Genetic Algorithm Based Optimization of Convolutional Neural Network for Respiratory Disease Detection

Vishwachetan D [a], Nandini S B [b] Pranjal Shrivastava and Nihal Jahagirdar [c]
Department of Computer Science and Engineering, Ramaiah Institute of Technology, Bangalore, India

Keywords: X-Ray Imaging, AI, Xception Neural Network, Pulmonary Disease, Genetic Algorithm.

Abstract: The pandemic Covid 19 in the year 2019 highlighted the need for advanced diagnostic methodologies to address a spectrum of pulmonary diseases. Although the major method of COVID-19 detection is still conventional PCR testing, the combination of AI and X-ray imaging presents a promising path toward a thorough diagnosis of pulmonary illness. Here, we provide a new optimization framework based on the Xception neural network architecture and Genetic Algorithm (GA) for precise pulmonary disease detection from X-ray pictures, including coronavirus and pneumonis (viral, bacterial). By utilising deep learning and convolutional neural networks, the main aim of this paper to improve the accuracy and efficiency of diagnosis. Using GA, we explore the vast design space of deep CNN architectures, encompassing parameters such as network depth, layer count, and type. Utilising an extensive dataset of X-ray pictures, the suggested Xception-based neural network is rigorously assessed repeatedly through GA-driven optimization. The result highlight how well the improved model distinguishes lung disorders achieved with AI-driven approaches.

1 INTRODUCTION

The technology based on deep learning algorithms has transformed traditional medical image diagnosis and prognosis in recent years. X-ray imaging is of the most readily available and used method for diagnosing lung conditions like pneumonia, whether bacterial or viral, as well as emerging threats like COVID-19. The use of deep neural networks (DNNs) for automated classification and detection tasks has seen significant advancements, delivering remarkable accuracy across various benchmarks. Among the notable CNN architectures is Xception, introduced in 2017 by François Chollet, known for its exceptional image recognition capabilities. Its innovative depth-wise separable convolutions enhance effective feature learning and extraction, making it a strong candidate for medical image analysis where accuracy and computational efficiency are paramount. However, fine-tuning Xception's architecture and parameters to meet specific medical imaging requirements remains a challenge. This study suggests a way to enhance the detection of pulmonary diseases, including COVID-19, viral, and bacterial pneumonia, in X-ray images. The approach involves a hybrid methodology that combines the optimization capabilities of genetic algorithms (GAs) with Xception's properties. Genetic algorithms, inspired by natural selection, serve as a powerful tool to identify and build optimal DNN configurations. This research seeks to determine whether using genetic algorithms to refine the Xception model's architecture and hyperparameters—specifically for lung disease detection from X-ray images—is effective. The objective is to improve the precision and robustness of disease diagnosis by iteratively adjusting neural network architectures using genetic algorithms, facilitating early identification and timely diagnosis or treatment. The optimized Xception-based DNN model will be evaluated on benchmark datasets through extensive experiments, including cases of COVID-19, viral pneumonia, bacterial pneumonia,

---

[a] https://orcid.org/0009-0009-9676-8279
[b] https://orcid.org/0009-0001-5708-2954
[c]

D, V., S B, N., Shrivastava, P. and Jahagirdar, N.
Genetic Algorithm Based Optimization of Convolutional Neural Network for Respiratory Disease Detection.
DOI: 10.5220/0013733200004664

In *Proceedings of the 3rd International Conference on Futuristic Technology (INCOFT 2025) - Volume 3*, pages 804-810
ISBN: 978-989-758-763-4

and healthy controls. The aim of this study is to advance the creation of reliable and efficient automated pulmonary disease diagnosis tools, supporting clinical decisions by healthcare professionals and ultimately improving patient outcomes.

2 LITERATURE REVIEW

Over the years, several techniques have been suggested for medical image analysis, shifting from traditional feature-based approaches to advanced machine learning techniques. Early investigations in medical image analysis focused on basic image processing techniques like thresholding, morphological operations, and edge detection. While these methods laid the groundwork for future research, they often fell short in meeting the accuracy and consistency required in clinical practice. Apostolopoulos and Mpesiana (Apostolopoulos, Mpesiana, et al. 2020) fine-tuned Convolutional Neural Networks (CNNs) for the automatic detection of COVID-19 from X-ray images, demonstrating that pre-trained models can enhance diagnostic accuracy. Similarly, Duran-Lopez et al. (Duran-Lopez, ominguez-Morales, et al. 2020) proposed COVID-XNet, a deep learning model designed to diagnose and localize COVID-19 in chest X-rays, aiming to improve both detection accuracy and efficiency. Sethy and Behera (Sethy, Behera, et al. 2020) investigated deep learning potential in medical imaging, using neural networks to extract features for the identification of COVID-19 in X- rays images. This study highlighted how deep learning techniques can streamline the analysis of X-ray data, offering a reliable solution for disease diagnosis and reducing unnecessary examinations.Narin et al. (Narin, Kaya, et al. 2021) applied various deep learning models for classification of coronavirus and normal cases, with their ResNet50 model achieving 98.0% accuracy in the best-case scenario. Zhang et al. (Zhang, Xie, et al. 2020) introduced another ResNet-based model that achieved an AUC of 0.952, effectively highlighting areas affected by pneumoniavari using Grad-CAM. Wang et al. (Wang, Lin, et al. 2020) proposed a deep CNN for classifying viral and bacterial infections and normal cases, achieving 83.5% accuracy.Image segmentation has also played a critical role in COVID-19 applications, including diagnostics (Chen, et al. 2019), (Wang, et al. 2021), (Jin, et al. 2020), (Song, et al. 2021). For example, Li et al. (Li, et al. 2020) used a U-Net architecture to segmentin lung images to differentiate COVID-19 from pneumonia acquired from the community using CT scans of the chest region. Jin et al. (Jin, et al. 2020) developed an AI system for rapid COVID-19 detection, where segmented CT slices serve as input for the classification model.Segmentation techniques also prove valuable in quantification tasks within medical applications (Jin, et al. 2020), (Shan, et al. 2021). A new model, XcepCovidNet, was introduced to identify features in X-rays of the chest region, utilizing transfer learning combined with hyperparameter tuning to address limitations in the training dataset (Juneja, Kumar, et al. 2024). Beyond X-rays, recent studies have turned to CT scans for the same purpose. For instance, Khan et al. (Khan, Shah, et al. 2020) developed CoroNet, a classification system consisting of four classes for COVID-19, achieving accuracies of 89.6% and 95% for chest X-ray (CXR) and CT scanned images, respectively. COVNet, designed by Li et al. (Li, et al. 2020), was based on ResNet50 and trained on a dataset of 4,356 images of CT scans of the chest region. Lastly, Joloudari et al. (Joloudariet, et al. 2023) proposed a deep learning-based global feature extractor for COVID-19 detection, further contributing to the research on using deep learning in medical image analysis.

3 PROPOSED METHOD

The suggested model in this paper makes use of three fundamental algorithms. The deep convolutional neural network Xception model is used for the detection of respiratory disease in lung X-ray images. The Genetic algorithm is then used to tune the hyperparameters of the Xception model to achieve the best possible architecture. This section explains the Genetic Algorithm, CNN, Xception model and finally the model suggested.

3.1 CNN

CNNs are a class of deep learning models for processing structured grid data, notably images and videos. Convolutional, pooling, and fully connected layers are how CNNs work to extract hierarchical characteristics from input data.Components of a CNN: Convolutional Layers: These layers are made up of filters, sometimes known as kernels, that execute convolutions by sliding across the input data. Each filter specializes in detecting specific features, such as edges or textures, by capturing spatial correlations. Activation functions such as ReLU introduce non-linearity. Pooling Layers:

Convolutional layer feature maps are down sampled by pooling layers, which reduces spatial dimensions without losing important information. For example, max pooling selects the greatest value within local regions, effectively shrinking feature map sizes. Feature maps are converted into vectors and run through one or more fully connected layers following a number of convolutional and pooling layers.The layers here handle classification or regression tasks by learning intricate relationships between extracted features and target labels. CNN architectures may incorporate additional elements such as dropout layers for regularization, batch normalization layers for accelerated convergence, and skip connections for improved gradient flow during training. Training CNNs involves optimizing parameters (weights and biases) using gradient-based optimization algorithms like SGD or its variants. During training, the network minimizes a loss function, quantifying the disparity between predicted outputs and ground truth labels.

3.2 Xception

Xception is a CNN architecture innovated by François Chollet, renowned for his contribution to the Keras deep learning library. Termed as "Extreme Inception," Xception builds upon the foundational concepts of the Inception architecture while introducing notable advancements. Central to Xception's design is the utilization of depth wise separable convolutions, a variant of conventional convolutional operations. This methodology effectively segregates spatial and channel-wise convolutions into distinct processes, resulting in a significant reduction in both parameters and computational complexity compared to conventional convolutions. Consequently, Xception achieves enhanced efficiency and model lightweightness. The architecture of Xception heavily draws from the Inception modules featured in the Inception v3 model. However, Xception distinguishes itself by replacing conventional convolutions within these modules with depth wise separable convolutions. This architectural refinement facilitates an improved utilization of computational resources, ensuring the modeling of intricate patterns and relationships across various scales. One notable advantage of Xception lies in its capability to capture both local and global dependencies within input data. The decomposition of the convolution operation into spatial and channel-wise components allows Xception to effectively model complex structures and correlations present in the data. Furthermore, Xception's architecture boasts expedited training and inference times, surpassing preceding CNN architectures. This attribute renders Xception particularly suitable for applications characterized by resource-constrained environments, where computational efficiency is paramount.

3.3 Genetic Algorithm

Genetic Algorithms (GAs) are commonly applied to optimization problems, such as tuning hyperparameters in machine learning. A population of potential solutions, known as people or chromosomes, is used by a genetic algorithm to solve problems. Each chromosome is a potential solution. In the context of hyperparameter tuning, these solutions usually correspond to different sets of hyperparameters for a machine learning model. As the algorithm runs through a sequence of stages called generations, selection, crossover, and mutation are applied to create a fresh set of potential solutions.

3.3.1 Initialization

The algorithm begins by generating an initial population of chromosomes, typically done throughrandomness or certain heuristics.

3.3.2 Evaluation

Each chromosome is assessed according to its fitness, which measures how effectively the solution performs the given optimization task. In hyperparameter tuning, the fitness is determined by how well the machine learning model performs when trained with the hyperparameters encoded in the chromosome.

3.3.3 Selection

Greater fitness values in chromosomes are preferable for reproduction, similar to the concept of "survival of the fittest."

3.3.4 Crossover

Selected chromosomes are paired to produce offspring through crossover or recombination. This process involves mixing genetic information from parent chromosomes to create new solutions. In hyperparameter tuning, crossover allows for the exploration of different hyperparameter combinations.

In hyperparameter tuning, where finding the ideal set of hyperparametersfor a machine learning model is the goal, GAs provide an effective way to navigate the large search space. Initially, a population of potential solutions (chromosomes) is created either randomly or using heuristics. Each chromosome encodes some hyperparameters for the machine learning model, which might include factors like learning rates, regularization strengths, or network architectures.Next, the fitness of each chromosome is assessed by training and testing the machine learning model using the hyperparameters it contains. Performance metrics, such as accuracy or loss, are used to assess fitness. Selection methods like tournament or roulette wheel selection are then applied to choose chromosomes to be reproduced based on their fitness, favoring those with greater values.Crossover and mutation are triedon these selected chromosomes to generate new offspring. Crossover mixes genetic information from parent chromosomes, enabling the exploration of new hyperparameter combinations. Mutation introduces random changes, maintaining diversity in the population and preventing early convergence to suboptimal solutions. The offspring replace the previous generation, with fitter individuals more likely to survive.Until a termination requirement is satisfied, for example, by reaching a certain number of generations or attaining adequate performance, this iterative process keeps going. Through this cycle of selection, crossover, and mutation, genetic algorithms efficiently search the hyperparameter space, slowly converging toward optimal or near-optimal configurations that enhance the ML model's performance.

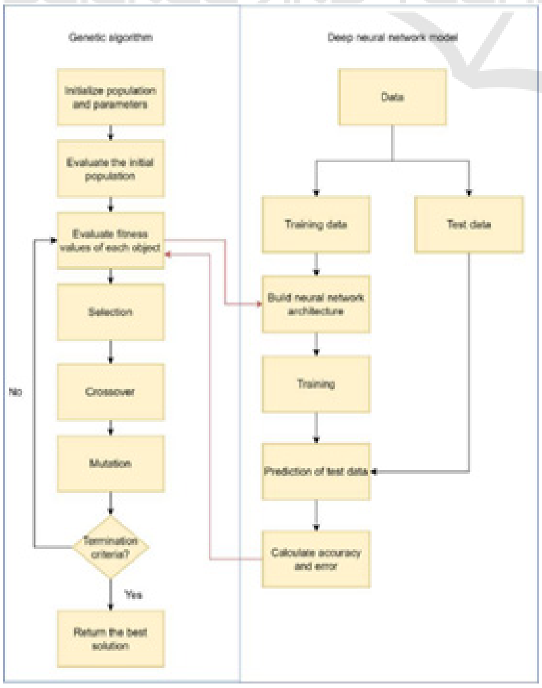

Figure 1: Flowchart of proposed model

This research employs Genetic Algorithm (GA) in conjunction with the Xception architecture, a sophisticated Convolutional Neural Network (CNN). First, a deep Xception network is constructed utilizing parameters computed from the Genetic Algorithm. Subsequently, these deep neural networks undergo training and evaluation using a dataset aimed at discerning COVID-19 presence in individuals. Each network's performance is assessed based on its error rate, with lower error rates indicative of more desirable solutions. These evaluated solutions undergo further refinement via iterations of the GA algorithm. With each iteration, novel networks are generated, leading to progressively improved outcomes. The rationale behind selecting the Xception model stems from its proven efficacy in prior research endeavours, consistently yielding commendable results. Its selection is particularly apt given its tailored focus on COVID-19 detection, aligning closely with the objectives of this study. Figure 1 explains the flowof the proposed model.

4 IMPLEMENTATION

4.1 Dataset for Respiratory Disease Detection Training

The dataset for respiratory disease detection training encompasses four primary classes: COVID, normal lung conditions, bacterial pneumonia, and viral pneumonia. Initially split into training, testing, and validation sets, the dataset underwent meticulous cleaning due to the presence of noise, including random letters and unnecessary watermarks, ensuring data integrity and reliability. With close to 1400 images solely for training purposes, augmentation techniques were employed to expand the dataset size, enhancing the model's ability to generalize and learn diverse patterns. This comprehensive dataset, meticulously curated and augmented, serves as a robust foundation for training and evaluating deep learning models aimed at accurate and effective respiratory disease detection and classification.

4.2 Model Implementation

For implementing the proposed model, after initializing the population, the evaluate_population function is called to train and evaluate each candidate solution (CNN) on the training and validation datasets. This step involves training the CNN model with the training data, then evaluating its execution using the validation data. The fitness scores are computed based on the performance metrics obtained during evaluation. Within the loop that iterates through generations, the training of CNN models is implicitly done during the evaluation step, as part of the evaluate_population function. The CNN models are trained with the training data before their performance is evaluated on the validation data. The training process typically involves using the training data on the CNN model, computing the loss, and then modifying the model's parameters. Algorithms for optimization such as SGD is used. The performance of each CNN model on the validation dataset is then evaluated using metrics like accuracy, loss, or other relevant measures, and these evaluations are used to compute the fitness scores. The process continues until a termination condition is reached, such as reaching a set number of generations.

5 RESULTS

Accuracy and categorical cross entropy loss were theperformance metrics used to find thebest architecture discovered by Genetic algorithm. One frequent metric used to assess a classification model's performance is accuracy. Out of all the anticipated classifications, it calculates the percentage of accurate classifications. Accuracy is mathematically calculated as shown in Figure 2.

The categorical cross-entropy loss function given in Equation 1serves as a pivotal evaluation metric in various machine learning tasks, particularly in classification problems where the output is represented in a categorical format. This metric quantifies the disparity between the true distribution of class labels and the predicted probabilities assigned by the model. By computing the logarithmic difference between the predicted probabilities and the actual class labels across all categories, the categorical cross-entropy loss penalises deviations from the true distribution, effectively guiding the model towards better classification performance. Its formulation makes it particularly suited for multi-class classification tasks, providing a continuous, differentiable measure of the model's performance that can be optimised through gradient descent methods.

$$\text{Loss} = -\sum_{i=1}^{n} y_i \cdot log\ \hat{y}_i \tag{1}$$

where, loss is the categorical cross-entropy loss, n is the output size, y is the correct probability distribution of class labels (one-hot encoded) and y^ is the estimated probability distribution of class labels.

The accuracy metric demonstrates a notable enhancement, with a consistent increase of 2-3%, indicating improved model performance in correctly classifying data points.

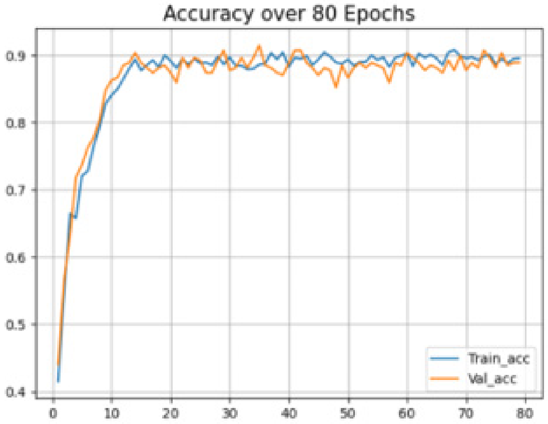

Figure 2: Accuracy obtained over 80 Epochs

The cross-entropy loss in Figure 3 exhibits a significant improvement, with a remarkable decrease of 30%, reflecting the model's enhanced ability to minimize discrepancies between predicted and true class probabilities. These advancements, coupled with the absence of significant spikes in loss, suggest the efficacy of hyperparameter tuning and the refinement of the dataset, contributing to a more stable and robust model performance with reduced noise interference.

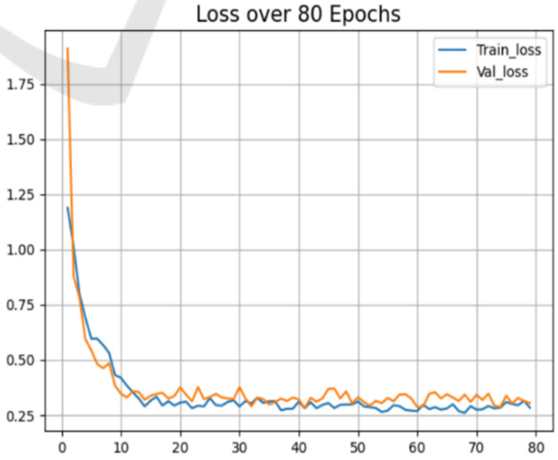

Figure 3: Categorical Cross Entropy Loss over 80 Epochs

The model trained was then tested across various images of lungs from the dataset to evaluate its performance. The accuracy on the test dataset was found out to be 0.79 and after optimization using GA, an accuracy of 0.85 was obtained. Categorical cross entropy loss of 0.79 was also reduced to 0.44 upon optimization by GA.

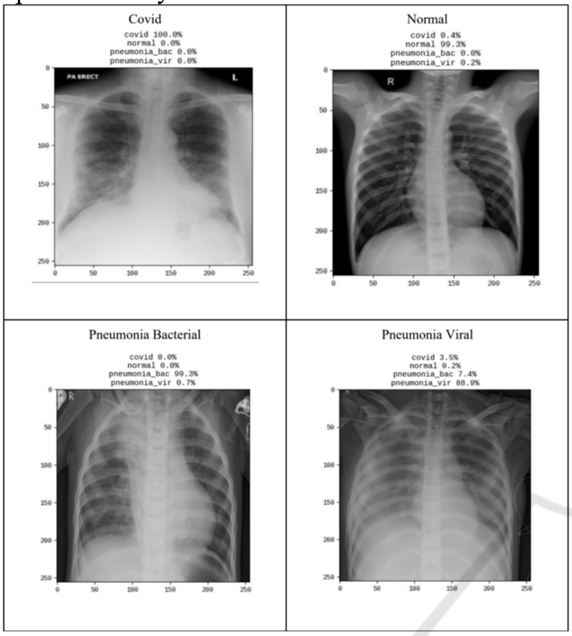

Figure 4: Predictions obtained for various classes

Figure 4 shows the results of classification along with their predicted conditions obtained from the trained model. The confusion matrices shown in Figure 5 and Figure 6 depict the accuracy of predictions of the model before and after optimizations using the Genetic Algorithm.

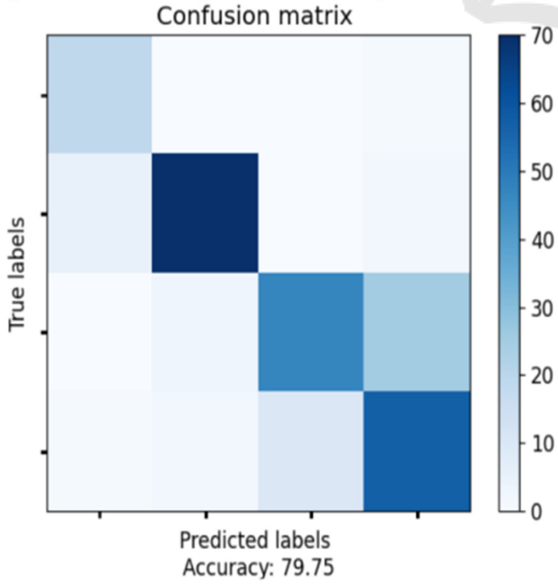

Figure 5: Confusion Matrix on test dataset before GA optimization

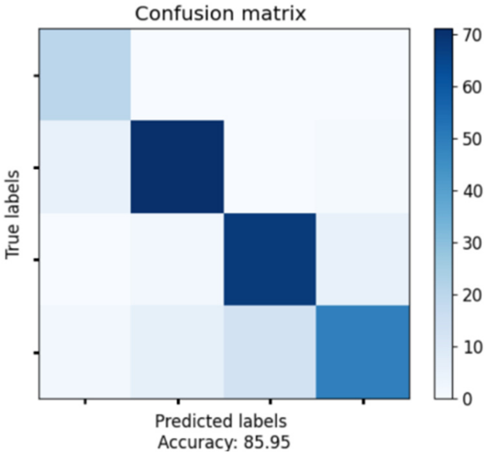

Figure 6 Confusion Matrix on test dataset after GA optimization

6 CONCLUSIONS

In conclusion, the enhancements made to the model have yielded significantly improved stability and robustness, as evidenced by the obtained results. Notably, there is a consistent enhancement in accuracy ranging between 2-3%, indicating the effectiveness of the implemented changes. Equally significant is the remarkable 30% decrease in Validation Loss, underscoring the model's improved generalisation capability. It's important to note that while training loss measures the performance of the model during the training phase, validation loss provides insight into how well the model generalises to unseen data, making it a crucial metric in assessing real-world performance. The fact that both training and validation accuracy and losses closely match underscores the absence of overfitting, signifying that the model has learned to generalise well to unseen data. This alignment between training and validation metrics further validates the reliability and efficacy of the model's performance. Overall, these results affirm the success of the enhancements implemented, paving the way for more reliable and accurate predictions in practical applications.

7 FUTURE WORK

For future work, a larger dataset should be assembled to encompass a more diverse set of lung diseases, potentially enhancing the model's ability to detect and classify a broader range of conditions. Additionally, fine-tuning the existing architecture with progressive techniques such as transfer learning or ensemble methods could potentially elevate the model's performance to even greater heights.

REFERENCES

D. Apostolopoulos and T. A. Mpesiana, “Covid-19: automatic detection from X-ray images utilizing transfer learning with convolutional neural networks,” Physical and Engineering Sciences in Medicine, vol. 43, no. 2, pp. 635–640, Apr. 2020, doi: 10.1007/s13246-020-00865-4. Available: https://doi.org/10.1007/s13246-020-00865-4

L. Duran-Lopez, J. P. Dominguez-Morales, J. Corral-Jaime, S. Vicente-Diaz, and A. Linares-Barranco, “COVID-XNet: A Custom Deep Learning System to Diagnose and Locate COVID-19 in Chest X-ray Images,” Applied Sciences, vol. 10, no. 16, p. 5683, Aug. 2020, doi: 10.3390/app10165683. Available: https://doi.org/10.3390/app10165683

P. K. Sethy, S. K. Behera, P. K. Ratha, and P. Biswas, “Detection of coronavirus Disease (COVID-19) based on Deep Features and Support Vector Machine,” International Journal of Mathematical Engineering and Management Sciences, vol. 5, no. 4, pp. 643–651, Aug. 2020, doi: 10.33889/ijmems.2020.5.4.052. Available: https://doi.org/10.33889/ijmems.2020.5.4.052

A. Narin, C. Kaya, and Z. Pamuk, “Automatic detection of coronavirus disease (COVID-19) using X-ray images and deep convolutional neural networks,” Pattern Analysis and Applications, vol. 24, no. 3, pp. 1207–1220, May 2021, doi: 10.1007/s10044-021-00984-y. Available: https://doi.org/10.1007/s10044-021-00984-y

Zhang J, Xie Y, Li Y, Shen C, Xia Y. COVID-19 Screening on Chest X-ray Images Using Deep Learning based Anomaly Detection. arXiv; 2020.

L. Wang, Z. Q. Lin, and A. Wong, “COVID-Net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images,” Scientific Reports, vol. 10, no. 1, Nov. 2020, doi: 10.1038/s41598-020-76550-z. Available: https://doi.org/10.1038/s41598-020-76550-z

O. Gozes et al., “Rapid AI Development Cycle for the Coronavirus (COVID-19) Pandemic: Initial Results for Automated Detection & Patient Monitoring using Deep Learning CT Image Analysis,” arXiv (Cornell University), Jan. 2020, doi: 10.48550/arxiv.2003.05037. Available: https://arxiv.org/abs/2003.05037

L. Li et al., “Using Artificial Intelligence to Detect COVID-19 and Community-acquired Pneumonia Based on Pulmonary CT: Evaluation of the Diagnostic Accuracy,” Radiology, vol. 296, no. 2, pp. E65–E71, Aug. 2020, doi: 10.1148/radiol.2020200905. Available: https://doi.org/10.1148/radiol.2020200905

J. Chen et al., “Deep learning-based model for detecting 2019 novel coronavirus pneumonia on high-resolution computed tomography,” Scientific Reports, vol. 10, no. 1, Nov. 2020, doi: 10.1038/s41598-020-76282-0. Available: https://www.nature.com/articles/s41598-020-76282-0

B. Wang et al., “AI-assisted CT imaging analysis for COVID-19 screening: Building and deploying a medical AI system,” Applied Soft Computing, vol. 98, p. 106897, Jan. 2021, doi: 10.1016/j.asoc.2020.106897. Available: https://doi.org/10.1016/j.asoc.2020.106897

C. Jin et al., “Development and Evaluation of an AI System for COVID-19 Diagnosis,” medRxiv (Cold Spring Harbor Laboratory), Mar. 2020, doi: 10.1101/2020.03.20.20039834. Available: https://doi.org/10.1101/2020.03.20.20039834

Y. Song et al., “Deep Learning Enables Accurate Diagnosis of Novel Coronavirus (COVID-19) With CT Images,” IEEE/ACM Transactions on Computational Biology and Bioinformatics, vol. 18, no. 6, pp. 2775–2780, Nov. 2021, doi: 10.1109/tcbb.2021.3065361. Available: https://doi.org/10.1109/tcbb.2021.3065361

Y. Cao et al., “Longitudinal Assessment of COVID-19 Using a Deep Learning–based Quantitative CT Pipeline: Illustration of Two Cases,” Radiology Cardiothoracic Imaging, vol. 2, no. 2, p. e200082, Apr. 2020, doi: 10.1148/ryct.2020200082. Available: https://doi.org/10.1148/ryct.2020200082

F. Shan et al., “Abnormal lung quantification in chest CT images of COVID-19 patients with deep learning and its application to severity prediction,” Medical Physics, vol. 48, no. 4, pp. 1633–1645, Mar. 2021, doi: 10.1002/mp.14609. Available: https://doi.org/10.1002/mp.14609

A. Juneja, V. Kumar, M. Kaur, D. Singh, and H.-N. Lee, “XcepCovidNet: deep neural networks-based COVID-19 diagnosis,” Multimedia Tools and Applications, Jun. 2024, doi: 10.1007/s11042-024-19046-6. Available: https://doi.org/10.1007/s11042-024-19046-6

A. I. Khan, J. L. Shah, and M. M. Bhat, “CoroNet: A deep neural network for detection and diagnosis of COVID-19 from chest x-ray images,” Computer Methods and Programs in Biomedicine, vol. 196, p. 105581, Nov. 2020, doi: 10.1016/j.cmpb.2020.105581. Available: https://doi.org/10.1016/j.cmpb.2020.105581

L. Li et al., “Using Artificial Intelligence to Detect COVID-19 and Community-acquired Pneumonia Based on Pulmonary CT: Evaluation of the Diagnostic Accuracy,” Radiology, vol. 296, no. 2, pp. E65–E71, Aug. 2020, doi: 10.1148/radiol.2020200905. Available: https://doi.org/10.1148/radiol.2020200905

J. H. Joloudariet al., “Developing a Deep Neural Network model for COVID-19 diagnosis based on CT scan images,” Mathematical Biosciences & Engineering, vol. 20, no. 9, pp. 16236–16258, Jan. 2023, doi: 10.3934/mbe.2023725. Available: https://doi.org/10.3934/mbe.2023725