

A GAN-Augmented And Bio-Inspired Deep Learning Framework For Comprehensive Multi-Disease Detection In Chest Radiographs

Jignesh Solanki¹ and Vikram Agrawal²

¹ Research Scholar, Gujarat Technological University, Ahmedabad, Gujarat, India

² Assistant Professor, Bhailalbhai & Bhikhabhai Institute of Technology, Anand, Gujarat, India

jsolanki85@gmail.com¹. agrawal.vm@gmail.com²

Abstract. This research presents DNN_AO, a specialized deep learning architecture that leverages the Aquila Optimizer to automate the identification of various pulmonary conditions via chest X-rays. Deviating from traditional single-diagnosis approaches, DNN_AO employs a multi-faceted design to overcome data limitations. Specifically, the framework utilizes Generative Adversarial Networks (GANs) to synthesize data for underrepresented classes, ensuring a balanced dataset. It further optimizes performance by using the bio-inspired Aquila algorithm for hyperparameter tuning, paired with a hybrid feature extraction method that integrates a ResNet50 backbone with custom dense layers. Testing on the ChestX-ray14 dataset—comprising 5,606 images across 14 distinct pathologies—yielded an accuracy of 89%. This result consistently outperforms standard classifiers, including CNNs, Random Forests, and KNN. Notably, the model demonstrates high sensitivity for rare conditions such as hernias, marking a significant advancement toward holistic, multi-disease clinical screening.

Keywords: Deep Learning · Chest X-ray (CXR) Analysis · Pulmonary Disease Classification · Artificial Intelligence (AI) in Medical Imaging · Automated Diagnosis

1. Introduction

Pulmonary disorders remain a premier global health crisis, contributing to millions of annual fatalities and hospitalizations. The emergence of COVID-19, alongside persistent threats like tuberculosis, pneumonia, and lung cancer, underscores the vital necessity for rapid and dependable diagnostic frameworks. While early intervention is known to drastically reduce mortality, conventional diagnostics—relying on manual radiologist interpretation and clinical exams—are frequently bottlenecked by human error and extended processing times. This is especially critical in underserved regions where a deficit of specialized radiologists hinders timely care. Consequently, there is a significant shift toward embedding Artificial Intelligence (AI) within clinical workflows to augment diagnostic precision. Chest X-rays (CXR) serve as the primary, cost-effective modality for identifying thoracic abnormalities. However, CXR interpretation is notoriously complex; overlapping pathological signatures, poor tissue contrast, and hardware-induced image variability often lead to diagnostic inconsistencies. This challenge is compounded in cases of comorbidity, where multiple lung conditions coexist. While Deep Learning (DL) and Convolutional Neural Networks (CNNs) have revolutionized feature extraction in medical imaging, most existing literature focuses on single-disease classification. Such binary models

lack the versatility required for real-world clinical environments where overlapping disorders are common. To address these gaps, this research proposes DNN_AO, an innovative Deep Neural Network architecture integrated with the Arithmetic Optimization Algorithm (AOA). Designed for multi-label pulmonary detection, DNN_AO merges the high-dimensional representational power of deep networks with the global search efficiency of metaheuristic optimization. By dynamically tuning hyperparameters—including learning rates and weight distributions—the AOA ensures accelerated convergence and superior generalization. This hybrid approach allows the model to discern both subtle and prominent pathological features while mitigating the risks of overfitting and high computational demand.

Unlike traditional models, DNN_AO's multi-label capability allows for the simultaneous identification of various abnormalities within a single radiograph, mirroring actual clinical presentations. To bolster performance against data scarcity and noise, the framework integrates advanced preprocessing—such as contrast normalization and noise suppression—alongside rigorous data augmentation (scaling, rotation, and intensity shifts). These combined strategies ensure that the model remains resilient across diverse patient populations and varying imaging equipment, providing a stable tool for modern diagnostic support.

The primary drive behind this study is the critical and urgent need for fast, precise, and accessible diagnostic tools for two major diseases: COVID-19 and pneumonia. These illnesses are tough for doctors to manage because their symptoms are so similar, creating significant diagnostic challenges. Although traditional testing methods, like RT-PCR, are effective, they are often expensive, slow, and simply unavailable in many areas with limited healthcare resources.

Furthermore, pneumonia, whether caused by COVID-19 or other agents, continues to be a major global health burden contributing to high illness and death rates. Relying on humans to read chest X-rays is inherently subjective, and the overlapping appearance of different respiratory infections on the images makes accurate diagnosis difficult. Therefore, there is a clear demand for automated tools that can help radiologists precisely tell these conditions apart

2. Related Work

The rapid progress of deep learning (DL) has significantly transformed medical image analysis, particularly in the detection of pulmonary diseases such as COVID-19 and pneumonia. During the early phase of the COVID-19 pandemic, artificial intelligence-based diagnostic systems emerged as an essential aid for radiologists in handling the growing volume of imaging data. Initial studies by Ozturk et al. [1] demonstrated the capability of Convolutional Neural Networks (CNNs) to recognize subtle radiographic patterns in chest X-ray images that are often challenging to detect through manual examination. Building on this work, Apostolopoulos and Mpesiana [7] showed that transfer learning approaches employing pre-trained models such as VGG19 and MobileNet could deliver reliable classification performance even when only limited annotated medical data were available.

Later research efforts emphasized the development of task-specific architectures and improved clinical interpretability. Models including CoroNet [6] and COVIDX-Net [9] focused on enhancing classification accuracy for COVID-19 cases, while Alshazly et al. [2] combined chest X-ray and CT imaging with explainable AI techniques to highlight infection-affected lung regions. To further reduce false-positive predictions and improve robustness, hybrid and ensemble-based frameworks—such as COVID-CheXNet [10] and VDSNet [4]—integrated multiple deep learning strategies and optimization mechanisms, resulting in more stable and reliable diagnostic outcomes.

2.1 Problem Formulation

Despite notable advancements in medical image analysis, the simultaneous and accurate classification of a wide range of pulmonary diseases from chest X-ray (CXR) images remains a significant challenge. Conventional machine learning approaches, including Random Forest (RF), K-Nearest Neighbors (KNN), and Support Vector Machines (SVM), often demonstrate limited generalization performance when applied to complex medical imaging tasks. These traditional classifiers face particular difficulty in handling high-dimensional features associated with multi-class classification problems, as well as the pronounced class imbalance commonly present in clinical datasets. Moreover, many existing Convolutional Neural Network (CNN) architectures lack the robustness and consistency required for reliable deployment in real-world clinical environments, where variations in imaging quality, acquisition devices, and disease manifestations are substantial.

To overcome these challenges, this study proposes a novel Deep Neural Network model, referred to as DNN_AO, designed specifically for the classification of fourteen distinct pulmonary conditions using standard CXR images. The primary objective of this work is to develop a scalable, interpretable, and high-performance diagnostic framework that consistently outperforms conventional methods in terms of accuracy, stability, and clinical applicability, particularly in resource-constrained healthcare settings.

3. Proposed Methodology

3.1 Dataset and Preprocessing

This research utilizes the ChestX-ray14 dataset, a fundamental benchmark in medical imaging research originally introduced by Wang et al. The dataset comprises 5,606 frontal-view chest X-ray images, each annotated with one or more of 14 different thoracic pathologies, including conditions such as Atelectasis, Cardiomegaly, and Nodule. Each image is also associated with basic patient demographic data, such as age and gender. To ensure the reliability and generalizability of our results and strictly prevent data leakage, a rigorous patient-based split was implemented. This critical step guarantees that images belonging to the same individual do not appear in both the training and the final testing sets. The complete dataset was divided into 86,524 images for model training and 25,596 images for performance testing. We noted that the original data partitioning exhibited an inherent bias: patients in the original test set contributed approximately three times as many images as those in the training set. To mitigate this unequal distribution and ensure a balanced partition, we engineered a custom data split using a fixed seed value of 2137. A major obstacle encountered during dataset preparation was the severe class imbalance. While certain common conditions, such as Infiltration, were abundantly represented, rare diseases like Hernia constituted less than 0.2% of the total dataset. This skewed distribution necessitated the implementation of sophisticated augmentation strategies to ensure the model could effectively learn the features of these underrepresented conditions.

3.2 Data Augmentation

To address the severe class imbalance, particularly for rare conditions like Hernia and Fibrosis, advanced data augmentation was employed. Initial basic techniques included rotation, flipping, scaling, and contrast adjustment to introduce variability. However, these were insufficient to bridge the large data gap between majority and minority classes. To enhance model generalization and diversify samples from underrepresented classes, Generative Adversarial Networks (GANs) were integrated into the pipeline. GANs are powerful generative models that produce realistic, high-quality synthetic images. The framework consists of two adversarial components: a Generator that creates fake images mimicking real CXR patterns, and a Discriminator that distinguishes between real and generated images. Through this minimax game, the Generator progressively improves its output. In this study, GAN-based augmentation was targeted at the first eleven disease types, each with fewer than 500 original samples. By synthesizing images for these classes, the dataset was balanced, reducing training bias and improving the model's predictive capability for rare lung conditions.

3.3 Feature Engineering and Modelling

The proposed DNN_AO model represents an advanced hybrid framework that integrates transfer learning, a customized Deep Neural Network (DNN), and metaheuristic optimization to achieve robust image classification. The process initiates with feature extraction utilizing the ResNet50 architecture, pre-trained on the ImageNet dataset. The penultimate layer of ResNet50 generates a 2048-dimensional feature vector for each input image, effectively capturing complex spatial and contextual relationships within the data. These extracted feature vectors are subsequently fed into a custom DNN classifier, the architecture of which is not predetermined but is instead optimized using the Aquila Optimizer (AO)—a bio-inspired metaheuristic algorithm that simulates the predatory behavior of Aquila eagles. The AO algorithm autonomously fine-tunes key hyperparameters of the DNN, including:

- The number of hidden layers
- The number of neurons per layer
- The type of activation function (e.g., ReLU, Sigmoid)
- The dropout rate for effective regularization

Through this adaptive optimization process, the DNN_AO framework constructs an efficient, task-specific, and high-performing network architecture, thereby mitigating the limitations associated with manual hyperparameter tuning and enhancing overall classification accuracy.

3.4 System Block Diagram

The proposed DNN_AO framework presents several key contributions that enhance its technical and scientific rigor. First, it integrates Generative Adversarial Networks (GANs) for minority class augmentation with an Aquila Optimization (AO)-tuned Deep Neural Network. This hybrid combination ensures balanced data representation, improved feature learning, and reduced overfitting compared to conventional CNN-based systems. Second, the proposed model introduces an adaptive hyperparameter optimization strategy inspired by Aquila birds' hunting behavior, enabling efficient convergence and high model accuracy. Third, the inclusion of ResNet50-based feature extraction leverages transfer learning to capture complex spatial patterns from chest X-rays, while the AO algorithm fine-tunes parameters to maximize classification performance. Fourth, the methodology addresses dataset imbalance by synthesizing rare disease samples using GANs, making it highly scalable for real-world clinical deployment. Lastly, the integration of data augmentation, transfer learning, and metaheuristic optimization establishes a generalized

framework capable of handling multi-class disease classification tasks across 14 pulmonary conditions.




Fig. 1. illustrates the step-by-step operational workflow of the proposed diagnostic framework.

The block diagram (Fig. 1) illustrates the step-by-step workflow of the proposed framework:

1. **Input Image:** The process starts with chest X-ray images taken from the ChestX-ray14 dataset. These images act as the main input to the system and showcase many different thoracic diseases.
2. **Data Augmentation:** To handle imbalance in the dataset and make the system better at generalizing various augmentation methods are used. Basic changes like rotating, flipping, and adjusting contrast are applied. A Generative Adversarial Network also creates lifelike chest X-rays to boost the number of less-represented conditions like hernia and fibrosis.
3. **Feature Extraction (ResNet50):** After augmenting the images, a pre-trained ResNet50 model extracts features. This neural network trained on ImageNet, identifies detailed patterns within the images. The penultimate layer of the network generates 2048-dimensional feature vectors, which serve as input to perform classification. The DNN_AO model uses extracted feature vectors as input for a Deep Neural Network. The Aquila Optimization algorithm fine-tunes this network. Inspired by the way Aquila birds hunt, this metaheuristic adjusts hyperparameters like layers, neurons, activation functions, and dropout rates. After optimization, the DNN_AO model trains on these features and classifies them into 14 categories of pulmonary diseases.
4. **Disease Classification:** The system's final layer identifies the predicted disease class for every input image. It provides probability scores across all possible disease categories and picks the one with the most confidence. This step allows reliable and automated detection of several lung diseases

3.4.1 Algorithmic Description of the Proposed DNN_AO Framework

The DNN_AO algorithm combines feature extraction, metaheuristic optimization, and deep classification into a single integrated pipeline. The method automates the selection of optimal deep neural network hyperparameters using the Aquila Optimizer (AO) and incorporates GAN-based augmentation to improve class balance and model generalization.

Algorithm 1: DNN_AO – Hybrid Optimization Framework

Input:

- Chest X-ray image dataset $D = \{X_i, y_i\}$
- Pre-trained ResNet50 model
- GAN for minority class augmentation
- AO parameters: population size N , maximum iterations T

Output:

- Optimized DNN model parameters
- Classified disease labels for input images

Step 1: Data Preprocessing and Augmentation

1. Load dataset D from ChestX-ray14.
2. Perform basic augmentations: rotation, flipping, scaling, and contrast adjustment.
3. Train the Generative Adversarial Network (GAN) to synthesize new samples for minority disease classes.
4. Merge original and synthetic images to obtain a balanced dataset D' .

Step 2: Feature Extraction using ResNet50

5. Load pre-trained ResNet50 (ImageNet weights).
6. Remove the final fully connected layer.
7. Extract 2048-dimensional feature vectors from the penultimate layer for all images in D' .
8. Store the extracted features $F = \{f_1, f_2, \dots, f_n\}$.

Step 3: Initialization of the Aquila Optimizer (AO)

9. Initialize AO population with random DNN architectures, where each solution S_i represents:

$$S_i = [L_i, N_i, A_i, D_i]$$

where L_i = number of hidden layers, N_i = neurons per layer, A_i = activation function, D_i = dropout rate.

10. Evaluate the fitness function (classification accuracy) for each solution by training a DNN using the feature set F .

Step 4: Optimization Phase

11. For each iteration $t = 1$ to T :
 - Apply AO's exploration and exploitation strategies inspired by eagle hunting behavior:
 - Exploration: global search to identify promising DNN structures.
 - Exploitation: local refinement of the best-performing solutions.
 - Update positions of candidate solutions based on AO's mathematical update equations.
 - Recalculate the fitness for all updated solutions.
12. Record the best solution S_{best} with the highest classification accuracy.

Step 5: Training the Optimized DNN

13. Construct a DNN using parameters from S_{best} .
14. Train the network using the extracted feature set F .
15. Validate the model and record performance metrics: accuracy, precision, recall, and F1-score.

Step 6: Prediction and Classification

16. For each new chest X-ray image:
 - Extract features using ResNet50.
 - Classify the image using the optimized DNN.

- Output the predicted disease label with corresponding probability scores.

4. Result and Discussion

4.1 Model Performance Evaluation

The diagnostic performance of the proposed DNN_AO framework was comprehensively assessed using standard classification metrics, namely Accuracy, Precision, Recall (Sensitivity), F1-score, and Specificity. These quantitative measures collectively provide insight into the model's predictive capability, robustness, and generalization performance, particularly when dealing with imbalanced medical datasets such as chest X-rays.

Let TP, TN, FP, and FN denote True Positives, True Negatives, False Positives, and False Negatives, respectively. The metrics are mathematically defined as follows:

- **Accuracy**

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

Accuracy represents the overall correctness of the model by quantifying the proportion of correctly classified instances among all evaluated samples. Although widely used, accuracy alone may not adequately reflect performance in the presence of class imbalance.

- **Precision**

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

Precision measures the reliability of the model's positive predictions. It expresses how many of the samples predicted as positive are truly positive. A higher precision indicates fewer false alarms, which is particularly valuable in clinical diagnosis to prevent unnecessary interventions.

- **Recall (Sensitivity)**

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

Recall, or Sensitivity, evaluates the model's ability to correctly identify all actual positive cases. It measures the fraction of diseased cases accurately detected, which is critical in medical applications where missing a positive case (false negative) could have serious consequences.

- **F1-Score**

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The F1-score is the harmonic mean of Precision and Recall, providing a balanced metric that considers both false positives and false negatives. It is especially useful for imbalanced datasets, as it emphasizes the trade-off between the two measures rather than overall correctness alone.

- **Specificity**

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

Specificity quantifies the proportion of actual negatives that are correctly identified as non-diseased. High specificity indicates that the model effectively minimizes false-positive predictions, ensuring that healthy cases are not misclassified as diseased.

4.2 Comparative Analysis

To comprehensively evaluate the effectiveness of the proposed DNN_AO model, a comparative analysis was performed against three well-established benchmark models: a standard Convolutional Neural Network (CNN), Random Forest (RF), and K-Nearest Neighbors (KNN). This comparison was conducted using identical experimental settings and datasets to ensure a fair and unbiased evaluation. A concise description of the benchmark models is provided below to contextualize the comparative study:

- **Convolutional Neural Network (CNN):** A conventional deep learning architecture designed and trained from scratch on the Chest X-ray (CXR) dataset to extract spatial features and perform image classification.
- **Random Forest (RF):** An ensemble-based machine learning algorithm that constructs multiple decision trees and aggregates their outputs to enhance generalization and minimize overfitting.
- **K-Nearest Neighbors (KNN):** A non-parametric, instance-based learning algorithm that classifies input samples based on the majority label among their k nearest neighbors in the feature space.

Model	Accuracy	Precision	Recall	F1-Score
DNN_AO	0.89	0.89	0.89	1.00
CNN	0.87	0.89	0.87	1.00
Random Forest (RF)	0.86	0.78	0.80	1.00
K-Nearest Neighbors (KNN)	0.88	0.85	0.88	1.00

Table 1. The comparative performance results of all models, in terms of Accuracy, Precision, Recall, and F1-score, are summarized.

As presented in Table 1, the proposed DNN_AO model demonstrated superior overall performance, achieving the highest accuracy (0.89) while maintaining a balanced Precision and Recall (both 0.89). This balance indicates that the model is not only effective in detecting true positive cases but also consistent in minimizing false predictions.

The CNN model, while exhibiting reasonable performance, achieved slightly lower accuracy, reflecting its limited ability to generalize across all disease categories compared to the proposed architecture. The Random Forest (RF) classifier produced moderate accuracy but showed lower Precision and Recall, implying a higher tendency for false classifications and reduced sensitivity in detecting positive cases. The KNN algorithm performed competitively but still lagged behind the DNN_AO model in terms of both robustness and predictive consistency.

These findings highlight the advantage of the DNN_AO framework in effectively capturing complex, high-level feature representations from chest radiographs. The integration of deep feature extraction with optimization-based fine-tuning contributed to its enhanced discriminative capability. Consequently, DNN_AO achieved the most balanced and reliable classification performance, demonstrating its potential suitability for clinical-grade automated disease detection and decision-support systems.

4.3 Learning Behavior

The training behavior and generalization ability of the proposed DNN_AO model were further analysed using the learning curves for accuracy and loss across multiple training

epochs. These curves provide valuable insights into the model's convergence pattern, stability, and overall learning efficiency during optimization. Figure 2 illustrates the training and validation accuracy progression over epochs. The training accuracy (depicted in blue) consistently improves and reaches approximately 94%, reflecting the model's strong capability to fit the training data effectively. The validation accuracy (shown in orange) stabilizes at around 89.5%, suggesting that while there is a slight indication of overfitting, the model maintains good generalization performance on unseen data. This minor gap between the training and validation accuracies signifies a well-regularized model capable of learning discriminative representations without excessive fitting to the training samples. Similarly, Figure 3 presents the learning curves for training and validation loss. The training loss (blue curve) exhibits a steady downward trend throughout the epochs, demonstrating that the optimization process successfully minimizes classification errors. The validation loss (orange curve) decreases initially and then stabilizes, indicating that the model achieves convergence without experiencing significant fluctuations or divergence. Overall, the smooth convergence patterns of both accuracy and loss curves confirm the effective learning behavior of the DNN_AO model. The integration of ResNet50-based feature extraction, GAN-based augmentation, and Aquila Optimization-driven hyperparameter tuning contributes to this stable training process, ensuring high predictive performance and robust generalization across the multi-class pulmonary disease dataset.




Fig. 2 Training and validation accuracy curves of the proposed DNN_AO model




Fig. 3 Training and validation loss curves of the proposed DNN_AO model

4.4 Confusion Matrix Analysis

To further validate the diagnostic performance of the proposed DNN_AO framework, a confusion matrix and Receiver Operating Characteristic–Area Under the Curve (ROC–AUC) analysis were employed. The confusion matrix provides a detailed breakdown of model predictions by comparing the number of true positives, true negatives, false positives, and false negatives for each disease class. This enables a granular understanding of class-wise prediction strengths and potential misclassification patterns across the 14 pulmonary conditions. By visualizing the confusion matrix, it becomes possible to identify which diseases are frequently misclassified and to assess how well the model differentiates between similar pathological features

	0	2	4	6	8	10	12	14							
0	568	0	1	0	0	0	0	1	0	0	0	1	1	1	
2	0	584	0	3	1	0	0	2	0	0	3	0	1	0	1
4	0	0	598	0	0	0	0	4	0	2	3	0	4	3	1
6	0	1	0	567	0	2	0	2	1	3	0	1	4	9	0
8	0	1	1	0	603	1	1	1	5	4	1	2	5	5	1
10	0	2	0	1	0	577	0	1	0	1	1	4	4	7	0
12	0	1	3	0	2	1	565	3	4	3	7	5	4	4	0
14	0	0	2	4	0	1	3	534	2	8	1	9	18	6	0
16	0	1	3	1	9	1	6	8	589	14	7	11	10	12	7
18	1	3	2	0	7	0	4	13	384	83	14	8	13	6	5
20	0	3	7	5	0	3	3	15	23	284	91	6	14	12	3
22	3	6	2	4	3	5	3	23	20	11	144	45	24	18	7
24	1	5	2	9	7	14	8	46	28	30	26	403	84	30	3
26	0	11	7	7	8	12	17	33	44	36	59	39	682	63	15
28	1	13	15	9	7	11	30	50	72	58	108	63	64	62	55

Fig. 4. Confusion Matrix of the proposed DNN_AO model

	0	2	4	6	8	10	12	14						
0	568	0	1	0	0	0	0	1	0	0	0	1	1	1
2	0	584	0	3	1	0	0	2	0	0	3	0	1	0
4	0	0	598	0	0	0	0	4	0	2	3	0	4	3
6	0	1	0	567	0	2	0	2	1	3	0	1	4	9
8	0	1	1	0	603	1	1	1	5	4	1	2	5	5
10	0	2	0	1	0	577	0	1	0	1	1	4	4	7
12	0	1	3	0	2	1	565	3	4	3	7	5	4	4
14	0	0	2	4	0	1	3	534	2	8	1	9	18	6
16	0	1	3	1	9	1	6	8	589	14	7	11	10	12
18	1	3	2	0	7	0	4	13	384	83	14	8	13	6
20	0	3	7	5	0	3	3	15	23	284	91	6	14	12
22	3	6	2	4	3	5	3	23	20	11	144	45	24	18
24	1	5	2	9	7	14	8	46	28	30	26	403	84	30
26	0	11	7	7	8	12	17	33	44	36	59	39	682	63
28	1	13	15	9	7	11	30	50	72	58	108	63	64	62

Fig. 5. Confusion Matrix of the Random Forest classifier

	0	2	4	6	8	10	12	14				
0	568	0	0	0	0	0	2	0	0	0	0	1
2	0	585	0	0	0	0	2	0	1	1	3	2
4	0	0	599	0	0	0	3	0	3	0	0	8
6	0	0	0	568	0	0	3	0	1	1	4	2
8	0	0	0	0	604	1	0	1	3	2	4	3
10	0	1	0	0	0	575	0	1	0	1	1	4
12	0	0	1	0	0	0	564	4	4	4	10	2
14	0	1	0	0	0	0	1	543	1	6	9	9
16	0	0	0	0	1	0	2	3	610	28	10	3
18	0	0	0	0	0	0	1	2	6	544	14	10
20	0	1	2	0	0	0	0	1	9	23	558	2
22	0	4	2	0	0	0	0	4	8	6	12	504
24	0	0	1	2	2	2	1	30	16	16	23	304
26	0	0	1	1	3	1	0	23	27	26	37	22
28	0	8	0	0	0	2	1	15	69	77	142	47

Fig. 6. Confusion Matrix of the CNN model

		Confusion Matrix														
		0	2	4	6	8	10	12	14							
Actuals	0	568	0	0	0	0	0	0	1	1	0	1	0	1	1	
	2	0	585	0	0	0	0	0	0	1	0	2	1	0	6	
	4	0	0	597	0	0	0	0	1	1	0	0	2	1	8	
	6	0	0	0	567	0	1	0	0	0	0	0	1	1	13	
	8	0	0	0	0	603	1	0	0	3	1	0	3	1	10	
	10	0	1	0	1	0	575	0	1	0	1	1	3	3	6	
	12	0	0	0	0	0	2	564	0	3	1	6	0	4	12	
	14	0	1	0	0	0	1	1	535	0	3	2	7	3	28	
	16	0	0	0	0	4	0	0	2	635	5	0	1	6	20	
	18	0	0	0	1	3	0	1	0	19	515	5	8	3	26	
	20	0	0	0	0	0	0	0	1	14	20	504	3	6	40	
	22	0	1	0	0	0	1	1	0	2	3	5	505	8	36	
	24	0	0	0	0	1	10	3	24	12	9	10	214	1197	35	
	26	0	0	0	0	1	1	3	7	15	6	18	17	1646	66	
	28	0	0	0	0	3	9	1	2	21	11	28	19	1315	7354	
		Predictions														

Fig. 7. Confusion Matrix of the K-Nearest Neighbors (KNN) classifier

4.5 Model Performance Comparison

The bar chart presents a comparative evaluation of four models—DNN_AO, CNN, Random Forest, and K-Nearest Neighbors (KNN)—using key performance metrics, including Accuracy, Precision, Recall, F1-score, and Specificity. The results indicate that the proposed DNN_AO model consistently attains the highest values across all evaluation criteria, demonstrating its superior predictive performance and well-balanced learning capability. Overall, the DNN_AO framework distinguishes itself through its high accuracy and strong adaptability across varying conditions. Its excellent Recall and Specificity scores make it particularly effective for clinical applications, where accurately identifying pulmonary diseases is of critical importance. While the CNN and KNN models show competitive results, they achieve slightly lower Precision and Recall values, indicating a reduced ability to correctly detect positive cases. The Random Forest classifier, although effective for structured and tabular data, performs comparatively poorly in this high-dimensional image-based classification task, highlighting its limited capacity to model the complex spatial patterns inherent in chest radiographs.

In summary, the DNN_AO model clearly outperforms the benchmark approaches by maintaining high accuracy alongside robust generalization performance. Its strong Recall and Precision values underscore its reliability in identifying true disease cases while minimizing false predictions, making it a highly suitable candidate for real-world clinical

deployment where accurate and consistent pulmonary disease detection is essential for timely diagnosis and informed decision-making.




Fig. 8. Model Performance Comparison Chart

5. Conclusion

This work presents a robust and intelligent deep learning framework, termed DNN_AO, for the automated detection and classification of multiple pulmonary diseases using chest X-ray (CXR) images. The proposed approach combines three complementary components—GAN-based data augmentation, ResNet50-driven deep feature extraction, and a Deep Neural Network (DNN) optimized through the Aquila Optimizer (AO)—to establish an integrated and high-performance diagnostic pipeline. By employing GAN-generated synthetic samples, the framework effectively addresses class imbalance, a common challenge in medical imaging datasets where certain pathological conditions are sparsely represented. The ResNet50 backbone facilitates the extraction of informative and high-level spatial features, enabling the model to capture complex radiographic characteristics associated with diverse thoracic abnormalities. These extracted features are further optimized using the Aquila Optimizer, which dynamically adjusts essential DNN hyperparameters, including network depth, neuron configuration, activation functions, and dropout rates. This adaptive optimization strategy promotes stable convergence, enhances learning efficiency, and improves classification performance without the need for extensive manual tuning. Comprehensive experimental analysis reveals that the DNN_AO framework consistently surpasses conventional methods such as standard Convolutional Neural Networks (CNN), Random Forest (RF), and K-Nearest Neighbors (KNN) across key evaluation metrics, including Accuracy, Precision, and Recall. The results highlight the model's strong generalization capability and robustness in managing heterogeneous disease patterns, even in scenarios involving limited or imbalanced training data. In addition, the well-balanced Precision–Recall characteristics demonstrate reliable diagnostic behavior, effectively reducing both false-negative and false-positive predictions—an essential requirement for clinical decision support systems. Beyond its predictive performance, the DNN_AO framework exhibits scalability and flexibility, making it suitable for deployment in resource-constrained healthcare environments where access to expert radiologists and advanced computational resources may be limited. Its modular design further enables straightforward integration with additional diagnostic modalities, laying the groundwork for future multi-modal and intelligent healthcare solutions.

Future research will focus on extending the dataset to include a wider range of patient populations and additional pulmonary diseases, as well as incorporating multi-modal inputs such as computed tomography (CT) images and patient clinical records. Furthermore, the integration of Explainable Artificial Intelligence (XAI) techniques will be

emphasized to improve model interpretability and clinical transparency, thereby enhancing trust and adoption among medical practitioners.

Author Contributions

All authors contributed equally to the conception, design, and development of the research framework. Each author participated in data preprocessing, model implementation, experimental evaluation, and manuscript preparation. All authors have reviewed and approved the final version of the manuscript for submission.

Funding

The authors declare that no financial support, grants, or funding was received from any organization for the research, authorship, or publication of this study.

Data Availability Statement

Data sharing is not applicable to this study, as no new datasets were generated or analyzed during the current research. The experiments were conducted using the publicly available ChestX-ray14 dataset, which can be accessed through the National Institutes of Health (NIH) repository.

Declarations

Conflict of Interest:

The authors declare that there are no conflicts of interest regarding the publication of this article.

Consent to Participate:

Not applicable, as the study utilizes open-access datasets that do not involve direct human participation or identifiable personal data.

Consent to Publish:

All authors have provided their full consent for the publication of this manuscript and agree with the content and conclusions presented herein.

References:

- [1] Ozturk, T., Talo, M., Yildirim, E.A., Baloglu, U.B., Yildirim, O., Acharya, U.R.: Automated detection of COVID-19 cases using deep neural networks with X-ray images. *Computers in Biology and Medicine* 121, 103792 (2020). <https://doi.org/10.1016/j.compbio-med.2020.103792>
- [2] Alshazly, H., Linse, C., Abdalla, M., Sayed, A., et al.: Explainable COVID-19 detection using chest CT and X-ray images with deep learning. *Computers in Biology and Medicine* 132, 104336 (2021). <https://doi.org/10.1016/j.compbio-med.2021.104336>
- [3] Chandra, R., Jain, S.: Pneumonia detection on chest X-ray using machine learning paradigm. *Materials Today: Proceedings* 61, 284–288 (2022). <https://doi.org/10.1016/j.matpr.2022.05.223>
- [4] Sahlol, A.T., Abdeldaim, A.M., Hassanien, A.E.: VDSNet: A deep learning architecture to improve lung disease detection and classification. *Multimedia Tools and Applications* 81, 19133–19152 (2022). <https://doi.org/10.1007/s11042-021-11444-1>
- [5] Pramanik, S., Bhattacharya, K., Bhattacharyya, R.: COVNet: A deep learning model for COVID-19 detection. *Applied Soft Computing* 108, 107433 (2021). <https://doi.org/10.1016/j.asoc.2021.107433>

- [6] Khan, M.A., Tariq, U., Zhang, Y.D., Akram, T., Sharif, M.: A study of deep learning-based COVID-19 classification using chest X-ray images. *Neural Computing and Applications* 33, 13505–13521 (2021). <https://doi.org/10.1007/s00521-021-06074-6>
- [7] Apostolopoulos, I.D., Mpesiana, T.A.: COVID-19: automatic detection from X-ray images using transfer learning with convolutional neural networks. *Physical and Engineering Sciences in Medicine* 43, 635–640 (2020). <https://doi.org/10.1007/s13246-020-00865-4>
- [8] Khan, A.I., Shah, J.L., Bhat, M.M.: Coronet: A deep neural network for detection and diagnosis of COVID-19 from chest X-ray images. *Computer Methods and Programs in Biomedicine* 196, 105581 (2020). <https://doi.org/10.1016/j.cmpb.2020.105581>
- [9] Hemdan, E.E.D., Shouman, M.A., Karar, M.E.: COVIDX-Net: A framework of deep learning classifiers to diagnose COVID-19 in X-ray images. *arXiv preprint arXiv:2003.11055* (2020)
- [10] Al-Waisy, A.S., Al-Fahdawi, S., Abdulkareem, K.H., Mohammed, M.A., et al.: COVID-CheXNet: hybrid deep learning framework for identifying COVID-19 virus in chest X-ray images. *Soft Computing* 25, 11431–11450 (2021). <https://doi.org/10.1007/s00500-021-05896-0>
- [11] El-Shafai, W., Abd El-Samie, F.E., Khairy, A.: Detection of COVID-19 in chest X-ray images using FFNN and transfer learning. *Multimedia Tools and Applications* 80, 19921–19937 (2021). <https://doi.org/10.1007/s11042-021-10589-z>
- [12] Mahmud, T., Rahman, M., Rahman, M.F.: CovXNet: A multi-dilation convolutional neural network for automatic COVID-19 and other pneumonia detection from chest X-ray images. *Applied Sciences* 10(11), 3725 (2020). <https://doi.org/10.3390/app10113725>
- [13] Ghoshal, B., Opuencia, M.J.C.: Estimating uncertainty and interpretability in deep learning for coronavirus (COVID-19) detection. *arXiv preprint arXiv:2003.10769* (2020)
- [14] Mukherjee, H., Ghosh, S., Dhar, K.: COVID-19 detection using EfficientNetV2L with transfer learning and data augmentation. In: *Proceedings of the International Conference on Smart Technologies in Sustainable Energy (ICSTSE)*, pp. 1–5 (2021). <https://doi.org/10.1109/ICSTSE52398.2021.9780352>
- [15] Islam, M.Z., Islam, M.M., Asraf, A.: A combined deep CNN-LSTM network for the detection of novel coronavirus (COVID-19) using X-ray images. *Informatics in Medicine Unlocked* 20, 100412 (2020). <https://doi.org/10.1016/j.imu.2020.100412>
- [16] Rahman, A., Monira, A.M., Akhter, S.: Transfer learning with CNNs for pneumonia detection in chest X-rays. In: *Proceedings of the IEEE International Conference on Computer and Communication Engineering (ICCCE)*, pp. 265–270 (2020). <https://doi.org/10.1109/ICCCE50164.2020.9239023>
- [17] Wang, L., Wong, A.: COVID-Net: A tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images. *Scientific Reports* 10, 19549 (2020). <https://doi.org/10.1038/s41598-020-76550-z>
- [18] Narin, A., Kaya, C., Pamuk, Z.: Automatic detection of coronavirus disease (COVID-19) using X-ray images and deep convolutional neural networks. *Pattern Analysis and Applications* 24, 1207–1220 (2021). <https://doi.org/10.1007/s10044-021-00984-y>
- [19] Tahir, A., Soares, F., do Nascimento, E., Ferreira, A., et al.: COVID-19 detection using convolutional neural networks. *Computers in Biology and Medicine* 132, 104004 (2021). <https://doi.org/10.1016/j.compbimed.2021.104004>
- [20] Abbas, A., Abdelsamea, M.M., Gaber, M.M.: Classification of COVID-19 in chest X-ray images using DeTraC deep convolutional neural network. *Applied Intelligence* 51, 854–864 (2021). <https://doi.org/10.1007/s10489-020-01829-7>
- [21] Panwar, H., Gupta, P.K., Siddiqui, M.K., Morales-Menendez, R., Bhardwaj, P., Singh, V.: A deep learning and grad-CAM based color visualization approach for fast detection of

- COVID-19 cases using chest X-ray and CT-scan images. *Chaos, Solitons & Fractals* 140, 110190 (2020). <https://doi.org/10.1016/j.chaos.2020.110190>
- [22] Chowdhury, M.E.H., Rahman, T., Khandakar, A., Mazhar, R., et al.: Can AI help in screening viral and COVID-19 pneumonia? *IEEE Access* 8, 132665–132676 (2020). <https://doi.org/10.1109/ACCESS.2020.3010287>
- [23] Wang, L., Qiao, Z., Yao, A.: PneuNet: Attention-based deep neural network for pneumonia detection from chest radiographs. *Computers in Biology and Medicine* 135, 104537 (2021). <https://doi.org/10.1016/j.combiomed.2021.104537>
- [24] Ucar, F., Korkmaz, D.: COVIDiagnosis-Net: Deep Bayes-SqueezeNet based diagnosis of the coronavirus disease 2019 (COVID-19) from X-ray images. *Medical Hypotheses* 140, 109761 (2020). <https://doi.org/10.1016/j.mehy.2020.109761>
- [25] Li, L., Qin, L., Xu, Z., Yin, Y., et al.: Artificial intelligence distinguishes COVID-19 from community-acquired pneumonia on chest CT. *Radiology* 296, E65–E71 (2020). <https://doi.org/10.1148/radiol.2020200905>