SCIENCE PHOTONS Publishers

Research Article ISSN: 3143-4371

Innovative Insights in Digital Health

Improving Lung Disease Detection with DenseNet121 and CLAHE Integrated Screening

Raúl Isea*

Fundación IDEA, Hoyo de la Puerta, Baruta, Venezuela.

Received Date: 24 Apr 2026;
Accepted Date: 29 Jun 2026;
Published Date: 18 Jul 2026.

***Correspondence:** Raúl Isea, Fundación IDEA, Hoyo de la Puerta, Baruta, Venezuela.*

Citation: Raúl Isea. Improving Lung Disease Detection with DenseNet121 and CLAHE Integrated Screening. Innov Insights Digit Health. 2026; 2(2): 1-6.

ABSTRACT

Accurate diagnosis of pulmonary lesions, pneumonia and COVID-19 in the early stages is difficult due to overlapping symptoms and major logistic limitations of advanced molecular testing techniques. In this study, a computational approach is proposed in which the DenseNet121 architecture and the Contrast Limited Adaptive Histogram Equalization (CLAHE) are combined for analysis of full frame thoracic digital radiograph. The system reduces selection bias and computational latency while evaluating uncropped thoracic fields, preserving peripheral lung textures and minimizing interference from surrounding structures such as the rib cage and heart. Evaluation against an independent blind dataset of 50 images yielded a classification accuracy of 92%. The framework showed robust screening ability for infectious diseases with a recall of 1.00 for both COVID-19 and Viral Pneumonia, with F1-scores of 0.92 and 0.97, respectively. Interpretability validation with Gradient-weighted Class Activation Mapping (Grad-CAM) confirmed the network's focus on basal and peripheral pulmonary features, although visual inspection indicated a potential for shortcut learning due to radiopaque non-anatomical markers. This architecture is not intended to replace standard clinical diagnostic procedures; instead, it serves as a scalable triage layer to optimize clinical resource allocation, reduce specialist fatigue, and provide a screening tool medical centers and resource-limited clinics.

Keywords: *Artificial Intelligence; Deep Learning; DenseNet121; CLAHE; COVID-19; Pneumonia; Respiratory Pathology; Screening Systems.*

Introduction

Lung cancer and complex respiratory pathologies remain major causes of death worldwide, representing a significant clinical challenge, as the diseases often remain asymptomatic until advanced stages [1]. Molecular immunodiagnostics such as B-cell receptor (BCR) repertoire profiling and autoantibody detection against specific tumor-associated B-cell epitopes have improved early detection strategies. Bioinformatics platforms can now actively analyze these immunogenomic signatures and can provide highly sensitive detection mechanisms for micro-metastatic or early-stage malignancies. However, these biomolecular pathways are subject to substantial operational bottlenecks. This technique depends on advanced next-generation sequencing (NGS) technology and highly trained laboratory personnel and incurs substantial financial costs and turnaround times that hinder bulk screening in resource-limited or high-throughput triage settings.

Non-invasive affordable diagnostic modalities are needed for clinical infrastructure. Digital chest radiography (CXR) is a practical option for primary patient contact at a fraction of the cost of molecular infrastructure. The development of automated computational tools for the evaluation of radiographs provides a democratic alternative for healthcare systems, allowing underfunded or rural hospitals without advanced molecular laboratories to deploy effective screening protocols. This computational approach is not intended to replace high-tier bioinformatics or definitive oncological diagnostics, but to establish a primary-contact filter to optimize clinical resource allocation by prioritizing high-risk patients rapidly.

The need for a scalable triage layer was evident during pandemic records, where overlapping clinical presentations often resulted in the confusion of COVID-19 cases with acute pneumonia [2,6]. Early intervention directly improves patient survival, but manual interpretation of high-volume imaging is still hindered by a shortage of specialists and human fatigue. Moreover, the identification of pathological signs is challenging in human thoracic anatomy because bone densities and heart silhouettes can mask subtle le-

sions [3]. Traditional computer-aided approaches are often limited by error-prone lung segmentation steps [4]; in this study, non-segmented framework is evaluated. The proposed system processes the uncropped thoracic frame through DenseNet121 and Contrast Limited Adaptive Histogram Equalization (CLAHE) [7] to enhance micro-textural interstitial patterns [5]. This architecture addresses the challenge of distinguishing static bone geometry from active pathological textures in the global context of the chest, thus reducing diagnostic latency in high-pressure medical environments.

2. Methodology

The computational framework was developed in Python in a GPU-accelerated environment [8] by means of TensorFlow and Keras [9,13] for the deep-learning routines, OpenCV [10] and NumPy [11,12] for image pre-processing and numerical operations. To improve statistical reliability and clinical significance, we used a large dataset of 1,750 original chest radiographs divided into four diagnostic categories: normal (n = 416), COVID-19 (n = 433), viral pneumonia (n = 469), and lung opacity (n = 432).

A stringent data segmentation procedure was used to achieve an unbiased performance evaluation. Prior to any algorithmic alteration, balancing, or augmentation, precisely 50 original photos (about 12 per class) were isolated using stratified random selection to function as an independent, blind test set. This initial split is important to ensure the test data is not at all affected by any preprocessing or optimization steps, thus providing a real baseline for generalization and preventing data leakage. The remaining 1,700 radiographs were stratified and oversampled to account for class imbalances, resulting in a balanced pool of 1,820 instances, which were then split into a training subset (n = 1,456) and an internal validation subset (n = 364) at an 80/20 ratio.

To overcome the problem of structural overlap in unsegmented radiographs, all images were standardized to 160x160 pixels and processed using Contrast Limited Adaptive Histogram Equalization (CLAHE) with a clip limit of 1.2 and an 8x8 grid [14,18]. This preprocessing normalizes exposure variability across heterogeneous X-ray acquisition devices and enhances subtle parenchymal features [17], such as ground-glass opacities, while minimizing the visual interference of the rib cage and cardiac silhouette. Post-CLAHE images were converted to a 0-1 range and mapped to a three-channel RGB configuration to meet input requirements. The primary architecture used was DenseNet121, due to its dense connectivity pattern which helps to optimize feature propagation and alleviate the vanishing gradient problem [16]. Predictive optimization used a two-stage transfer learning approach [19]. In the first phase, we loaded the pre-trained ImageNet weights [20], froze the base layers, and used a learning rate of 0.001 to stabilize the classification head. To avoid memorizing and overfitting, we used a Global Average Pooling layer and a 0.4 Dropout layer on the data. In phase two, the architecture was unfrozen and the learning rate was decreased to 0.0001 to fine-tune the deep parameters and adapt to the morphological nuances of the respiratory pathologies. The final h5-serialized model will be made publicly available in a GitHub repository after completion of this study.

3. Results and Discussion

The first training phase showed robust convergence, with accuracy increasing from the baseline metrics to a validation plateau of 95.33% at the final fine-tuning epochs, where optimal weights were saved in HDF5 format. The network was evaluated after model convergence on the independent external blind testing subset (n = 50) to assess generalizability to unseen datasets without anatomical cropping or explicit lung segmentation.

3.1. Impact of CLAHE Preprocessing

The high-accuracy performance of the non-segmented network is directly due to local contrast optimization. Figure 1 shows the changes introduced by CLAHE to diagnostic clarity:

- Haze Effect Removal: CLAHE corrects for local contrast gradients due to inter-device calibration differences, distinguishing pulmonary parenchyma from background artifacts, and revealing internal tissue depth.
- Bone Transparency: The rib cage appears as high-density structures on plain films and obscures underlying pulmonary tissue. These structures become optically transparent following CLAHE, allowing the DenseNet121 layer to extract features directly from the skeletal geometry.

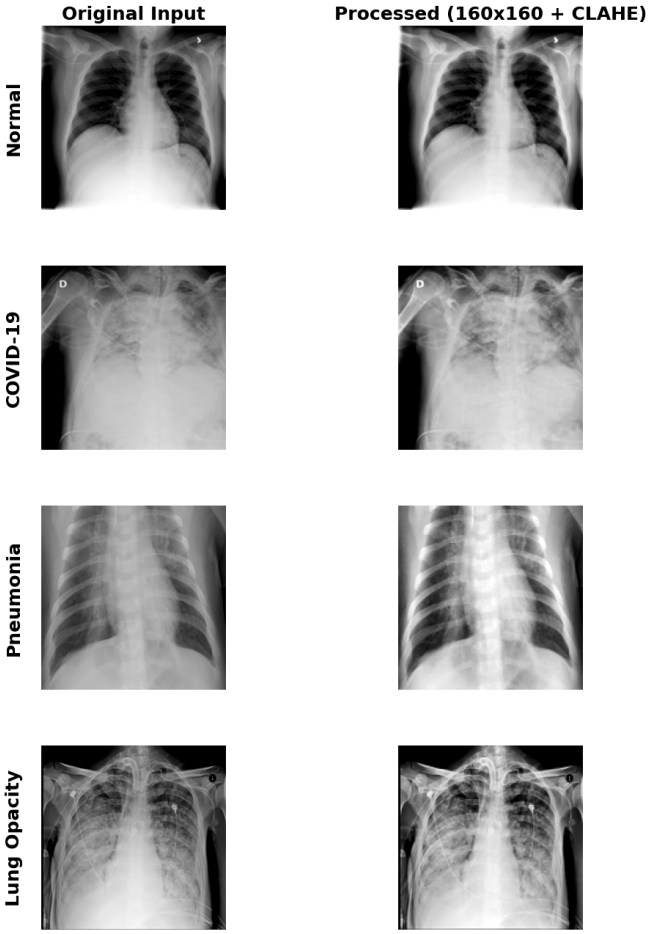

Figure 1. *Chest radiographs (raw and after CLAHE processing). The local contrast enhancement removes inter-device haze, increases the transparency of the skeleton and amplifies subtle alveolar textures for optimal feature extraction without segmentation.*

- Alveolar Texture Enhancement: Diffuse ground-glass opacities, usually perceived as subtle grey shadows, are upgraded into clear, rugose textures with distinct boundaries.

The model does not guess between raw greyscale gradients; instead, it processes stabilized texture maps, not uncorrected intensity variations, within the deep network. Texture-centric learning retains accuracy across the entire thoracic field without complex U-Net segmentation loops.

3.2. Quantitative Performance Metrics

The model achieved good clinical generalizability with an overall accuracy of 0.92 on the blind test set. The framework resulted in an F1 score of 0.92, recall of 1.00, and precision of 0.86 for COVID-19 screening such that no active viral infections were missed. Likewise, the Viral Pneumonia cohort exhibited an F1-score of 0.97, reflecting a good discriminative ability against the common structural patterns. The system demonstrated consistent performance in Lung Opacity cases with precision of 0.91 and recall of 0.83. For the Normal cohort, the network's precision and recall were 1.00 and 0.83 respectively; this conservative triaging behavior is by design, as it prioritizes safety and forwards ambiguous cases for human review rather than risk false negative classifications in healthy individuals [15-17].

As seen in the last testing matrix (Figure 2), the detection of active infectious patterns is still very reliable. The achieved absolute recall of 1.00 for both COVID-19 and Viral Pneumonia demonstrates the effectiveness of the model as an automatic screening tool for emergency settings.

3.3. Interpretability and Model Limitations

Gradient-weighted Class Activation Mapping (Grad-CAM) confirmed the clinical reasoning of the model (Figure 3). The network consistently targets the basal and peripheral lung fields, similar

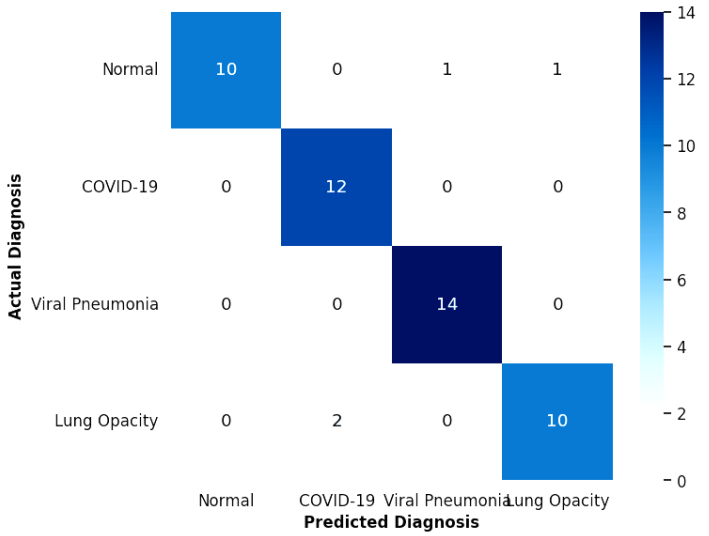

Figure 2. *Confusion matrix of DenseNet121 architecture tested on the independent blind test set (n=50). The framework achieves an absolute recall of 1.00 on active infectious profiles (COVID-19 and Viral Pneumonia).*

to the typical clinical distribution of ground glass opacities and consolidations, a hallmark of SARS-CoV-2 infections. This alignment confirms that the framework is attending to real pathological markers and not to global image artifacts.

But visual inspectability also revealed serious limitations regarding non-pulmonary noise. In one specific false positive case a healthy control image was misclassified as Lung Opacity. The Grad-CAM layer showed that the network had focused only on a radiopaque anatomical marker in the upper corner and the artificial black padding borders of the frame. This indicates a predilection for shortcut learning, with peripheral artifacts serving as surrogates for pathologic density change.

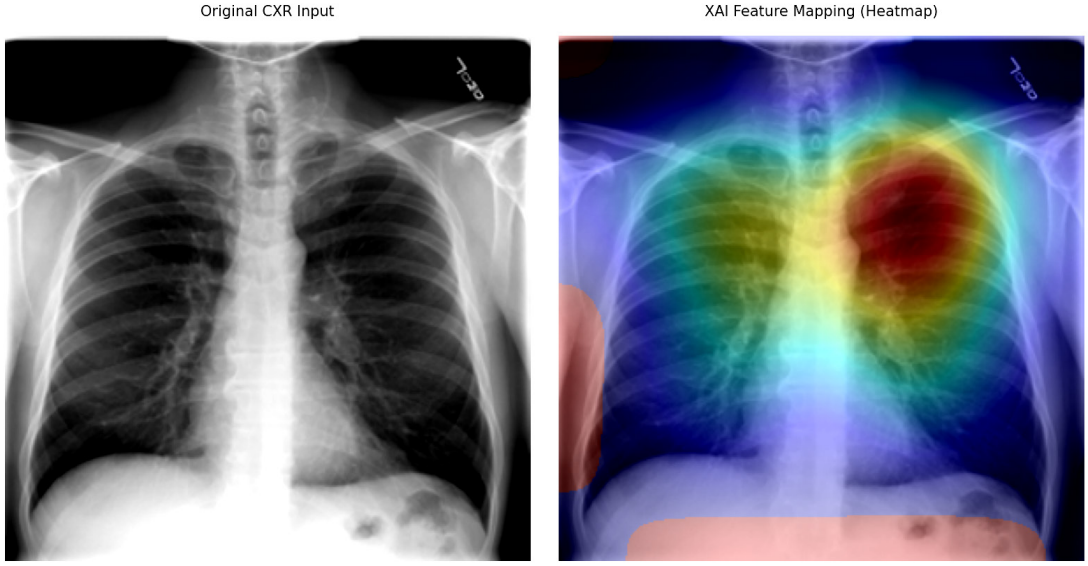

Figure 3. *Visual activation maps of Grad-CAM. The heatmaps confirm the network's specific focus on basal and peripheral parenchymal features indicative of viral infections but highlight vulnerability to external radiopaque markers (shortcut learning).*

Hence, the non-segmented approach significantly improves the processing speed and inference latency. However, the addition of the central boundary crop is a necessary refinement for future deployments to separate the parenchymal area from institution-specific markers. These results show that in our paper we focus on clinical transparency and offer a scalable first-contact triage tool, not a black box.

4. Conclusions

The proposed approach is a non-segmented pipeline for pulmonary disease screening. The DenseNet121 and CLAHE combination provides robust diagnostic performance without the computational overhead of lung segmentation algorithms. Processing the full-frame thoracic radiograph avoids the inadvertent loss of peripheral pathological features, which is often associated with input cropping. Local contrast enhancement combined with dense feature propagation enables the identification of subtle interstitial patterns otherwise masked by the structural geometry of the ribs and cardiac silhouette.

Testing on an isolated blind subset (n = 50) confirmed the generalizability of the model across all four diagnostic categories. The recall for both COVID-19 and Viral Pneumonia is 1.00, which demonstrates the network as a reliable emergency triage tool. Furthermore, Grad-CAM activation mapping provided the necessary clinical transparency to ensure the predictions are driven by valid parenchymal features while successfully identifying the network's vulnerability to non-anatomical markers. This interpretability is an important step in the validation of automated models prior to clinical deployment.

This non-segmented pipeline is more efficient than U-Net architectures, but the peripheral artifact vulnerability demonstrates the need for a central-focus crop to achieve institutional generalizability as a refinement. The framework offers an accessible fast triage layer that can be added to conventional diagnostic pipelines, optimizing resource allocation and minimizing diagnostic delays in high-pressure medical environments and low-resource clinics.

Acknowledgments

The authors would like to thank anonymous reviewers and editorial board of Innovative Insights in Digital Health for their constructive recommendations. They provided direct feedback for the improvement of the technical robustness, methodological transparency and final presentation of this computational framework.

Appendix A: Computational Implementation

We implemented the full computation pipeline for data set orchestration, feature enrichment and biphasic network optimization in a Python environment. This section describes the main algorithmic blocks and the structure of the serialized deployment file for ensuring scientific reproducibility and meeting the validation criteria without compromising the algorithmic opacity.

A.1. Initialization and Contrast-Limited Preprocessing Routine

The first phase to overcome the problem is the mathematical normalization of baseline exposure variations of heterogeneous X-ray source devices; using Contrast Limited Adaptive Histogram Equalization (CLAHE).

Python

```
import os
import cv2
import numpy as np
from tensorflow.keras.applications import DenseNet121
from tensorflow.keras import layers, models
from tensorflow.keras.optimizers import Adam
from tensorflow.keras.callbacks import EarlyStopping
from sklearn.model_selection import train_test_split
from sklearn.metrics import confusion_matrix, classification_report
def load_and_preprocess_pipeline(routes_dict):
    X, y = [], []
    # Configure CLAHE to avoid noise amplification across local grids
    clahe = cv2.createCLAHE(clipLimit=1.2, tileGridSize=(8, 8))
        for route, label in routes_dict.items():
        if os.path.exists(route):
            files = [f for f in os.listdir(route) if f.lower().endswith((‘.png’, ‘.jpg’, ‘.jpeg’))]
            for name in files:
                # Read in grayscale to compute localized histogram equalization
                img = cv2.imread(os.path.join(route, name), cv2.IMREAD_GRAYSCALE)
                if img is not None:
                    img_clahe = clahe.apply(img)
                    # Remap to 3-channel RGB to satisfy DenseNet tensor constraints
                    img_rgb = cv2.cvtColor(img_clahe, cv2.COLOR_GRAY2RGB)
                    img_resized = cv2.resize(img_rgb, (160, 160))
                    X.append(img_resized / 255.0)
                    y.append(label)
    return np.array(X), np.array(y)
```

A.2. Data Partitioning and Anti-Leakage Protocol

Strict stratified separation isolation protocol is applied before performing any oversampling operations, to remove potential vectors of data leakage or selection bias.

Python

```
# Step A: Isolate exactly 50 original images for independent blind testing
X_temp, X_test, y_temp, y_test = train_test_split(
    X_total, y_total, test_size=50, stratify=y_total, random_state=42
)
# Step B: Stratified oversampling restricted exclusively to the training pool
max_c = max([np.sum(y_temp == i) for i in range(4)])
X_res, y_res = [], []
for i in range(4):
    idx = np.where(y_temp == i)[0]
    extra = np.random.choice(idx, max_c, replace=True)
    for e in extra:
        X_res.append(X_temp[e])
        y_res.append(y_temp[e])
# Step C: Divide balanced instances into 80/20 train/validation segments
X_train, X_val, y_train, y_val = train_test_split(
    np.array(X_res), np.array(y_res), test_size=0.2, stratify=np.array(y_res), random_state=42
)
```

A.3. Network Architecture and Biphasic Optimization Logic

The classification pipeline is transitioned from frozen feature extraction to deep parameter fine-tuning through a transfer learning pipeline with a DenseNet121 backbone.

Python

```
# Feature Extractor Setup
base_model = DenseNet121(weights='imagenet', include_top=False, input_shape=(160, 160, 3))
base_model.trainable = False

model = models.Sequential([
    base_model,
    layers.GlobalAveragePooling2D(),
    layers.Dropout(0.4), # Rigorous regularization layer against memorization
    layers.Dense(4, activation='softmax')
])

early_stop = EarlyStopping(monitor='val_accuracy', patience=4, restore_best_weights=True)

# Phase 1: Classification Head Stabilization
model.compile(optimizer=Adam(learning_rate=0.001), loss='sparse_categorical_crossentropy', metrics=['accuracy'])
model.fit(X_train, y_train, epochs=15, validation_data=(X_val, y_val), batch_size=32, callbacks=[early_stop])

# Phase 2: Deep Parenchymal Fine-Tuning
base_model.trainable = True
model.compile(optimizer=Adam(learning_rate=0.0001), loss='sparse_categorical_crossentropy', metrics=['accuracy'])
model.fit(X_train, y_train, epochs=20, validation_data=(X_val, y_val), batch_size=16, callbacks=[early_stop])

# Model Serialization
model.save('diagnostico_pulmonar_DenseNet.h5')
```

A.4. Hierarchical Structure of the Serialized .h5 Container

The final compiled output of the system is packaged in a container file using the Hierarchical Data Format (HDF5). This container, rather than being a black-box script that has not been verified, organizes the mathematical parameters of the model into specific layers:

1. **Root Group Attribute Meta-Graph:** models the absolute structural topology of the sequential network. It depicts the tensor dimensions moving from the standard 160x160x3 RGB input matrix to the end classification outputs.
2. **model_weights Dataset Cluster:** Contains the final optimized bias and kernel tensors which are continuous multidimensional arrays. It constrains the frozen parameter distributions with

the tuned weights of the fine-tuning in Phase 2, acquired from the deep convolutional blocks (conv5_block16_concat).

3. **optimizer_weights Group:** Internal tracking parameters of the Adam optimizer, i.e. moving averages of the previous gradients and squared gradients. This layer preserves the exact mathematical state of convergence, ensuring the stability and reproducibility of future inference routines.

Conflicts of Interest

The authors declare no conflict of interest and received no specific funding for this work.

References

1. Siegel RL, Miller KD, Wagle NS, Jemal A. Cancer statistics, 2023. CA Cancer J Clin. 2023; 73: 17-48.
2. Isea R. Predicción computacional cuantitativa de epítopos de células B. VacciMonitor. 2015; 24: 93-97.
3. Aparicio G, Blanco F, Blanquer I, Bonavides C, Chaves JL, et al. Developing Biomedical Applications in the Framework of EELA. In: Udoh E, Wang FZ, editors. Handbook of Research on Grid Technologies and Utility Computing: Concepts for Managing Large-Scale Applications. Hershey, PA: IGI Global Scientific Publishing; 2009: 206-218.
4. Ilzins O, Isea R, Hoebeke J. Can Bioinformatics Be Considered as an Experimental Biological Science? arXiv. 2016;1607.04836.
5. Isea R, Mayo-Garcia R, Restrepo S. Reverse vaccinology in Plasmodium falciparum 3D7. arXiv. 2016;1611.01692.
6. Marginean CM, Popescu M, Vasile CM, Cioboata R, Mitrut P, et al. Challenges in the Differential Diagnosis of COVID-19 Pneumonia: A Pictorial Review. Diagnostics (Basel). 2022; 12: 2823.
7. Wang X, Peng Y, Lu L, Lu Z, Bagheri M, et al. ChestX-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. Proc IEEE Conf Comput Vis Pattern Recognit. 2017: 2097-2106.
8. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015). 2015: 234-241.
9. Türk F, Kökver Y. Detection of Lung Opacity and Treatment Planning with Three-Channel Fusion CNN Model. Arab J Sci Eng. 2024; 49: 2973-2985.
10. Gong H, Guo W, Neil MAA. GPU-accelerated real-time reconstruction in Python of three-dimensional datasets from structured illumination microscopy with hexagonal patterns. Philos Trans A Math Phys Eng Sci. 2021; 379: 20200162.
11. Domínguez C, Heras J, Pascual V. IJ-OpenCV: Combining ImageJ and OpenCV for processing images in biomedicine. Comput Biol Med. 2017; 84: 189-194.
12. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, et al. Array programming with NumPy. Nature. 2020; 585: 357-362.
13. Mohammad N, Muad AM, Ahmad R, Yusof MYPM. Accuracy of advanced deep learning with TensorFlow and Keras for classifying teeth developmental stages in digital panoramic imaging. BMC Med Imaging. 2022; 22: 66.
14. Sandler M, Howard A, Zhu M, Zhmoginov A, Chen LC. MobileNetV2: Inverted residuals and linear bottlenecks. Proc IEEE/CVF Conf Comput Vis Pattern Recognit. 2018: 4510-4520.
15. Okyere-Gyamfi S, Asante M, Peasah KO, Missah YM, Akoto-Adjepong V. Contrast limited adaptive histogram equalization (CLAHE) and colour difference histogram (CDH) feature merging capsule network (CCFMCapsNet) for complex image recognition. PLoS One. 2025; 20: e0335393.
16. Nirupama, Virupakshappa. MobileNet-V2: An Enhanced Skin Disease Classification by Attention and Multi-Scale Features. J Imaging Inform Med. 2025; 38: 1734-1754.
17. Prashant268. Chest X-ray COVID-19 & Pneumonia Dataset [Internet]. Kaggle; 2023. Available from: https://www.kaggle.com/datasets/prashant268/chest-xray-covid19-pneumonia
18. Zuiderveld K. Contrast limited adaptive histogram equalization. In: Heckbert PS, editor. Graphics Gems IV. San Diego, CA: Academic Press; 1994: 474-485.
19. Alzubaidi L, Zhang J, Humaidi AJ, Al-Dujaili A, Duan Y, et al. Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. J Big Data. 2021; 8: 53.
20. Deng J, Dong W, Socher R, Li LJ, Li K, et al. ImageNet: A large-scale hierarchical image database. Proc IEEE Conf Comput Vis Pattern Recognit. 2009: 248-255.