Pneumonia Disease Classification Models with Ensemble Transfer Learning Approach

Akinbo Racheal S.
Computer Science Department
Federal University of
Technology Akure, Nigeria

Olatubosun Olabode
Data Science Department
Federal University of
Technology Akure, Nigeria

Daramola Oladunni A.
Information Technology
Department, Federal University
of Technology Akure, Nigeria

Emmanuel O. Ibam
Department of Information Systems,
Federal University of Technology, Akure, Nigeria

ABSTRACT
Pneumonia, a lung disease that affects both younger children and older adults, causing respiratory problems. Pneumonia disease has been known to cause several deaths, and there is a need for a faster method of diagnosing the disease for early treatment. The use of machine learning techniques in various fields has delivered excellent results as it has been applied to address challenges encountered in society. It has been utilized in several domains like medical, environment, security, industrial, business, finance and many more. Machine learning applications are characterized by the ability to process massive amounts of data, identify patterns, relationships, and provide logical interpretations in the identification, classification or prediction of an outcome. One aspect of machine learning called transfer learning utilizes deep convolution. In this research, the transfer learning technique is applied to classify pneumonia. The radiologist's readings could be subject to inter-class variability, and delayed results can occur. The methodology involved the application of five transfer learning techniques, namely, MobileNet-V2, GoogLeNet, AlexNet, DenseNet, and VGG-19. The individual experimental results show an accuracy of 0.79, 0.96, 0.96, 0.83, and 0.91, respectively. The results revealed that GoogLeNet and AlexNet have the best accuracy, while VGG-19 followed closely with 0.91, DenseNet and MobileNet-V2 record the lowest. GoogLeNet and AlexNet have a high precision of 0.99. The weighted ensemble applied gives an approximate accuracy of 0.97. This research shows an improvement on the other models reviewed and can be deployed to fast-track the detection of pneumonia.

General Terms
Image Processing, Medical Images, Pattern Recognition Artificial Intelligence and Machine Learning in Medicine, Digital Medicine

Keywords
Pneumonia, detection, classification, transfer learning, deep learning, ensemble.

1. INTRODUCTION
Pneumonia is an infection that inflames the lung tissue and is a significant global health issue.1. It may cause symptoms like fever, coughing, and breathing difficulties. Various pathogens, including bacteria, viruses, and fungi, can cause it, with viral pneumonia being a specific concern due to the ongoing COVID-19 pandemic caused by the novel coronavirus (SAR-CoV-2). Timely and accurate diagnosis of pneumonia is vital to its effective treatment and management. It is crucial to take it seriously. A descriptive cross-sectional study by 2 in 2023 examined the incidence of pneumonia among children below five and the triggers that influence it. 2 attended a health care center in the city of Lagos, Nigeria's Amuwo Odofin Local Government Area. The study employed a two-stage sampling strategy, and EPI INFO version 7.0 was used for data analysis. Logistic regression and chi-square tests were utilized in the analysis. 42 youngsters (12.7%) out of the 330 tested developed pneumonia, according to the report. This indicates that about 1 in 8 of the study's under-5 participants got pneumonia. Prematurity has been demonstrated to raise the possibility of pneumonia; the study found. 2

Pneumonia is a common illness that strikes both adults and children worldwide. The highest death rates were found in the Asian continent and sub-Saharan Africa, according to statistics of World Health Organization’s 3 in 2022 study. Unbelievably, 4 stated in 2020 that over two million Nigerian children may die from pneumonia in the following ten years if more wasn't done to prevent it.

2. RELATED WORKS
5 carried out a thorough survey of computer-aided detection methods for chest ailments in 2018. The study explored general image preprocessing techniques, including contrast enhancement, segmentation, and bone suppression, focusing on datasets from chest X-rays. The study's intention was to detect several disorders associated with the human chest, with an emphasis on the algorithm's basic principles, the data that was used, evaluation metrics, and the outcomes. Nevertheless, no real-world implementation was carried out while talking about the artificial intelligence-based CAD system in work done using the chest radiography. An investigation into the possibility of using machine learning finding and identifying pneumonia on a chest X-ray picture automatically was carried out by 6 and associates. They employed an ensemble of convolutional networks called Mask R-CNN and RetinaNet to identify pneumonia. The ensemble did, however, receive poor F-1, recall, and precision ratings of 0.77, 0.79, and 0.75, respectively. These scores suggest that the findings weren't appropriate for use in real-world implementation.

An adversarial approach was presented by [7] to improve and broaden the performance of a pneumonia classifier derived from chest radiographs. They removed the reliance regarding dataset's origin and the classifications with X-ray perspective by using a deep learning system built on adversarial optimization. Regretfully, the obtained outcomes (AUC 74.7%) are insufficient for field implementation.

The goal of [8] study from 2021 was to identify cases of pneumonia within paediatric patients' chest X-ray images. Using four distinct pre-trained models, the study used transfer learning strategy founded on automated convolutional neural networks. The categorization in binary form gave an accuracy ranging from 83% - 86%, according to the data. It is crucial to remember that the study's tiny dataset prevented the results from being broadly applied to bigger populations.

In 2022, [9] did a study that examined the development of deep learning in connection with COVID-19 and pneumonia. The study's goals were to evaluate and elucidate the quantitative and qualitative developments in deep learning techniques. Describing several deep-learning model designs and comparing their efficacy for identical tasks. Reviews of several deep learning techniques for pneumonia prediction, but without any practical applications.

An ensemble model for convolutional neural networks was created in 2023 by [10] to help with the pneumonia diagnosis revealed on a radiograph of the chest. The model was designed to be extremely precise, lightweight, and transportable. Three unique models using three different kernel sizes were employed in the architecture. With a f1-score of 88.56% and 99.23% is a high recall percentage. The model performed well. However, the design yielded reduced accuracy and precision because of overfitting and a lack of large datasets. However, the recall score of the model was good.

Numerous studies have been conducted on the classification and detection of pneumonia, as reported in the literature. However, there are some gaps that still need to be addressed. The following are the limitations of the studies reviewed, particularly in [10, 2] of incorrect diagnosis, overfitting problems, limited data and few models used, poor prediction, and sensitivity performance.

2.1 Objectives

Using Transfer Learning techniques, this research aims to construct an Ensembled Classification. Chest X-ray images are used to predict the occurrence of pneumonia, using models. The performance of the models will be assessed utilizing standard metrics.

3. METHODOLOGY

This work aims to develop transfer learning models that predict pneumonia and identify lung anomalies from an X-ray of the chest (radiograph) images. Chest X-rays (radiographs) are used in the suggested work's implementation, and five transfer learning models—GoogLeNet, DenseNet, MobileNet-V2, VGG19, and AlexNet—are used in the experimental evaluation. The Lung X-Ray pictures are then classified and the Normal/Pneumonia class is predicted using an ensemble classifier. The preprocessing module processes the Chest X-ray (CXR) picture before being sent for additional analysis.

The dataset utilized in this study is arranged with each image category into two subfolders (Pneumonia/Normal) within the train and test folders. 11,064 JPEG X-ray pictures from the Kaggle database, divided separated into two groups: pneumonia and normal. make up the dataset. Originally, before examining all chest radiographs and chest X-ray pictures were initially screened to remove any low-quality or illegible scans in order to maintain quality control.

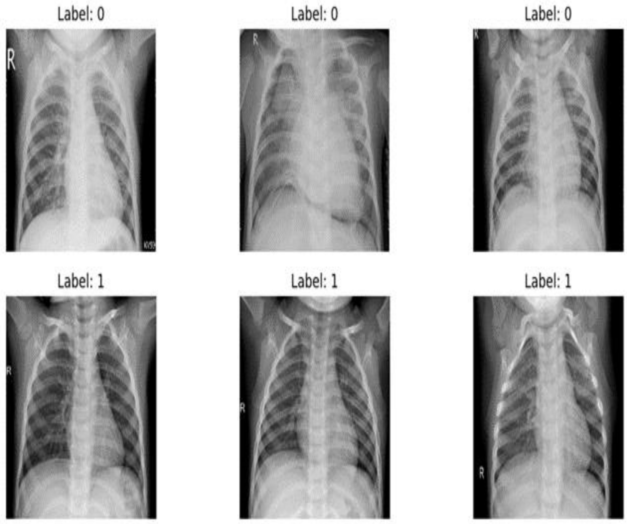

Fig 1: Images Dataset

3.1 Components of the Architecture

i. Data Collection and Description

To conduct research on pneumonia classification, a significant amount of data is required. This data is collected from Kaggle's publicly available database. The JPEG-formatted photos in the collection are categorized into two groups: Pneumonia and Normal. To enhance performance, the images are augmented, resized, and normalized. The Kaggle dataset was chosen for this research because it contains a variety of datasets, making it an ideal resource for conducting comparative analysis and advancing scientific research. The sample of x-ray image data collected is seen in figure 1. below and the architecture of the system in Figure 2.

i Pre-Processing

During the first phase of machine learning, raw data is generated which needs to be pre-processed to reduce complexities. The dataset consists of JPEG images categorized into Pneumonia and Normal classes. The pixel values range from 0 to 1.

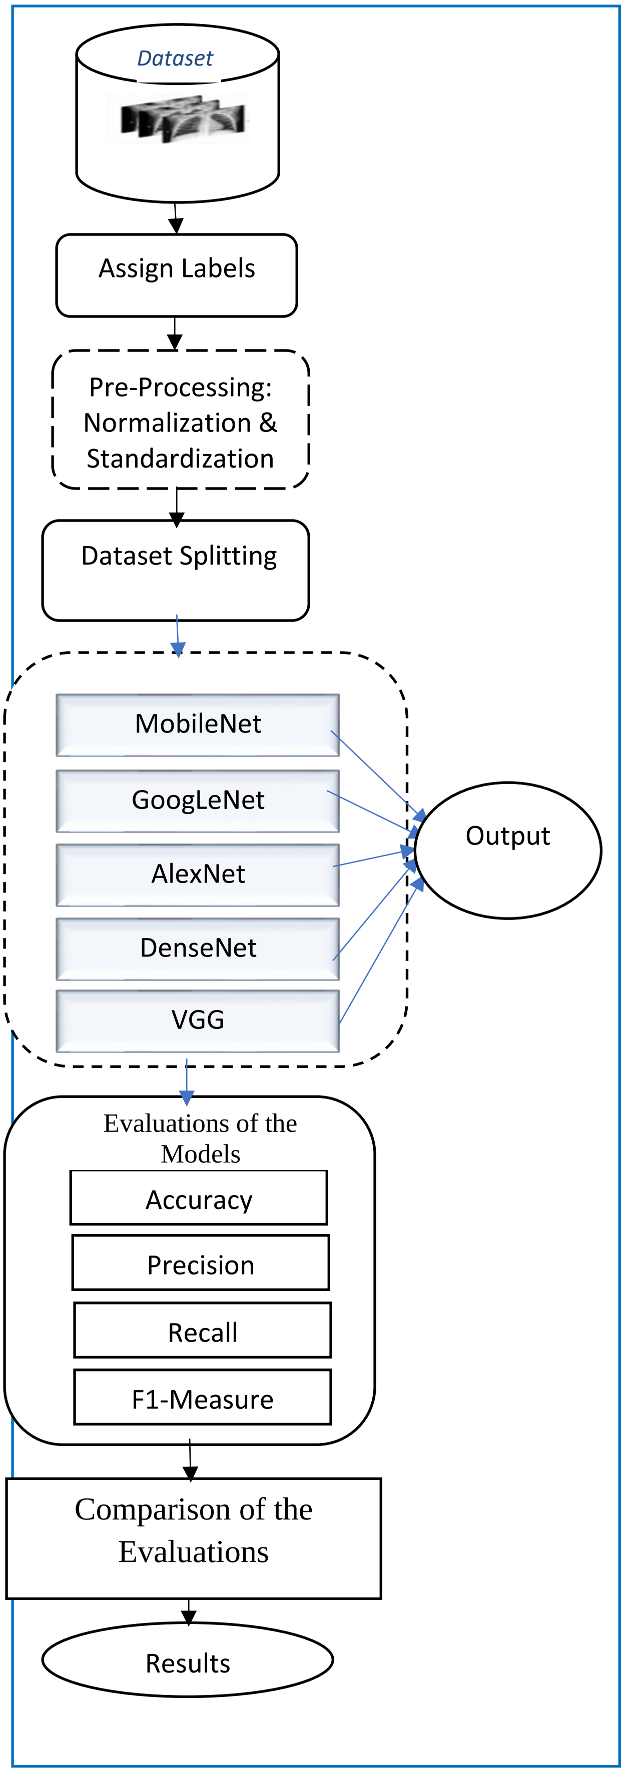

Fig 2: Model Architecture Table 1: Class Encoding

Class	Encoding
Normal	**0**
Pneumonia	**1**

ii. Model Training/Splitting
The dataset contains information about normal and pneumonia cases. The data will be split into two datasets:

1. Training dataset: The dataset in question will be used to train the model. and it will comprise the bulk of the data. Ratio of training data will be 75% of 11,064 which is 8,298.

2. Test dataset: The generalizability of the model to fresh data is evaluated using the test data; it is withheld during training. Test data constitutes 25% of 11,064 totalling 2,766 of dataset.

Table 2. The Class Ratio

Train	Test	Total
8,298	2766	11,064

3.2. Transfer Learning Classification Models

In this study, we classified and predicted pneumonia from X-ray pictures of the lungs using five different classical deep-learning architectures. MobileNet-V2, GoogLeNet, AlexNet, DenseNet, and VGG-19 were the names of these designs. We trained only the classifier heads and froze the weights of the feature extractors (network bodies) to assess their performance. We continued with full training after choosing the most accurate networks for the second step. Train the classifier and the feature extractor, feature extractor and the classifier, required unfreezing the weights of the entire network.

(1) MobileNet-V2: The Inverted Residual Block and the two primary components of the architecture are the Bottleneck Residual Block. Residual bottleneck block aims to expand the model's capacity by introducing non-linearity to the residual connections, the Inverted Residual Block is intended to maximize the model's efficiency by lowering the computing cost of the residual connections. The 1x1 Convolution and the 3x3 Depthwise Convolution are the two types of convolution layers used in the MobileNet V2 architecture. The input tensor's number of channels can be changed using the former, and its characteristics can be extracted using the latter. [11]

(2) GoogLeNet: It is a complex neural network that utilizes deep convolutional layers. It is a part of the Inception Network, which was developed by researchers at Google. The network comprises of 22 layers including 9 inception blocks, which are grouped into three. A max-pooling layer is put in between the groups to help reduce the dimensionality of the input. The network's head includes a global average pooling layer that produces estimates. GoogLeNet accepts images with RGB colour channels and sizes of 224 x 224 pixels. The activation function used throughout all the convolutional layers is Rectified Linear Units (ReLU). [12]

(3) AlexNet: This particular kind of neural network is made especially for problems involving picture classification. Five convolutional layers, two normalization layers, three layers of max-pooling one SoftMax layer, and three completely connected layers make up this structure. A convolution filter and a non-linear activation function called "ReLU" are present in every convolutional layer. The pooling action is carried out by the max-pooling layers, and as layers that are entirely connected are present, input size is fixed. According to [13], the input size is defined as 224x224x3. This indicates that the input image has three colour channels (red, green, and blue) and measures 224 pixels wide by 224 pixels tall.[13]

(4) DenseNet: Densely Connected Convolutional Network, as it is known, is an architecture in which all layers are connected to all other layers. With L(L+1)/2 direct links in a network, there are L levels. Every layer takes as input the feature maps of every layer that came before it, and every layer that came after it took its input from those feature maps. The architecture is made up of multiple dense blocks, a classification layer, and transition layers in order of appearance. Using 64 filters of size 7x7 and a stride of 2, the first convolution block is a Max Pooling layer with 3x3 max pooling. [14]

(5) VGG-19: The University of Oxford Visual Geometry Group built the deep convolutional neural network. It stands out due to its 19-layer depth, which increases both its computational complexity and cost compared to VGG-16, its predecessor. The convolutional and max-pooling layers of the VGGNet-19 are arranged in a sequence to process images with a pixel size of 224 x 224. While the max-pooling layers minimize the spatial dimensions of the output feature maps, the convolutional layers retrieve features from the input image. There are nineteen layers in total in the network: three fully-connected layers and sixteen convolutional layers. There are sixteen convolutional layers and three completely connected layers in this. [15]

3.3 Training

To avoid overfitting, we used the Early Stopping regularization strategy during the network training procedure. We used a ten-epoch patience criterion and monitored validation loss. We employed cross-entropy covers both the first and second stages of network training. This was computed using the following formula: The number of classes, denoted by c, is two in this case. The ternary indicator (ground-truth label) is represented by yi, and the c-th class's SoftMax probability is represented by pi. We included a tiny positive constant, ε, to prevent an undefined case of log (0).

3.4 Ensemble Classification

By merging the predictions of several models, the weighted average ensemble technique is a helpful machine-learning tool that can improve a model's performance. This method applies the weighted ensemble classifier to five models of transfer learning. Figure 2 illustrates the model architecture and Figure 3, ensemble architecture. For the Normal and Pneumonia Classes, the binary number is represented by the class encoding in table 1.

3.5 Analysis of Methodology

1.The Convolution layer: Applying the n × n convolution kernel to using an identically-sized m × m image, we create a new image that is referred to as the feature map. map of features. Calculating the feature map's dimensions, we utilize the formula that follows: ((m+n)/l) + 1, where l is the convolution step size that we choose. Finally, output is given as:

$$D_k\, \theta_k = (\mathcal{X}1,\ \mathcal{X}2,\ \dots\dots, \mathcal{X}m)T \quad (1)$$

the input layer (l) receives an image as its input (X). The image has three colour channels (red, green, and blue), and a convolution filter with a set of kernels is applied to each of these channels. The number of split attributes is represented by D, and $\theta_k$is the kth sample from the vector. $The\ input\ of\ the\ layer\ L\ is\ X.$ A convolution filter with a set of kernels is applied to each of the three-color channels of an input image. This is represented with:

$$x_j^l = \mathrm{f}\left(\sum_{i\in M} x^{l-1} \ * k_{ij}^l + b_j^l\right) \quad (2)$$

2. Pooling Layers: To prevent overfitting, a dropout layer is used in conjunction with a pooling layer to limit feature mappings and system parameters. The activation of max pooling can be calculated using the following formula:$a_{b.k=max_{d=1}^{r}(h_{b,(c-1)(n+d)}}$ (3)

where ab, c is denoted by subsampling factor n, d, and pooling scale r.

3. Activation Function Layer- (Rectified Linear Unit): To convert input values into a non-negative range, one layer of the Rectified Linear Unit (ReLU) is the activatinfunction It applies to all input values in when f(x) = max(0, x)**.** In simpler terms, the ReLU layer turns all negative values to zero by applying the following function:

$$y\ = \max(0,\ x) \quad (4)$$

4. Learned Feature Maps: Various features found in an image, such as bends during convolution, pooling, and ReLU activation, edges, and vertical and horizontal lines, are identified using feature detectors.

5. Fully Connected Layer: Finally, the function that determines the completely linked layer "l" is:

$$x^1 = \mathrm{f}\,(u^l),\quad u^l = w^l\ x^{l-1} + b^l \quad (5)$$

where the activation function is denoted by
where f (.) In this case, the sigmoid function is used, and the training sample error loss is as follows:

$$E^n = \frac{1}{2}\sum_{k=1}^{c}\left(t_k^n\ -\ y_k^n\right)^2 - \frac{1}{2}\|t^n - y^n\|_2^2 \quad (6)$$

where the number of classes in the category is represented by c.

6. SoftMax layer: The SoftMax function is calculated using the following equation:

$$F(x_i) = \frac{e^{\,x_i}}{\sum_{j=0}^{k} e^{x_j}} \quad (7)$$

where $i = 0,\ 1,\ 2,\ \dots,\ k,\qquad x_i\ =\ x_1,\ \ x_2, \dots,\ x_k$, the input value and x_j represent the output of the values.

7. Back Propagation: In back propagation, the sensitivity of layer I can be written as follows:

$$\delta^i\left(w^{i+1}\right)\mathrm{T}\ \delta^{i+1} \cdot f^1\left(u^i\right) \quad (8)$$

The weights in the output layer w have a base value.

8. Flattened layer: Vectors can be flattened to create a fully linked layer for the final classification model, which includes all pixel data in a single layer.

$$u_{flatten}^i = \mathrm{flatten}(u_{ij}^l \quad (9)$$

9. Weighted Ensemble

$w^{(i)} = \sum_{x \epsilon A}(i) tan\ h(x)$ (10)

$Ensj = \frac{\sum_{iw}(i)\ X\ P_{j(i)}}{\sum_i w(i)}$ (11)

$y = Argmax\ (ens) \sum_{j=1}^{m} wjpij$ (12)

$wj\ is\ the\ weight\ of\ an\ individul\ learner\ that\ can\ be\ assigned\ to\ the\ jth\ classifier$ =1.

3.6 Evaluation Metrics

To evaluate the classification models' effectiveness, the performance indicators listed below are employed.

i. Precision: This represents a score that assesses how well a classifier predicts the good outcomes. It displays the percentage of accurately identified positive cases, or "true positives," relative to all positive predictions, or "false positives" and "true positives. The precision score is determined using the following formula: TP / (TP + FP), where FP stands for false positives and TP denotes true positives.

ii. Recall: Also known as true positive rate or sensitivity, it gauges how well the model can identify positive instances among all of the real positive examples. TP represents the quantity of correct positive predictions. and the number of wrong negative predictions is represented by FN. The calculation of it involves dividing TP by the total of TP and FN.

iii. F1-score: This is an indicator of how well a model classifies cases. To help balance the two measures, The recall and precision harmonic means are employed in its calculation. F1-score offers a solitary figure that symbolizes the model's overall efficacy in accurately categorizing situations. To find the F1 score, multiply the total precision and recall by two and divide the result by the total.

iv. Accuracy: One way to calculate a model's accuracy is to take divide the total by the sum of these values to get the False Negatives (FN), False Positives (FP), True Negatives (TP), and True Negatives (TN). Accuracy is computed using the formula Accuracy = (TP + TN) / (TP + TN + FP + FN).

It is crucial to remember that for every given input, all outputs must add up to one and be positive.

3.7 Model Training and Evaluation

There are training and testing versions of the dataset following preprocessing. Numerous algorithms for transfer learning including GoogLeNet, DenseNet, MobileNet, VGG and AlexNet are used for model training. The model's effectiveness in accurately detecting and classifying pneumonia is assessed using the use of metrics including F1-score, recall, accuracy, and precision. Python libraries such as numpy, pandas, and matplotlib along with Keras in TensorFlow are used in this classification project.

3.8. Hyperparameters and Parameters

Hyperparameters are essential for maximizing the effectiveness of the models that are employed. These hyperparameters are carefully selected to achieve the best results. Below are the key hyperparameters and parameter options considered for the performance metrics: For any input, the outputs must all be positive and they must sum to unity.

Table 3. Hyperparameters

Parameters	Option
The Cost function	Categorical cross entropy
Learning Rate as (Lr)	0.0000001
The Optimizer	ADAM
Number of Epoch	100

Table 4. Model Parameter

Model	Parameters
MobileNet	6,981,698
GoogLeNet	57,300,096
AlexNet	25,719,386
DenseNet-121	11,236,930
VGG -19	23,175,234

3.9 Experimental Result

The machine learning models data are into training and testing dataset on pneumonia images. Grayscale images from specified directories for test, train, and validation sets are loaded and resized to 200x200 pixels. The data is organized into tuples containing the resized image and its corresponding class label (0 for 'NORMAL' and 1 for 'PNEUMONIA'). The dataset is then shuffled to introduce randomness. Furthermore, the code uses Matplotlib to create a bar chart that illustrates the distribution of samples between the dataset's "NORMAL" and "PNEUMONIA" classifications. Python programming, an NVIDIA Tesla K80 graphics card, 12GB of VRAM, a Linux operating system with TensorFlow, Numpy, Pandas, OpenCV, Sklearn, Keras, Matplotlib, Seaborn, and X-ray images are used in the implementation of the categorization of viral pneumonia disease.

(i) **Confusion Matrices of the Transfer Learning Classifiers**

Confusion matrix shows TP, TN, FP, and FN for each class to detail model performance.

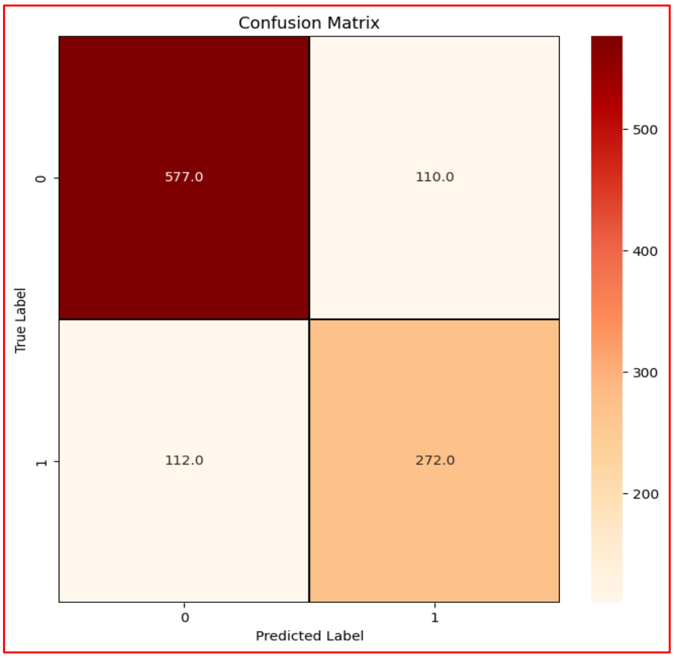

Fig 3: Confusion Matrix Result of MobileNet-V2

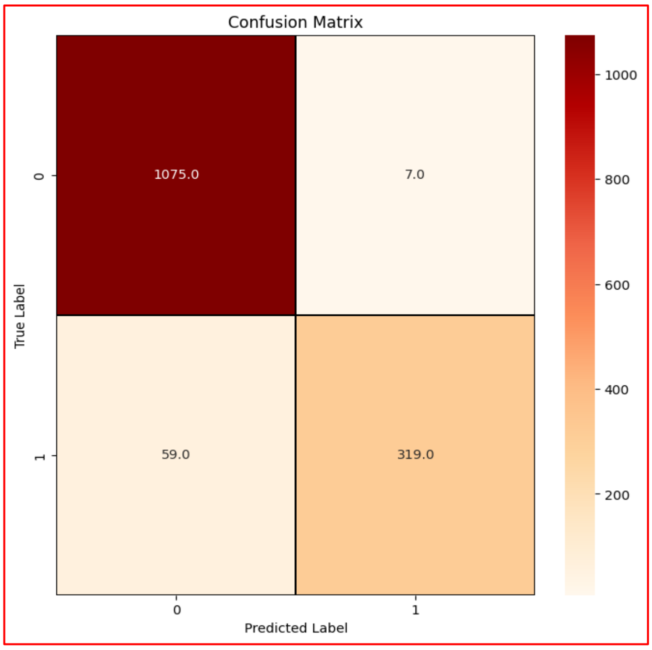

Fig 5: Confusion Matrix Result of AlexNet

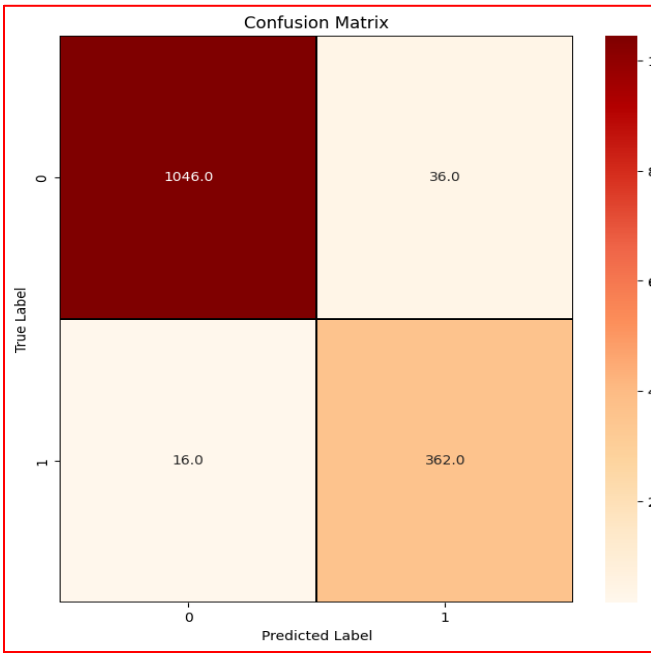

Fig 4: Confusion Matrix Result of GoogLeNet

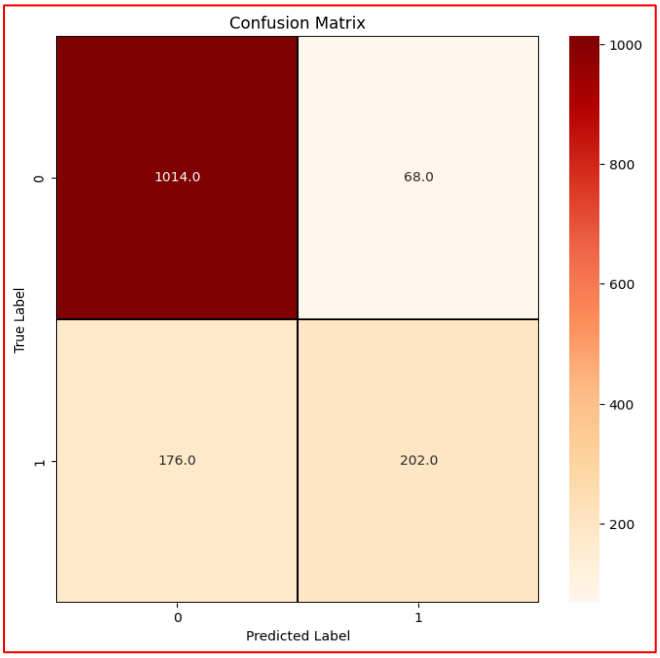

Fig 6: Confusion Matrix Result of DenseNet

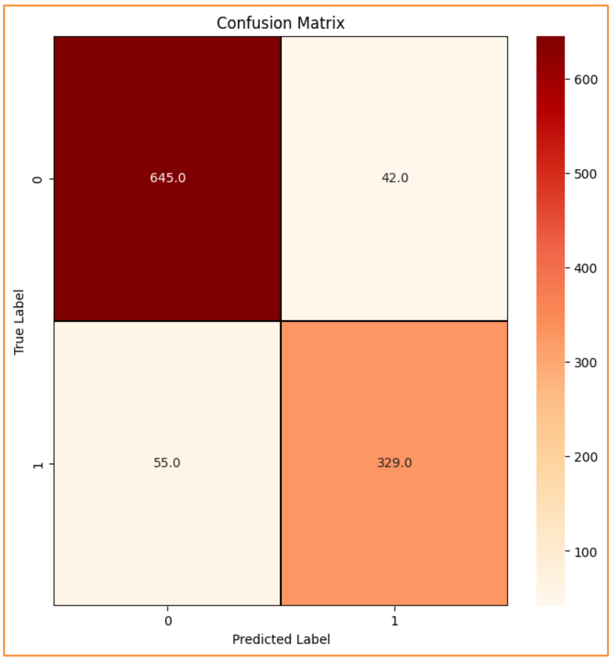

Fig 7: Confusion Matrix Result of VGG-19

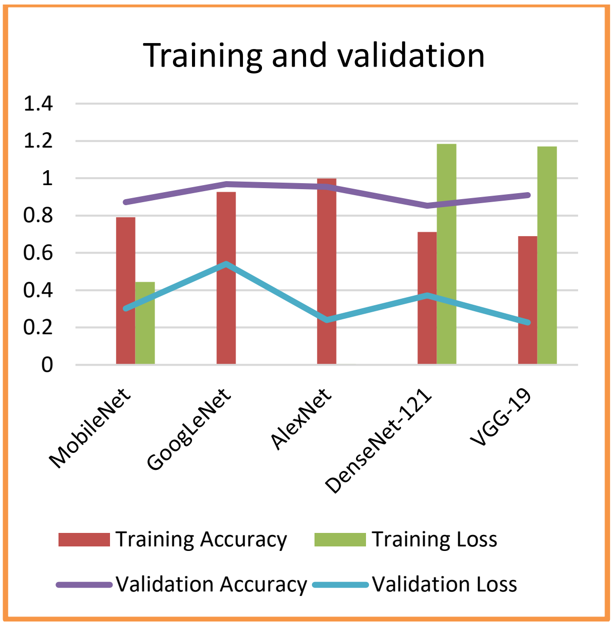

Fig 8: Training and Validation Metric Performances

Table 5. Performance Metrics Result

TL Models	Accuracy	Precision	Recall	F1-Score
MobileNet-V2	0.79	0.84	0.84	0.84
GoogLeNet	0.96	0.99	0.94	0.96
AlexNet	0.96	0.99	0.95	0.97
DenseNet-121	0.83	0.85	0.94	0.89
VGG-19	0.91	0.92	0.94	0.93

Table 6. Training and Validation Results

Training/Validation	MobileNet	GoogLeNet	AlexNet	DenseNet-121	VGG-19
Training Accuracy	0.7906	0.9263	0.9989	0.7128	0.6899
Training Loss	0.4447	0.0008	0.0044	1.1849	1.1701
Validation Accuracy	0.8712	0.9685	0.9545	0.8527	0.9094
Validation Loss	0.3012	0.5410	0.2407	0.3731	0.2271

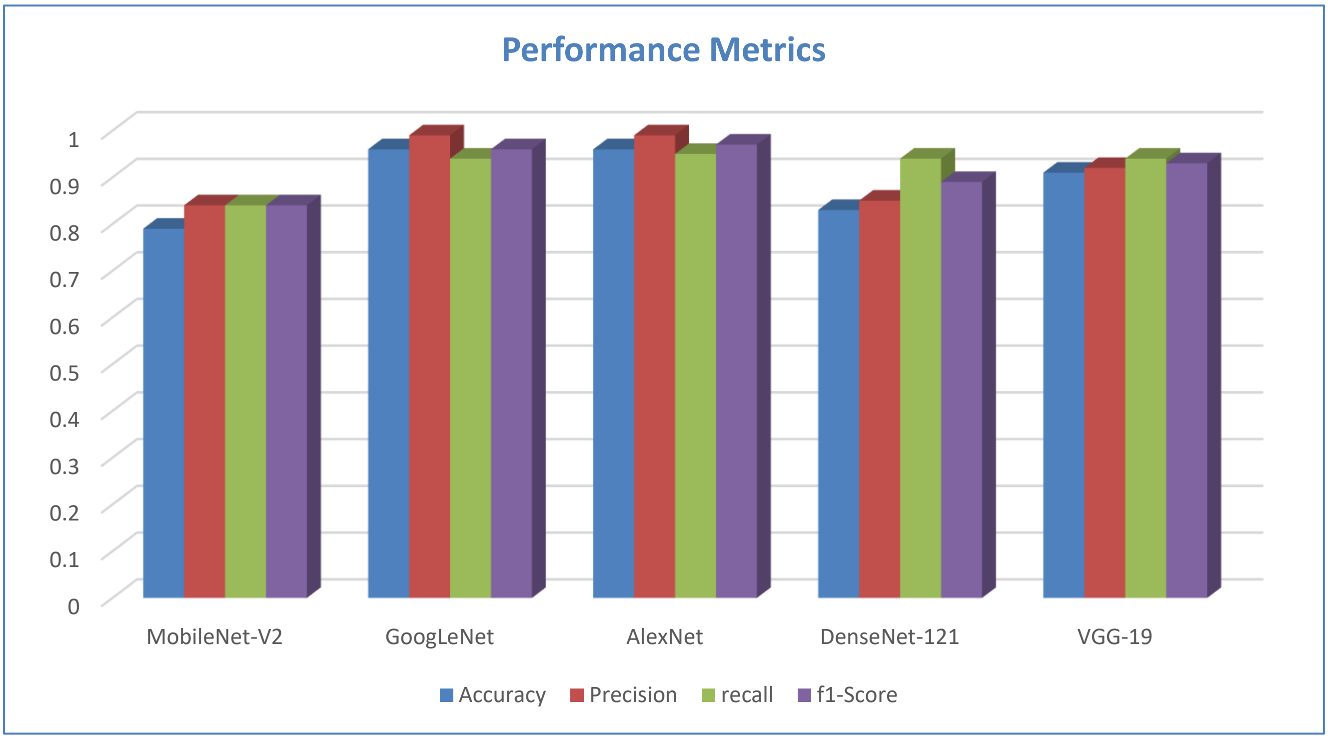

Fig 9: Evaluation Metric Scores Chart

Table 7. Weighted Average Ensemble

S/N	Weighted Ensemble	Score
1	Accuracy	0.9669
2	Precision	0.9446
3	Recall	0.9902
4	F1-Score	0.9304

4. DISCUSSION

The experiment was implemented using Python Programming language with its library and reports of the classification, confusion matrix, accuracy, recall, Fi-score and Precision were recorded. Table 5, present the accuracy the five classifiers. GoogLeNet and AlexNet recording 0.96 and 0.96 accuracy, followed by VGG-19 with 0.91, DenseNet, 0.83 and MobileNet having the lowest which 0.79. Table 7, display the weighted averaging ensemble recall, f1-score, accuracy, and precision of five classifiers to 0.9902, 0.9304, 0.9669 and 0.9446 respectively.

Overall, the research found that each architecture had its own strengths and weaknesses regarding performance for pneumonia classification and prediction. By analysing their performances, the models' results were able to provide insights into the applicability of these architectures for similar image processing tasks.

5. CONCLUSION

Pneumonia is a dangerous infection that can affect people of all ages, but it is especially hazardous for elderly individuals over 70 years old and children under five years old. Early diagnosis and timely treatment are crucial to reducing the mortality rate associated with this disease. In this research, we proposed five transfer learning models to diagnose pneumonia. Four performance indicators were used to assess our system's performance: F1-score, recall, accuracy, and precision. With a training accuracy of 97% and recall of 99%, ensemble classifier produced remarkable results, and our models fared well. Comparing this to the precision attained by current methods and procedures, the difference is substantial. We believe that this framework will be a valuable tool for doctors and radiologist in the diagnosis and treatment of pneumonia.

6. ACKNOWLEDGMENTS

(i)The School of Computing, at The Federal University of Technology Akure, is acknowledged by the authors, for the lecturer's insightful criticism that helped to shape the work.
(ii) Mr Godwin Akinbo, for his consistent encouragement and support in making this work a reality.

Funding: The authors certify that they did not receive any funding for this article's research, writing, or publishing.

7. REFERENCES

[1] Nguyen, T. Q., Nguyen, T. N., Nguyen, T. T., & Nguyen, T. H. 2022. Deep learning for chest X-ray diagnosis: A comparative review. Journal of Medical Imaging and Health Informatics, 12(5), 1045-1053.

[2] Igweonu-Nwakile C. O, Adejimi A.A, Roberts A. A, Oluwole E. O, Oridota O. E, Oyeleye O. A & Onajole A. T. 2023, Prevalence of Pneumonia and Its Determinants among Under-five Children attending a Primary Health Care Clinic in Amuwo Odofin Local Government Area, Lagos, Nigeria. Journal of Community Medicine and Primary Health Care. 35 (1) 40-49

[3] World Health Organisation (WHO). 2022. Pneumonia Children.https://www.who.int/news-room/fact-sheets/detail/pneumonia. Accessed 8th, January, 2024.

[4] UNICEF Nigeria Pneumonia. 2020. https://www.unicef.org/nigeria/press-releases/two-million-children-nigeria-could-die-next-decade-unless-more-done-fight-pneumonia Accessed 5th, January, 2024.

[5] Qin, C., Yao, D. & Shi, Y. Computer-aided detection in chest radiography based on artificial intelligence: a survey. 2018. BioMed Eng,17,113. https://doi.org/10.1186/s12938-018-0544-y____Accessed 19th, March 2024.

[6] Sirazitdinov I, Kholiavchenko M, Mustafaev T, Yixuan Y, Ramil Kuleev R, & Ibragimov B. 2019. Deep neural network ensemble for pneumonia localization from a large-scale chest x-ray database. Computers and Electrical Engineering, 78, 388–399, https://doi.org/10.1016/j.compeleceng.2019.08.004

[7] Janizek J., Erion G., DeGrave A. & Lee S. 2020. An adversarial approach for the robust classification of pneumonia from chest radiographs. Proceedings Of the ACM Conference on Health, Inference, And Learning. 69-79. https://doi.org/10.1145/3368555.3384458

[8] Salehi M, Mohammadi R, Ghaffari H, Sadighi N, Reiazi R. 2021. Automated detection of pneumonia cases using deep transfer learning with paediatric chest X-ray images. Br J Radiol. 2021 May 1;94(1121):20201263. doi: 10.1259/bjr.20201263. Epub 2021. PMID: 33861150; PMCID: PMC8506182.

[9] Shah A & Shah M. 2022. Advancement of deep learning in pneumonia/ Covid-19 classification and localization: a systematic review with qualitative and quantitative analysis. Chronic Dis Transl Med. 8:154-171. doi:10.1002/cdt3.17

[10] Bhatt H, & Shah M. 2023. A Convolutional Neural Network ensemble model for Pneumonia Detection using chest X-ray images. Journal of Healthcare Analytics, 3, 100176

[11] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. & Chen, L.C. 2018 MobileNetV2: Inverted Residuals and Linear Bottlenecks." In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 4510-4520). IEEE.

[12] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D. & Rabinovich, A.2015. Going deeper with convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 2015.1–9.

[13] Gayazahmad. 2020. Comparison And Architecture of Pre-Trained Model AlexNet. 2020 https://medium.com/@gayazahmad62/comparison-and-architecture-of-pre-trained-model-alexnet-22b5be5e6ff6 Accessed 10th, February, 2024.

[14] Arora A.2020. DenseNet Architecture Explained with PyTorch Implementation from Torch Vision. Densely Connected Convolutional Networks. https://amaarora.github.io/posts/2020-08-02-densenets.html Accessed 8th, January, 2024.

[15] Bangar S. 2022. VGG-Net Architecture Explained.https://medium.com/@siddheshb008/vgg-net-architecture-explained-71179310050f Accessed February5th, 2024