)$$

This loss encourages the generator to produce images that are increasingly difficult for the discriminator to differentiate from real images.

On the other hand, the discriminator aims to correctly classify both real and synthetic images, and its loss is given by:

$$\mathcal{L}_D = -[\log D(x|y) + \log(1 - D(\hat{x}|y))] \quad (9)$$

This loss penalizes the discriminator when it incorrectly classifies real images as fake or fake images as real.

Strategic weight updates by the generator and discriminator occur after computing losses through optimization routines including SGD. The model constantly processes input and output until it meets criteria at which point both systems function equally well. By following these learning stages, the cGAN system produces convincing synthetic chest X-ray images that match known clinical traits including pneumonia. Our advanced model consumes input information and transforms it into useful output through mathematical foundations that protect its performance when working with unbalanced medical images and limited AI training data. Our proposed cGAN model operates under the following production logic.

Algorithm cGAN Model

- 1: *Initialize*
- 2: $G(z|y; \theta_g)$ as generator with parameters θ_g
- 3: $D(x, y; \theta_d)$ as discriminator with parameters θ_d
- 4: λ -regularization coefficient.
- Define generator layers:**
- 5: *Dense layer with Leaky ReLU activation*
- 6: *Batch Normalization*
- 7: *Transpose Convolution layers with ReLU activation*
- 8: *Final Convolution layer with tanh activation for image output*
- Define discriminator layers:**
- 9: *Convolution layers with leaky ReLU layers*
- 10: *Batch Normalization*
- 11: *Dropout layers for regularization*
- 12: *Final Dense Layer with Sigmoid Activation Function*
- 13: *for each training epoch:*
- 14: *for each batch $\{z_i, y_i, x_i\}$*
- 15: *sample z_i from $p_z(z)$, $i = 1$ to m*
- 16: *sample $\{x_i, y_i\}$ from $p_{data}(x, y)$, $i = 1$ to m*

Generate synthetic data:

17: $z_{y_i} = \text{concat}(z_i, y_i)$ for each i

18: $\hat{x} = G(z_{y_i}; \theta_g)$ using defined generator layers

Compute discriminator output:

19: For real data: $D_{real} = D(\text{concat}(x_i, y_i); \theta_d)$

20: For fake data: $D_{fake} = D(\text{concat}(x_i, y_i); \theta_d)$

Calculate discriminator loss:

21: $L_{D_{real}} = -\log(D_{real})$

22: $L_{D_{fake}} = -\log(1 - D_{fake})$

23: $L_D = (L_{D_{real}} + L_{D_{fake}})/2$

24: Add L2 regularization:

25: $L_{D_{total}} = L_D + \lambda ||\theta_d||^2$

update discriminator parameters:

26: $\theta_d \leftarrow \theta_d + \alpha \nabla_{\theta_d} L_{D_{total}}$

Calculate generator loss:

27: $L_G = -\log(D_{fake})$

Add L2 regularization:

28: $L_{G_{total}} = L_G + \lambda ||\theta_g||^2$

update generator parameters:

29: $\theta_g \leftarrow \theta_g + \alpha \nabla_{\theta_g} L_{G_{total}}$

30: optionally update learning rates α, β based on the decay schedule

31: Repeat steps 13 until convergence or $\text{epoch}_{maximum}$

32: Evaluate model performance

33: Compute performance metrics on validation set

34: Adjust hyperparameters or model architecture

The pseudocode gives an overview of the entire adversarial training cycle of the suggested cGAN framework. It starts with the setting of the Generator and Discriminator parameters and trains the two networks repeatedly in each training epoch. The Generator synthesizes conditioned on the class label X-ray images from sampled noise in batches, using a Generator before the Generator. These sampled images, along with true images, are input into the Discriminator, which calculates adversarial loss to identify the authentic and synthetic data. The Generator parameters are then updated to increase its classification boundary, and the Discriminator parameters are again updated with the gradient of the adversarial loss to enhance its capacity to fool the Discriminator. The process is repeated until convergence, allowing the Generator to produce high-fidelity, label-consistent synthetic images. This can be achieved through this iterative process of optimization that is essential for achieving stable training, preventing mode collapse, and generating synthetic samples that can be used to balance the pneumonia dataset.

The structured cGAN method creates artificial intelligence training to generate clinical-quality synthetic chest X-ray images. Initially, the model commences by setting up the foundational components: The model consists of a generator and a discriminator with their separate parameter settings. The generator transforms noise data and disease labels into images while the discriminator checks original and generated images to confirm accuracy within the given labels. During each training cycle, noise vectors are drawn from a predefined noise distribution, and real data samples, consisting of images paired with their corresponding labels, are gathered from the dataset. The generator merges each noise vector with its respective label,

creating a conditioned input that is then transformed into a synthetic image. These images are designed to mimic the appearance of real X-ray images while reflecting the specific features associated with the input label, such as indications of pneumonia.

The proposed medical cGAN also deviates from the traditional formulation through several major architectural enhancements designed to address the synthesis of chest X-rays. First, attention blocks are added after the first convolutional layers of the Generator, allowing the model to focus on specific regions of the lungs associated with pneumonia and then upsample them. Second, it employs better normalization; specifically, it uses batch normalization, which is strategically implemented in both the Generator and Discriminator to stabilize convergence and maintain radiological contrast. Lastly, to align the activation pipeline with medical image intensities, Sigmoid is used at the output layer, and non-linear activations, maximized to identify fine pathological structures, are used in the intermediate layers. Combining all these architectural modifications can enhance fidelity, structural realism, and feature presentation that are disease-relevant, making the proposed cGAN unique compared to the standard version.

4. Results And Discussion

This research seeks to boost chest X-ray pneumonia diagnosis using a new cGAN as an image-detection method. Our model was created to address medical image analysis problems, particularly class imbalance and insufficient labeling information. We have selected the NIH Chest X-ray Database as the data source for this work. Our selection involves this data collection because it offers detailed labels and contains a substantial number of X-ray images covering many patients from different groups in varying disease stages. The dataset shows numerous imaging options demonstrating how patients vary by age and health conditions while exhibiting different image quality. These examples illustrate the wide range of medical images the proposed cGAN model needs to produce reliable results for different types of patients. The results are compared with existing works [26] – [30]. The dataset used in this study comprises 1,200 chest X-ray images, evenly distributed into 600 pneumonia cases and 600 normal cases, ensuring a balanced class representation at the source level. For experimental evaluation, the dataset was partitioned using a 70:15:15 split, resulting in 840 images for training, 180 for validation, and 180 for testing. Among these, the training set contained 420 pneumonia and 420 normal images, which served as the only portion subjected to cGAN-based augmentation.

The validation and test sets were deliberately kept unchanged to preserve the integrity of unbiased performance measurement. Prior to model training, all images underwent a unified preprocessing pipeline that included grayscale conversion, resizing to 256×256 pixels, contrast enhancement, intensity normalization, and noise suppression, ensuring uniform input quality. This explicit reporting of image counts and preprocessing steps provides transparency in dataset handling and validates the reliability of the subsequent cGAN generation and classifier evaluation.




Figure 3. Sample Images from Prescribed Dataset

The large number of clinical samples makes the results predictable and applicable to multiple medical settings. The cGAN model is thoroughly assessed against five advanced medical image processing methods seeking comparison. The chosen methods reflect their current usage and matching purposes making the performance evaluation strong. The evaluation methods include testing the model's accuracy and its ability to detect specific situations while also measuring how well it distinguishes true results from false ones. The specific set of evaluation metrics helps us measure different properties of how well the model diagnoses while showing its success in balancing detection rates against false positives. Our findings prove the cGAN model excellence and show how artificial intelligence can enhance medical diagnosis systems. Model input samples appear as grayscale images in Figure 4 before data processing with the model starts. X-ray images become easier to analyze when doctors convert them to grayscale to see structural details while saving the system's computing power. With clear input data the cGAN performs better because it can focus on only important medical traits. The augmentation threshold was reached by the creation of synthetic images one after another and the analysis of their quality in terms of structure by SSIM and FID scores. Synthetic samples that passed the quality check indicator were only kept and generation was propagated until the count of validated synthetic pneumonia images matched that of normal images. This ensure that the final dataset was balanced in classes, and still the image fidelity was not reduced.




Figure 4. Gray Scale Images of the Input Samples

Figure 5 displays the original image, the cGAN-created synthetic output and a difference image based on one case study. The synthetic X-ray consistently matches real X-ray data while preserving essential diagnostic features according to the cGAN. Because of its high accuracy the synthetic image minimizes the detection of small discrepancies between real and synthetic images. Our approach helps synthetic data integrate with genuine medical images without distorting them.




Figure 5. Real Image, Synthetic and the difference Image (Sample 1)

Figure 6 shows another sample that confirms the cGAN's ability to produce good synthetic images for all given input cases. The model demonstrates good stability, as it produces similar high-quality results for many samples. This confirms its reliability for dataset expansion in medical image diagnostics.




Figure 6. Real Image, Synthetic and the Difference Image (Sample 2)

Figure 7 shows a third case where our model creates a synthetic duplicate that matches the original image precisely including subtle details. Medical imaging requires the system to spot small differences that impact diagnosis, so this detail detection matters.




Figure 7. Real Image, Synthetic and the Difference Image (Sample 3)

Figure 8 traces the small differences in pixel values between original and produced images. The cGAN produces images that stay close to actual data because the generated image heatmap shows low variation. Our high-level detail creation methods produce medical images that help expand training data while maintaining diagnostic accuracy.

The heatmap shows that cGAN-generated X-ray images of the chest share fundamental structures and textures with actual X-ray scans. The significant differences in the heatmap indicate that the cGAN successfully generated lung opacities and infiltrates with minimal added noise. The detailed approach produces synthetic images that look real and serve meaningful diagnostic functions. The model shows its ability to address imbalance issues by creating medical-grade variations within its generated samples. The heatmap in Figure 10 helps confirm that the synthetic images accurately reflect real-world clinical data which ensure better diagnostic accuracy in medical practice. The maximum difference is apparent in the pixel difference heat map in the lung areas, especially around the mid-to-lower lobes where opacities of pneumonia are most likely to be found. The more intense deviations between real and synthetic images are observed in these areas, indicating that the generator is more oriented towards the reconstruction of pathological patterns like infiltrates and consolidation, which are associated with greater structural complexity.




Figure 8. Pixel Difference Heatmap

Our research compares the accuracy results of cGANs with DCNNs, Attention Mechanisms, Hybrid CNN-LSTM models, and standard machine learning methods, as shown in Figure 9. The accuracy measurement helps us determine whether the model correctly identifies both pneumonia and non-pneumonia cases, demonstrating its diagnostic quality. Our cGAN model proves superior in accuracy compared to other methods because it addresses the issues of data imbalance and limited variation. Our results indicate improved training outcomes due to cGAN's capability to generate authentic and diverse synthetic medical images for reinforcement.




Figure 9. Comparative Analysis of Accuracy of cGAN

The proposed cGAN-based model achieves a detection rate of approximately 97%, demonstrating superior sensitivity in identifying pneumonia cases compared to all baseline methods. Table 1 shows how the proposed cGAN model outperforms DCNNs, and other advanced methods including Attention Mechanisms Hybrid CNN-LSTM models and traditional machine learning techniques for specific case identification.

Table 2. Consolidated Performance Comparison Table

Model	Accuracy	Sensitivity	Specificity	F1-Score
cGAN	0.95	0.96	0.94	0.945
Enhanced CNN [26]	0.91	0.94	0.91	0.93
Multi-task Learning [27]	0.93	0.95	0.90	0.935
Advanced GANs [28]	0.90	0.89	0.88	0.89
Federated Learning [29]	0.91	0.90	0.89	0.90
Explainable AI [30]	0.94	0.95	0.93	0.94

The cGAN model shows a better true positive rate than other tested approaches because it efficiently recognizes pneumonia cases under any situation. The cGAN produces high-quality synthetic images that both increase dataset diversity and balance class distribution, making it better at detecting pneumonia. The cGAN trains better by using a wider range of pneumonia patterns across both typical and uncommon cases, which allows it to recognize different types of pneumonia more accurately than other models do. The advanced design of the cGAN architecture selects and focuses on prominent pneumonia clues to capture every detail in its image generation process. Profiling tests how well the system can prevent healthy individuals from being identified as pneumonia patients. The cGAN model outperforms all other methods by correctly identifying and testing fewer normal chest X-ray instances. Additionally, the F1 score of our proposed cGAN model stands against DCNN, Attention Mechanisms, Hybrid CNN-LSTM models, and basic machine learning in terms of performance. The F1 score offers an ideal performance measure because it combines precision and recall, helping us evaluate exactly how well the model detects pneumonia from X-rays. The approach delivers better

results in datasets where one class appears much less often than others, especially in medical scans used to find pneumonia.

Figure 10 shows the Area Under the Curve results for our cGAN approach alongside DCNNs, Attention Methods, CNN-LSTM combinations, and basic machine learning algorithms. Through the AUC analysis, our model demonstrates how well it detects pneumonia and non-pneumonia cases over many decision threshold points. The model shows strong separation skills when its AUC value reaches higher numbers because it prevents both incorrect positive and negative results at different decision points.




Figure 10. Comparative Analysis of AUC of cGAN

The cGAN model evaluation through Figure 11 shows how its performance matches with different machine learning approaches including DCNNs and Attention Mechanisms. The cGAN system proves its efficiency by producing quality images at the same level rate as compared to other systems despite its advanced setup and generation power. The model design keeps the workload low for both making synthetic images and classifying medical photos. The design enhances performance through training improvements and structural optimizations that simplify the model without losing its diagnostic precision or stability.




Figure 11. Comparative Analysis of Computational Efficiency

The new framework demonstrates better speed while performing data synthesis tasks when measured against leading industry models. This section presents more details about our

experimental procedures and lists the hardware tools used. Our model, trained on an NVIDIA Tesla V100 GPU with 32GB of memory, utilized TensorFlow 2.0 for the backend. Our cGAN system required about 12 hours of training on 50 epochs of 10,000 samples and could produce one image in 0.02 seconds. Other techniques took more time to complete because they employed more complicated hybrid CNN-LSTM and attention-based models. The competitive performance of the cGAN stands out as it achieved better results than other models across all evaluated performance metrics. Our method successfully combines efficient processing speed and accurate test results to create a workable tool for diagnosing pneumonia. The cGAN shows clear potential for real-world use because it provides precise results while remaining affordable to run in various healthcare settings.




Figure 12. Training vs. Validation Accuracy Analysis

Figure 12 shows the accuracy of training and validation over 30 epochs. The training accuracy displays a gradual and steady increasing trend, reaching 97% in the final epoch. The validation accuracy also exhibits an overall positive trend with slight oscillations, providing evidence of sound generalization without overfitting. This behavior demonstrates that the proposed cGAN-augmented dataset enables the classifier to effectively learn discriminative pneumonia features and achieve a high-performance operating point.




Figure 13. Training vs. Validation Loss Analysis

Figure 13 shows the loss curves correspondingly. The loss during the training process drops in the very first epochs and slowly levels off as the model approaches convergence. The validation loss has a comparable downward trend with minor variations/oscillations, which are common in medical imaging tasks because of inter-sample variations. The lack of difference between the two curves proves the existence of stable learning dynamics and supports the validity of the optimized classifier that was trained using the balanced dataset.

The SSIM analysis shows how synthetic chest X-ray images resemble genuine chest X-ray images, as presented in Table 3. The SSIM results indicate that real X-rays maintain a stable score above 0.91, while synthetic images maintain a score between 0.87 and 0.90. Synthetic images accurately retain important structural elements that X-ray specialists use for diagnosis, such as detecting lung lesions and pulmonary infection signals.

Table 3. SSIM Analysis

Sample Index	Real Image SSIM	Synthetic Image SSIM
1	0.92	0.88
2	0.94	0.89
3	0.93	0.88
4	0.91	0.87
5	0.95	0.89

SSIM is calculated using luminance, contrast, and structural comparison between two images, as follows:

$$SSIM(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (10)$$

In Equation (10), μ_x and μ_y are the mean intensities, σ_x^2 and σ_y^2 are the variances, and σ_{xy} is the covariance between images x (real) and y (synthetic). Constants C_1 and C_2 stabilize the division. The high SSIM scores validate that the cGAN effectively generates images with structural integrity and diagnostic relevance. By ensuring that synthetic images closely mimic real ones, SSIM supports the objective of addressing class imbalance without compromising image quality and improving classification performance.

Table 4. Class Distribution Analysis

Class	Before Augmentation	After Augmentation
Pneumonia	200	800
Non-Pneumonia	1000	1000

The class distribution table in Table 4 shows that the augmentation method achieved its goal because it generated balanced datasets. The training process benefits from balanced classes and realistic synthetic images, as they produce better results demonstrated through different

performance indicators. These tests verify that the cGAN system meets its goals by connecting synthetic data production to real world medical practice.

The findings of the proposed cGAN framework indicate that the approach can produce high-quality synthetic pediatric chest X-ray images that approximate the underlying properties of the original dataset. The synthetic images were able to achieve high SSIM and FID scores, which means that the model did not alter the structural patterns, intensity distributions, or disease-associated visual cues of high fidelity. These validated synthetic samples were then combined with the original data to create a balanced training set, which was effective in addressing the limited sample size of the dataset and minimizing the possibility of model bias due to small or skewed medical datasets. The visual inspection also confirmed that the images produced by cGAN were anatomically realistic and did not produce any artifacts. The validated results of the generative outputs, through quantitative and qualitative means, provide clear evidence that the suggested cGAN model improves the diversity of the datasets and contributes to a more adequate analysis of images without affecting the integrity of the validation and testing partitions. This clear display of the results is a direct indication of the success of the generative model and enhances the overall quality of the outcomes section.

5. Conclusion

The proposed cGAN framework improves the detection of pneumonia by generating high-quality synthetic chest X-ray images to balance medical samples and increase diversity. This significantly enhances diagnostic results. The model surpasses the performance of traditional approaches in terms of precision and speed within clinical scenarios and can be optimized for stability in GAN training and fidelity of generated images. Its application cannot yet be fully promoted due to some obstacles, such as unbalanced datasets and reliance on a single benchmark, which complicate its transfer into different contexts of medical practice. Future research will outline its effectiveness with various clinical datasets and assess its potential in diagnosing tuberculosis and COVID-19.

Declaration

- Funding – The author(s) did not receive support from any organization for the submitted work.
- Conflicts of Interest - The author(s) has no relevant financial or non-financial interests to disclose.
- Ethics Approval – The paper is an original contribution of research and is not published elsewhere in any form or language.
- Consent Statement – All authors mentioned have contributed towards the research work, drafting of the paper as well as have given consent for publishing of this article.
- Availability of Data & Material – The benchmark dataset the NIH Chest X-ray Database is used to evaluate the performances (<https://www.kaggle.com/datasets/nih-chest-xrays/data>). The data repository is open access.

- Consent to publication – all authors listed above have consented to get their data and image published.

References

- [1] Alapat, Daniel Joseph, Malavika Venu Menon, and Sharmila Ashok. "A Review on Detection of Pneumonia in Chest X-ray Images Using Neural Networks." *Journal of biomedical physics & engineering* 12, no. 6 (2022): 551.
- [2] Siddiqi, Raheel, and Sameena Javaid. "Deep Learning for Pneumonia Detection in Chest X-ray Images: A Comprehensive Survey." *Journal of imaging* 10, no. 8 (2024): 176.
- [3] Tang, Yu-Xing, You-Bao Tang, Yifan Peng, Ke Yan, Mohammadhadi Bagheri, Bernadette A. Redd, Catherine J. Brandon et al. "Automated Abnormality Classification of Chest Radiographs Using Deep Convolutional Neural Networks." *NPJ digital medicine* 3, no. 1 (2020): 70.
- [4] Rahmat, Taufik, Azlan Ismail, and Sharifah Aliman. "Chest X-ray Image Classification Using Faster R-CNN." *Malaysian Journal of Computing (MJoC)* 4, no. 1 (2019): 225-236.
- [5] Li, Yuanyuan, Zhenyan Zhang, Cong Dai, Qiang Dong, and Samireh Badrigilan. "Accuracy of Deep Learning for Automated Detection of Pneumonia Using Chest X-ray Images: A Systematic Review and Meta-Analysis." *Computers in Biology and Medicine* 123 (2020): 103898.
- [6] Singh, Sukhendra, Manoj Kumar, Abhay Kumar, Birendra Kumar Verma, Kumar Abhishek, and Shitharth Selvarajan. "Efficient Pneumonia Detection Using Vision Transformers on Chest X-rays." *Scientific reports* 14, no. 1 (2024): 2487
- [7] Asif, Sohaib, Ming Zhao, Fengxiao Tang, and Yusen Zhu. "A Deep Learning-Based Framework for Detecting COVID-19 Patients Using Chest X-rays." *Multimedia Systems* 28, no. 4 (2022): 1495-1513.
- [8] Singh, Sukhendra, Manoj Kumar, Abhay Kumar, Birendra Kumar Verma, and S. Shitharth. "Pneumonia Detection with QCSA Network on Chest X-ray." *Scientific Reports* 13, no. 1 (2023): 9025.
- [9] Duong, Linh T., Phuong T. Nguyen, Ludovico Iovino, and Michele Flammini. "Automatic Detection of Covid-19 from Chest X-ray and Lung Computed Tomography Images Using Deep Neural Networks and Transfer Learning." *Applied Soft Computing* 132 (2023): 109851.
- [10] Singh, Sukhendra, and B. K. Tripathi. "Pneumonia Classification Using Quaternion Deep Learning." *Multimedia Tools and Applications* 81, no. 2 (2022): 1743-1764.
- [11] Mabrouk, Alhassan, Rebeca P. Diaz Redondo, Abdelghani Dahou, Mohamed Abd Elaziz, and Mohammed Kayed. "Pneumonia Detection on Chest X-Ray Images Using Ensemble of Deep Convolutional Neural Networks." *Applied Sciences* 12, no. 13 (2022): 6448.

- [12] Wu, Huaiguang, Pengjie Xie, Huiyi Zhang, Daiyi Li, and Ming Cheng. "Predict Pneumonia with Chest X-Ray Images Based on Convolutional Deep Neural Learning Networks." *Journal of Intelligent & Fuzzy Systems* 39, no. 3 (2020): 2893-2907.
- [13] Chakravarthi, Sangapu Sreenivasa, Shaik Nagoor Meeravali, Mohammad Aazmi Irfan, S. Sountharajan, and E. Suganya. "Pneumonia Detection Using Chest X-Rays: A Comprehensive Review." In *International Conference on Computational Intelligence in Data Science*, Cham: Springer Nature Switzerland, 2024, 292-305.
- [14] Reshan, Mana Saleh Al, Kanwarpartap Singh Gill, Vatsala Anand, Sheifali Gupta, Hani Alshahrani, Adel Sulaiman, and Asadullah Shaikh. "Detection of Pneumonia from Chest X-ray Images Utilizing Mobilenet Model." In *Healthcare*, vol. 11, no. 11, MDPI, 2023, 1561.
- [15] Sitaula, Chiranjibi, and Mohammad Belayet Hossain. "Attention-based VGG-16 Model for COVID-19 Chest X-ray Image Classification." *Applied Intelligence* 51, no. 5 (2021): 2850-2863.
- [16] Sharma, Shagun, and Kalpna Guleria. "A Systematic Literature Review on Deep Learning Approaches for Pneumonia Detection Using Chest X-ray Images." *Multimedia Tools and Applications* 83, no. 8 (2024): 24101-24151.
- [17] Barakat, Natali, Mahmoud Awad, and Bassam A. Abu-Nabah. "A Machine Learning Approach on Chest X-rays for Pediatric Pneumonia Detection." *Digital Health* 9 (2023): 20552076231180008.
- [18] El Houbay, Enas MF. "COVID-19 Detection from Chest X-ray Images Using Transfer Learning." *Scientific Reports* 14, no. 1 (2024): 11639.
- [19] Jangam, Ebenezer, Aaron Antonio Dias Barreto, and Chandra Sekhara Rao Annavarapu. "Automatic Detection of COVID-19 from Chest CT scan and Chest X-rays Images Using Deep Learning, Transfer Learning and Stacking." *Applied Intelligence* 52, no. 2 (2022): 2243-2259.
- [20] Zhang, Dejun, Fuquan Ren, Yushuang Li, Lei Na, and Yue Ma. "Pneumonia Detection from Chest X-ray Images Based on Convolutional Neural Network." *Electronics* 10, no. 13 (2021): 1512.
- [21] Verma, Devvret, Chitransh Bose, Neema Tufchi, Kumud Pant, Vikas Tripathi, and Ashish Thapliyal. "An Efficient Framework for Identification of Tuberculosis and Pneumonia in Chest X-ray Images Using Neural Network." *Procedia Computer Science* 171 (2020): 217-224.
- [22] Stephen, Okeke, Mangal Sain, Uchenna Joseph Maduh, and Do-Un Jeong. "An Efficient Deep Learning Approach to Pneumonia Classification in Healthcare." *Journal of healthcare engineering* 2019, no. 1 (2019): 4180949.
- [23] Khasawneh, Natheer, Mohammad Fraiwan, Luay Fraiwan, Basheer Khassawneh, and Ali Ibnian. "Detection of Covid-19 from Chest X-ray Images Using Deep Convolutional Neural Networks." *Sensors* 21, no. 17 (2021): 5940.

- [24] Zouch, Wassim, Dhouha Sagga, Amira Echtioui, Rafik Khemakhem, Mohamed Ghorbel, Chokri Mhiri, and Ahmed Ben Hamida. "Detection of COVID-19 from CT and Chest X-ray Images Using Deep Learning Models." *Annals of Biomedical Engineering* 50, no. 7 (2022): 825-835.
- [25] Kora, Padmavathi, Chui Ping Ooi, Oliver Faust, U. Raghavendra, Anjan Gudigar, Wai Yee Chan, K. Meenakshi, K. Swaraja, Pawel Plawiak, and U. Rajendra Acharya. "Transfer Learning Techniques for Medical Image Analysis: A Review." *Biocybernetics and biomedical engineering* 42, no. 1 (2022): 79-107.
- [26] Al-Sheikh, Mona Hmoud, Omran Al Dandan, Ahmad Sami Al-Shamayleh, Hamid A. Jalab, and Rabha W. Ibrahim. "Multi-Class Deep Learning Architecture for Classifying Lung Diseases from Chest X-ray and CT Images." *Scientific Reports* 13, no. 1 (2023): 19373.
- [27] Wang, Linda, Zhong Qiu Lin, and Alexander Wong. "COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest X-ray Images." *Scientific reports* 10, no. 1 (2020): 19549.
- [28] Nahiduzzaman, Md, Md Omaer Faruq Goni, Md Robiul Islam, Abu Sayeed, Md Shamim Anower, Mominul Ahsan, Julfikar Haider, and Marcin Kowalski. "Detection of Various Lung Diseases Including COVID-19 Using Extreme Learning Machine Algorithm Based on the Features Extracted from A Lightweight CNN Architecture." *Biocybernetics and Biomedical Engineering* 43, no. 3 (2023): 528-550.
- [29] Pan, Zegang, Haijiang Wang, Jian Wan, Lei Zhang, Jie Huang, and Yangyu Shen. "Efficient Federated Learning for Pediatric Pneumonia on Chest X-ray Classification." *Scientific Reports* 14, no. 1 (2024): 23272.
- [30] Veeramani, Nirmala, Reshma Sherine SA, Sakthi Prabha S, Srinidhi S, and Premaladha Jayaraman. "NextGen Lung Disease Diagnosis with Explainable Artificial Intelligence." *Scientific Reports* 15, no. 1 (2025): 33052.
- [31] Szepesi, Patrik, and László Szilágyi. "Detection of Pneumonia Using Convolutional Neural Networks and Deep Learning." *Biocybernetics and biomedical engineering* 42, no. 3 (2022): 1012-1022.
- [32] An, Qiuyu, Wei Chen, and Wei Shao. "A Deep Convolutional Neural Network for Pneumonia Detection in X-ray Images with Attention Ensemble." *Diagnostics* 14, no. 4 (2024): 390.
- [33] Lafraxo, Samira, Mohamed El Ansari, and Lahcen Koutti. "A New Hybrid Approach for Pneumonia Detection Using Chest X-rays Based on ACNN-LSTM and Attention Mechanism." *Multimedia Tools and Applications* 83, no. 29 (2024): 73055-73077.
- [34] Motamed, Saman, Patrik Rogalla, and Farzad Khalvati. "Data Augmentation Using Generative Adversarial Networks (GANs) for GAN-Based Detection of Pneumonia and COVID-19 in Chest X-ray Images." *Informatics in medicine unlocked* 27 (2021): 100779.
- [35] Ranjan, Amit, Chandrashekhar Kumar, Rohit Kumar Gupta, and Rajiv Misra. "Transfer Learning Based Approach for Pneumonia Detection Using Customized VGG16 Deep

- Learning Model." In International Conference on Internet of Things and Connected Technologies, Cham: Springer International Publishing, 2021,17-28.
- [36] Arun Prakash, J., C. R. Asswin, Vinayakumar Ravi, V. Sowmya, and K. P. Soman. "Pediatric Pneumonia Diagnosis Using Stacked Ensemble Learning on Multi-Model Deep CNN Architectures." *Multimedia tools and applications* 82, no. 14 (2023): 21311-21351.
- [37] Ukwuoma, Chiagoziem C., Zhiguang Qin, Md Belal Bin Heyat, Faijan Akhtar, Olusola Bamişile, Abdullah Y. Muaad, Daniel Addo, and Mugahed A. Al-Antari. "A Hybrid Explainable Ensemble Transformer Encoder for Pneumonia Identification from Chest X-ray Images." *Journal of Advanced Research* 48 (2023): 191-211.