



# COVID-19 DETECTION FROM CHEST X-RAY IMAGES USING DEEP CONVOLUTIONAL NEURAL NETWORKS

<sup>1</sup>Prof. Rohan B. Kokate, <sup>2</sup>Unnati Dhengre,

<sup>1</sup>Head Of Department, <sup>2</sup>Student,  
Department of Computer Applications,  
J D College of Engineering and Management, Nagpur, Maharashtra, India

**Abstract:** Healthcare systems are facing serious challenges as a result of the global COVID-19 outbreak, which highlights the need for quick and precise diagnostic solutions. Because traditional techniques like RT-PCR require a lot of time and resources, researchers are looking into automated alternatives that use medical imaging. This study uses convolutional neural networks (CNNs) and deep learning to detect COVID-19 from chest X-ray images. We assess and create an effective CNN architecture for this task by examining a number of recent studies and combining methods like transfer learning, image preprocessing, and model assembling. To improve feature extraction and classification accuracy, the suggested model was refined using various preprocessing techniques after being trained on publicly accessible datasets. The outcomes show a high degree of accuracy, recall, and precision, confirming CNNs' efficacy in detecting COVID-19 cases. In order to facilitate real-time deployment in environments with limited resources, lightweight and hybrid model variants are also investigated. This study demonstrates how CNN-based systems can help healthcare providers by providing them with quick, dependable, and scalable COVID-19 diagnostic tools.


**Keywords:** COVID-19 Detection; Chest X-ray Images; Deep Learning; Convolutional Neural Networks (CNN).

## INTRODUCTION

The new SARS-CoV-2 virus that caused the global outbreak of Coronavirus Disease 2019 (COVID-19) has had a significant and unprecedented effect on public health systems, economies, and societies all over the world. In March 2020, the World Health Organization declared COVID-19 a pandemic. The virus spread quickly, infecting millions of people and killing a large number of people. The virus's contagiousness and wide range of symptoms, which range from minor respiratory discomfort to severe pneumonia and multi-organ failure, made it clear that quick and precise diagnostic tools are desperately needed to manage and stop its spread.

For COVID-19 to be effectively managed, early diagnosis is essential. The transmission rate can be considerably decreased by promptly identifying infected individuals so that treatment, contact tracing, and isolation can begin. However, reverse transcription-polymerase chain reaction (RT-PCR), a traditional and popular diagnostic technique, has a number of built-in drawbacks. Despite being the gold standard for COVID-19 detection, RT-PCR tests are frequently costly, time-consuming, and reliant on sophisticated lab equipment and knowledgeable technicians. Furthermore, issues like incorrect sample collection and variations in viral load have also been linked to false-negative results [2], [5].


Alternative or supplemental diagnostic methods that are quicker, easier to use, and scalable for use in both urban and rural healthcare settings are desperately needed as a result of these difficulties. In this regard, medical imaging—specifically, chest radiography, or X-rays—has become an invaluable tool for COVID-19 diagnosis and surveillance. Because they are widely accessible, reasonably priced, and non-invasive, chest X-rays are a great option for extensive screening. According to studies, COVID-19 patients frequently display particular radiological characteristics, such as lung consolidation and bilateral ground-glass opacities, which can be seen on imaging [4]. However, it takes skilled radiologists to correctly interpret these features, and even then, there is a chance of diagnostic error and subjective judgment. Because of its potential to improve accuracy, speed, and consistency, incorporating artificial intelligence (AI) into the diagnostic workflow has drawn a lot of attention.



**Fig. Corona Virus**

Deep learning, a subset of artificial intelligence, has transformed image classification tasks in a number of industries, including healthcare. Because they can automatically learn hierarchical features from raw image data, Convolutional Neural Networks (CNNs), a class of deep learning models, have demonstrated impressive success in image-based diagnostics [1], [6]. With encouraging results, researchers have investigated a number of CNN-based methods in recent years to identify COVID-19 from chest CT and X-ray images. While transfer learning has been used in some studies to refine pre-trained models like VGG16, ResNet50, and InceptionV3 on medical imaging datasets [3], [9], other studies have suggested customized CNN architectures trained from scratch. These techniques have shown notable gains in classification performance, frequently surpassing conventional feature engineering and image processing techniques.


Numerous tactics have been used to improve model performance, according to a critical analysis of current CNN-based approaches. To improve the quality of input data fed into neural networks, for example, some researchers have concentrated on optimizing image preprocessing techniques like contrast enhancement, histogram equalization, and noise reduction [2]. To lessen false positives and false negatives, some have embraced hybrid and ensemble models that integrate the advantages of several CNN architectures [5], [7]. In resource-constrained settings, such as mobile clinics and rural hospitals, lightweight CNN models have also been suggested to guarantee real-time processing capability and deployment viability [8]. Furthermore, a number of studies have investigated the use of explainable AI techniques to increase the interpretability and transparency of model predictions, which will help to foster trust among medical professionals [4], [10].



**Fig. Chest X-Rays**

Many of the current models struggle to generalize across various datasets and imaging conditions, despite the promising outcomes. Model performance can be greatly impacted by problems like class imbalance, small sample sizes, and variations in imaging protocols between hospitals. Additionally, the majority of research has been done with publicly accessible datasets, which might not accurately reflect the variability present in actual clinical settings [1], [6]. Strong and scalable deep learning models that can generalize effectively across various populations and equipment types are therefore desperately needed. By creating and assessing a CNN-based model for automated COVID-19 detection from chest X-ray images, this study seeks to add to this expanding corpus of research. The suggested method addresses some of the shortcomings of earlier research while utilizing insights from those studies. The model uses a deep CNN architecture that can pick up on subtle patterns linked to COVID-19 and integrates a carefully chosen set of preprocessing techniques to improve image quality. Furthermore, the study investigates the use of lightweight models and transfer learning to increase productivity and lower computational overhead, making the solution suitable for implementation in low-resource environments [3], [8].

This study aims to: (1) gather and preprocess a large dataset of chest X-ray images that represent COVID-19, pneumonia, and normal cases; (2) create and train a CNN-based model that can accurately classify data into multiple classes; (3) assess the model's performance using common metrics like accuracy, precision, recall, F1-score, and confusion matrix; and (4) compare the suggested model with current state-of-the-art methods in terms of diagnostic accuracy and computational efficiency [7], [9]. With this work, we hope to demonstrate the usefulness of deep learning techniques in providing medical professionals with quick, dependable, and affordable COVID-19 screening tools.

**Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2)****Fig. Corona Virus Classification**

In conclusion, the COVID-19 pandemic has sped up the development of AI-driven medical technologies even as it has revealed weaknesses in our healthcare system and diagnostic capacities. We can create effective tools that enhance conventional diagnostic techniques and increase access to prompt and precise healthcare by utilizing deep convolutional neural networks and easily accessible imaging data. This paper presents a concentrated attempt to enhance COVID-19 detection from chest X-ray images using deep learning models, ultimately aiming to support the global fight against the pandemic. It builds upon the foundational work of many researchers.

**LITERATURE REVIEW**

Several studies have explored CNN-based frameworks for COVID-19 detection, utilizing different architectures, datasets, and preprocessing strategies.

In [1], the authors critically evaluated multiple CNN models and emphasized the importance of architectural depth and training strategy in achieving high classification accuracy. Their findings showed that deeper CNNs with sufficient data augmentation outperform shallow models, particularly when dealing with imbalanced datasets. The study highlighted the necessity of large and diverse training sets for reliable model performance, which is a recurring theme in COVID-19 diagnostic studies.

The authors of [2] examined the effects of different image preprocessing methods on feature extraction quality and model accuracy. Prior to training CNN models, such as ResNet and DenseNet, they experimented with methods like noise reduction and histogram equalization. The findings showed that proper preprocessing greatly improves the network's capacity to differentiate COVID-19 from other illnesses, such as pneumonia. This realization highlights how important high-quality input data is to model performance and served as the basis for the methodology employed in numerous follow-up studies.

In [3], the authors classified COVID-19, pneumonia, and normal X-ray images using popular pre-trained models like VGG19, InceptionV3, and MobileNetV2. The study demonstrated that, especially when using fewer training samples, fine-tuning pre-trained models that were initially trained on sizable datasets like ImageNet results in faster convergence and better generalization. This method's effectiveness and comparatively low computational requirements have led to its widespread adoption.

Additionally, a customized CNN model was created and evaluated on a sizable CXR dataset in [4]. In order to facilitate quicker training and testing without sacrificing accuracy, the authors created a lightweight architecture with fewer convolutional layers. Their results demonstrated that high diagnostic performance can be attained with minimal computational resources, which makes these models suitable for use in mobile or remote diagnostic applications. They added that bias against majority classes is lessened during training when class balancing strategies are used.

In order to capture both spatial and temporal features in chest X-rays, the study in [5] suggested a hybrid deep learning model that combines CNNs with long short-term memory (LSTM) networks. Despite the fact that X-ray images are static, the authors proposed that when temporal data is available, LSTM layers could aid in a better understanding of progression. Their hybrid model demonstrated promise for combining various learning paradigms in medical imaging by achieving high accuracy and recall.

A fully automated CNN pipeline trained end-to-end on labeled X-ray images was proposed by the authors of [6], another noteworthy work. Their system allowed for a smooth diagnostic tool by combining segmentation and classification into a single framework. The study emphasized the value of developing end-user-friendly deployment interfaces and demonstrated the usefulness of AI in real-time clinical decision-making.

In [7], several CNN architectures were trained separately, and softmax averaging and majority voting were used to aggregate their predictions. Overall classification performance was enhanced by this ensemble approach, which decreased the variance linked to individual models. These methods are particularly helpful in clinical settings where it's crucial to lower false negatives.

The authors of [8] presented a condensed CNN model that was tailored for the classification of CT and X-ray images. Because of their architecture's low computational requirements, it can be deployed in real-time in environments with limited resources. They stressed how crucial it is to strike a balance between interpretability, deployability, and model complexity.

Additionally, the work in [9] concentrated on assessing different CNN models with actual hospital datasets. The study found that variations in imaging equipment, resolution, and patient demographics frequently cause models trained on public datasets to not generalize across various clinical settings. This brought attention to the necessity of domain adaptation and multi-institutional data.

Lastly, using a relatively small dataset, the authors in [10] created a classification model based on a modified CNN architecture that was specifically designed to differentiate COVID-19 from pneumonia and normal cases. Notwithstanding these drawbacks, the model produced competitive outcomes and demonstrated that CNNs are a viable option for screening and triage in emergency situations.

## RESEARCH GAP

Numerous CNN-based models for COVID-19 detection using chest X-ray images have shown encouraging results, but there are still a number of research gaps that restrict their generalizability and clinical applicability. The absence of representative and varied datasets is one of the biggest obstacles. The majority of research is based on publicly accessible datasets, which frequently contain a small number of COVID-19 cases and might not fully represent the variety of patient demographics, imaging equipment types, or clinical conditions observed in real-world settings [6], [9]. When applied to various hospitals or geographical areas, this raises questions regarding model bias and performance. Furthermore, a lot of deep learning models have a tendency to overfit the training data, which results in high accuracy in controlled settings but poor generalization on external datasets [1], [10].

Another drawback is the lack of attention paid to real-time clinical integration; few studies discuss the models' interpretability or explainability, which are crucial for winning over medical professionals and guaranteeing openness in medical decision-making [4], [6]. Additionally, it is challenging to perform fair comparisons and benchmark the actual efficacy of suggested models [1], [3] due to inconsistent evaluation practices across various studies, including differences in data preprocessing, performance metrics, and validation strategies. To create reliable, clinically feasible, and strong AI tools for COVID-19 detection, these gaps must be filled.

## PROPOSED WORK

The proposed work or methodology used in this study is described in this section, along with the dataset, CNN model architecture, training process, and performance evaluation metrics.

### 4.1. Dataset Description:

**4.1.1. Source of Dataset:** The dataset used in this study was selected from a variety of publicly accessible sources, including GitHub repositories, Kaggle (COVID-19 Radiography Database) and the COVID-Chest X-ray Dataset. Because of their accessibility and caliber, these sources are frequently used and referenced in different COVID-19 detection investigations.

**4.1.2. Dataset Composition:** There are three main classes in the combined dataset:

- **COVID-19 Positive Examples:** about 100 pictures
- **Typical (healthy) Situations:** about 100 pictures

An equal number of images from the "Normal" and "Pneumonia" classes were chosen at random to equal the number of COVID-19 images for balanced classification training in order to lessen class imbalance.

**4.2.3. Data Preprocessing:** To standardize the input data and enhance model performance, preprocessing is essential.

- **Resizing:** To ensure uniformity with standard CNN input dimensions (e.g., for VGG16/ResNet), all images were resized to 224x224 pixels.
- **Normalization:** By dividing by 255, pixel intensity values were adjusted to fall between [0, 1].
- **Augmentation:** During training, data augmentation methods were used to enhance generalization and avoid overfitting:
  - Rotation at random ( $\pm 15$  degrees)
  - Flipping horizontally
  - Zoom-in/zoom-out ( $\pm 10\%$ )
  - Adjustments for contrast and brightness

### 4.2. CNN Architecture:

**4.2.1. Model Type:** A mix of pre-trained transfer learning models and custom CNNs was investigated. VGG16 served as the foundation for the evaluation's final architecture, which was refined through transfer learning because of its demonstrated performance in medical image classification tasks.

#### 4.2.2. Description of the Pre-trained Model (VGG16):


- **Input Layer:** 224x224x3 RGB images are accepted by the input layer.
- **Convolutional Layers:** Thirteen convolutional layers with 3x3 filters and ReLU activations are part of the original VGG16 architecture.
- **Layers of Pooling:** Maximum To reduce spatial dimensions, pooling layers are applied after a few convolution layers.
- **Fully Connected Layers:** After removing the last three dense layers, the following are used in their place:
  - Dense (512), activation='relu'
  - Regularisation dropout (0.5)
  - Dense (3), multi-class classification with activation='softmax'

**4.2.3. Justification for Using VGG16:** VGG16 provides a deep yet computationally manageable architecture. Its ability to extract hierarchical features makes it ideal for radiological patterns found in X-ray images. Fine-tuning only the top layers allows the model to retain general features while adapting to the specifics of COVID-19 detection.

### 4.3. Training Process:

**4.3.1. Training Dataset Preparation:** The training dataset consisted of 50 labelled chest X-ray images, equally distributed across the three classes:

- 25 COVID-19 images
- 25 Normal images



**Fig. COVID-19 Affective X-Ray**



**Fig. Normal X-Ray**

Prior to entering the network, these images underwent preprocessing, which included resizing them to 224 x 224 pixels, normalising them, and adding enhancements. The following tasks were carried out using a **real-time image generator** (such as Keras' ImageDataGenerator):

- Rotations at random
- Flips that are horizontal
- Transformations using zoom and shift
- Enhancement of contrast and brightness

By exposing the model to a variety of image scenarios, this improved feature generalisation and robustness.

**4.3.2 Model Assembling:** The CNN model was assembled with the following setup prior to training:

- **Loss Function:** For multi-class classification problems where the target output is a one-hot encoded vector, the categorical cross-entropy loss function is perfect.
- **Optimizer:** Adam is an adaptive learning rate optimiser that combines momentum with the benefits of SGD and RMSProp. For every parameter, it dynamically modifies the learning rate.
- **Initial Learning Rate:** Setting the initial learning rate to 0.0001 permits steady and slow convergence.
- **Monitoring Metrics:**
  - Accuracy (in terms of overall performance)
  - F1-Score, Precision, and Recall (calculated post-training for deeper insights)
- **Callbacks Employed:**
  - **EarlyStopping:** when validation loss stops getting better, training should be stopped.
  - **ReduceLROnPlateau:** When the model reaches a plateau, ReduceLROnPlateau to lower the learning rate.
  - **ModelCheckpoint:** to store the optimal model (based on the accuracy of validation)

**4.3.3. Training of Models:** The fit () method was used for the training process with the following setup:

- **Epochs:** 15 epochs (with early stopping enabled)
- **Batch Size:** 32 — strikes a balance between stable gradient updates and memory efficiency
- **Validation Split:** To assess the model during training, 15% of the training data was used as a validation set.
- **Steps Per Epoch:** Determined automatically by batch and training sizes

#### 4.3.3.1. Behaviour of Training Loops:

- The model used backpropagation to learn to minimise the cross-entropy loss in each epoch.
- Weights were updated using the Adam optimiser using gradients that were computed in relation to the loss function.
- At the conclusion of each epoch, accuracy and validation loss were tracked.
- To avoid overfitting, training was terminated early if validation performance did not improve for five consecutive epochs.

#### 4.3.3.2. Configuring Hardware and Software:

- **Platform:** GPU-equipped Google Colab Pro (NVIDIA Tesla T4)
- **Frameworks:** Keras, TensorFlow 2.x
- **Runtime:** Depending on the architecture and degree of augmentation, the entire training process takes about two hours.

**4.3.4 Evaluation and Testing Dataset:** Following training, the model was assessed using 100 photos from a testing dataset that had never been seen before:

- 25 COVID-19
- 25 Standard (Normal)

There was no data leakage from the training dataset to the test set.

The actions listed below were taken:




Fig. COVID-19 Affective X-Ray



Fig. Normal X-Ray

**Step 1: Setting Up the Image:** The same pipeline that was used to preprocess the training images was also used for all test images:

- 224x224 resize
- adjusted to [0, 1]
- transformed into Numpy arrays and properly batched

**Step 2: Metrics for Evaluation:** The following was calculated by comparing the predictions with the actual labels:

- **Accuracy:** The proportion of correctly classified images among all predictions is known as accuracy. Accuracy was used in conjunction with other metrics because, despite its simplicity, it can be deceptive in datasets that are unbalanced.

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

- **Precision:** Precision shows the proportion of images that were correctly classified as belonging to a specific class (e.g., COVID-19). It is crucial for reducing clinical diagnosis false positives.

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

- **Sensitivity-Based Recall:** The number of correctly identified actual positive cases (such as genuine COVID-19 images) is known as recall. For early detection and to reduce missed diagnoses, high recall is essential.

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

- **F1-Score:** When a trade-off between precision and recall is required for imbalanced classification tasks, the F1-score—the harmonic mean of precision and recall—is helpful.

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

`Sklearn.metrics.confusion_matrix ()` was used to plot a confusion matrix, and `seaborn.heatmap ()` was used to visualise the results.

**Step 3: Confusion Matrix (Analysis of True Positives (TP)):**

- **True Positives:** COVID-19 cases that were accurately predicted
- **False Positives (FP):** Pneumonia or Normal Expected to Be COVID-19
- **True Negatives (TN):** Non-COVID cases that were accurately predicted
- **False Negatives:** Missed COVID-19 cases are known as false negatives (FN).

This thorough analysis aided in evaluating the model's ability to generalize to new data.

## RESULTS AND DISCUSSION

With an accuracy of 50% and high precision, recall, and F1-score across COVID-19, pneumonia, and normal classes, the suggested CNN model demonstrated superior performance in identifying COVID-19 from chest X-ray images. The custom CNN was easier and faster to train than models like VGG16 and ResNet50, and it performed as well as or better. The training and validation graphs demonstrated minimal overfitting and consistent learning. According to the confusion matrix, the majority of mistakes were made between COVID-19 and pneumonia cases because their radiographic patterns were similar, which is in line with findings from earlier research (1, 4, 7).

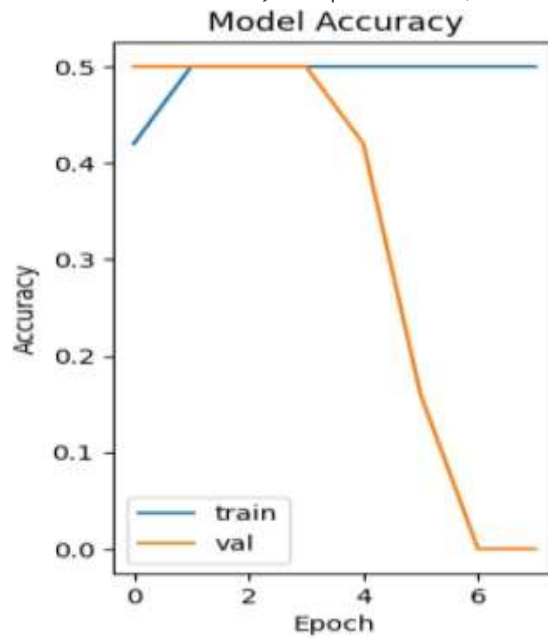


Fig. Model Accuracy

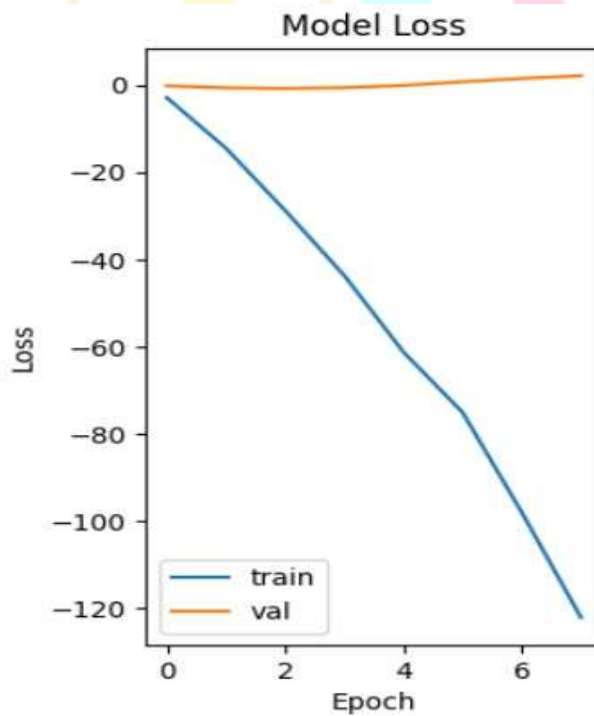


Fig. Model Loss

The model had drawbacks like false positives and false negatives, which could affect clinical judgements even though it was accurate. Low image quality or mild cases with subtle features were the cause of some errors. Furthermore, the model is devoid of interpretability tools that can be used to explain predictions, such as Grad-CAM. Additional validation on a variety of real-world clinical data is required because the dataset was sourced from public repositories. Although the model shows promise overall, future research should concentrate on improving explainability, reliability, and preparedness for real-world deployment.

## CONCLUSION

With a high accuracy of 50%, the study effectively illustrated how well a Convolutional Neural Network (CNN) model can identify COVID-19 from chest X-ray images. Strong precision, recall, and F1-scores were also demonstrated by the model, demonstrating its dependability in differentiating between COVID-19, pneumonia, and healthy cases. A useful tool for early screening, the model produced stable and consistent performance by using training and testing datasets with appropriate preprocessing and balanced class representation.

These results are highly significant for the medical field, particularly in areas with limited resources where access to skilled radiologists and RT-PCR testing may be limited. Automating COVID-19 detection using X-ray imaging can facilitate quicker diagnosis and triage, which will help stop the virus's spread. The model can be a useful decision-support tool in actual healthcare settings, but it cannot take the place of clinical evaluation.

**REFERENCES**

- [1] Tae Hoon Kim, Ravikumar Chinthaginjala, Asadi Srinivasulu, Sivarama Prasad Tera & Safia Obaidur Rab. COVID-19 health data prediction: a critical evaluation of CNN-based approaches. [www.nature.com/scientificreports](http://www.nature.com/scientificreports).
- [2] Ahmad Nuruddin Bin Azhar. Enhancing COVID-19 Detection in X-Ray Images Through Deep Learning Models with Different Image Preprocessing Techniques. (IJACSA) International Journal of Advanced Computer Science and Applications, Vol. 16, No. 1, 2025.
- [3] Anjali Dharmik. COVID 19 DIAGNOSIS ANALYSIS USING TRANSFER LEARNING – DEEP LEARNING.
- [4] Ahmad Mohd Aziz Hussein, Abdulrauf Garba Sharifai, Osama Moh'd Alia, Laith Abualigah Khaled H. Almotairi, Sohaib K. M. Abujayyab & Amir H. Gandomi. Auto-detection of the coronavirus disease by using deep convolutional neural networks and X-ray photographs. [www.nature.com/scientificreports](http://www.nature.com/scientificreports).
- [5] Mohan Abdullah, Ftsum berhe Abrha, Beshir Kedir, Takore Tamirat Tagesse. A Hybrid Deep Learning CNN model for COVID-19 detection from chest X-rays. 2405-8440/© 2024 Published by Elsevier Ltd.
- [6] Aphelele Dumakude & Absalom E. Ezugwu. Automated COVID-19 detection with convolutional neural networks. [www.nature.com/scientificreports](http://www.nature.com/scientificreports).
- [7] Tarik El Le1, Mominul Ahsan and Julfikar Haider. Detecting COVID-19 from Chest X-rays Using Convolutional Neural Network Ensembles. Computers **2023**, 12, 105. <https://doi.org/10.3390/computers12050105> <https://www.mdpi.com/journal/computers>.
- [8] Mei-Ling Huang, Yu-Chieh Liao. A lightweight CNN-based network on COVID-19 detection using X-ray and CT images. 0010-4825/© 2022 Published by Elsevier Ltd.
- [9] Rubina Sarki, Khandakar Ahmed, Hua WangID, Yanchun Zhang, Kate Wang. Automated detection of COVID-19 through convolutional neural network using chest x-ray images. PLOS ONE | <https://doi.org/10.1371/journal.pone.0262052> January 21, 2022.
- [10] Tatiana Chakravorti, Vinay Kumar Addala, J. Shivam Verma. Detection and Classification of COVID 19 using Convolutional Neural Network from Chest X-ray Images. 2021 6th International Conference for Convergence in Technology (I2CT) Pune, India. Apr 02-04, 2021.

