

PneumoniaXAttnNet: An Attention-Driven Xception Framework for Accurate Pneumonia Classification with XAI

Tanvir Ahmed
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
2021200000015@seu.edu.bd

Md. Abdur Rahman
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
2021200000025@seu.edu.bd

Md. Delower Hossain
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
2021200000102@seu.edu.bd

Yousuf Howlader
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
2020000000150@seu.edu.bd

Md. Mahi Uddin
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
2021000000009@seu.edu.bd

Md. Tajuddin
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
2021200000005@seu.edu.bd

MD Illeas Hossain
Technology & Computer Science
University of the Potomac
USA
illeas.hossain@student.potomac.edu

Md. Mijanur Rahman*
Computer Science and Engineering
Southeast University
Dhaka, Bangladesh
mijanur.rahman@seu.edu.bd

Abstract—Pneumonia is a considerable worldwide health challenge, demanding the advancement of swift and precise diagnostic techniques. Although deep learning models have demonstrated potential, their clinical adoption is often hindered by a lack of interpretability and an inability to focus on diagnostically salient regions. To overcome these limitations, this paper introduces PneumoniaXAttnNet, a novel deep learning framework that integrates a custom channel-spatial attention module with a fine-tuned Xception backbone to enhance feature representation by concentrating on critical areas within radiographic images. Our model outperforms prior benchmarks with an accuracy of 97.35% and an Area Under the Curve (AUC) of 0.9927 on the publicly available Chest X-Ray Images (Pneumonia) dataset. To visually confirm that the model's predictions are based on clinically significant pathological traits, including regions of lung opacity, we used a suite of explainable AI (XAI) techniques, specifically Grad-CAM, Simple Gradients, and SmoothGrad, in addition to quantitative measurements. By combining high diagnostic accuracy with transparent, verifiable decision-making, PneumoniaXAttnNet represents a significant step towards developing a robust and trustworthy computer-aided diagnostic tool for clinical practice.

Index Terms—Pneumonia Classification, Deep Learning, Convolutional Neural Networks (CNN), Attention Mechanism, Explainable Artificial Intelligence (XAI)

I. INTRODUCTION

Pneumonia continues to be a major global health concern, leading to a high annual rate of illness and mortality. This respiratory infection, which is among the most prevalent, can

cause serious problems, especially in high-risk groups including young children and people with weakened immune systems [1]. The development of medical imaging technologies has significantly advanced disease diagnosis, with tools like chest X-rays and CT scans becoming essential for identifying conditions like pneumonia [2].

The healthcare sector holds a wealth of pneumonia-related data that, when effectively analyzed, can provide valuable insights for early diagnosis and treatment. Various computer-aided diagnostic (CAD) systems now assist radiologists in identifying pneumonia, while data mining methods help in processing complex medical data to uncover meaningful diagnostic patterns.

In recent years, deep learning, a particular domain of machine learning, has gained notable attention for its capacity to recognize complex patterns from extensive datasets through multi-layered neural networks. This method enables the construction of models that capture high-level data features and has shown remarkable success in radiology by surpassing the performance of traditional machine learning techniques in certain applications. In many cases, deep learning models rival or exceed the diagnostic capabilities of human experts [3]. Beyond healthcare, deep learning is widely applied in domains such as speech recognition, image processing, language understanding, graph analytics, and smart transportation systems [4].

The power of deep learning lies in its ability to automati-

cally extract intricate features, deliver highly accurate results which often exceeding human-level performance in medical imaging that support end-to-end learning with little human intervention, and efficiently manage massive datasets, making it ideal for modern big data challenges. Recently, the field has experienced rapid progress, moving from theoretical innovation to real-world deployment. New model architectures and optimization strategies like VGGNet, ResNet and Transformer have significantly improved tasks in both image classification and natural language processing [5].

However, many current deep learning models used for pneumonia diagnosis struggle to effectively focus on the most critical regions of medical images. Additionally, their decision-making procedures are sometimes opaque, which makes it challenging for clinicians to completely trust or understand the outcomes. Because healthcare professionals need models that are both accurate and explicable, this restriction has hampered clinical acceptance.

We address these issues by introducing a new deep learning framework for classifying pneumonia that is based on a modified Xception architecture and enhanced by a specially designed Channel-Spatial Attention (CSA) module. This attention mechanism allows the model to identify both “what” and “where” to focus within an image, leading to improved feature learning and classification accuracy, while offering better interpretability for clinical use.

This research work’s main contributions are:

- We introduce PneumoniaXAttnNet, a new attention-driven model that integrates a Channel-Spatial Attention module into a fine-tuned Xception backbone [6]. This strategy optimizes the model’s attention to critical regions in chest X-rays and improves classification accuracy relative to traditional CNN models.
- We utilize explainable AI (XAI) techniques, such as Grad-CAM [7], Simple Gradients, and SmoothGrad, to elucidate and comprehend the model’s decisions. These strategies emphasize the particular image regions that affect the diagnosis, enhancing clinical confidence and rendering the system more applicable for practical medical use.

II. RELATED WORK

In recent years, there has been a significant rise in the utilization of deep learning in healthcare, especially for the analysis and classification of medical imaging data, including oncological pathology slides, diabetic retinopathy scans, pneumonia chest X-rays, tuberculosis images, and microbial specimens. Joint endeavors among pathologists, radiologists, and computer scientists have resulted in the creation of computer-assisted diagnosis tools for ailments such as pneumonia, cancer, and tuberculosis [8].

Numerous deep learning methodologies have been presented for pneumonia diagnosis utilizing chest X-ray and CT imaging. Gu and Lee [9] introduced a transfer learning framework leveraging general-purpose pretrained models such as ResNet, DenseNet, and MobileNet for pneumonia classification. To

adapt grayscale X-rays to RGB-based models, they replicated the single-channel images across three channels. DenseNet161 achieved 93.6% accuracy, demonstrating that transfer learning from real-world images can effectively overcome limitations of small medical datasets while reducing training time and computational cost.

Similarly, Akbar et al. [10] evaluated multiple neural network models, including custom CNNs and pretrained architectures such as MobileNet, DenseNet, VGG16/19, InceptionV3, ResNet, NASNetMobile, and EfficientNet, utilizing a publicly accessible collection of chest X-ray pictures. They applied data augmentation techniques (rotation, zoom, shift), resizing to 224×224 pixels, and normalization. Employing the Adam optimizer, binary cross-entropy loss, and dropout, EfficientNet-B0 attained superior performance with 94.13% accuracy, 93.50% precision, 92.99% recall, and 93.14% F1-score, demonstrating the efficacy of advanced pretrained CNNs for automated pneumonia diagnosis.

Widiarto et al. [11] proposed an ensemble CNN framework combining InceptionResNetV2 and MobileNetV2, utilized a dataset of around 5,800 chest X-ray pictures. The ensemble model achieved 92.47% accuracy, surpassing the individual models (89.58% for InceptionResNetV2 and 90.38% for MobileNetV2), demonstrating that ensemble strategies can effectively leverage the complementary strengths of multiple architectures.

Shannar el al. [12] introduced a deep learning-based pneumonia classification system utilizing the VGG19 architecture on Kaggle dataset chest X-ray pictures. The study addressed both unbalanced and balanced datasets, where undersampling was applied to mitigate class imbalance. The pretrained VGG19 model was fine-tuned using Adam and stochastic gradient descent optimizers after the system integrated pre-processing operations like resizing, normalization, and dataset balance. Experiments revealed that the model improved pneumonia classification performance with VGG19 with undersampling, achieving 86% accuracy on the imbalanced dataset and 94% accuracy on the balanced dataset.

While these studies demonstrate the promise of deep learning in pneumonia classification, they also reveal notable gaps. Many models are trained and tested on relatively small or imbalanced datasets, limiting generalizability. In addition, few approaches integrate explainability mechanisms, which are essential for clinical adoption. Moreover, real-time applicability and external validation on diverse populations remain underexplored, highlighting the need for robust and deployable solutions.

III. METHODOLOGY

Our approach to pneumonia classification is based on the fine-tuned convolutional neural network (CNN) architecture, enhanced with an integrated attention mechanism to improve localization of salient pathological features. This approach, which we term PneumoniaXAttnNet, is designed to accurately classify chest radiographs while providing insights into its decision-making process. The general structure is depicted in

Fig. 1. This section details the dataset preparation, the specific components of our model architecture, and the explainability techniques employed.

A. Dataset and Preprocessing

We used the publicly available Chest X-Ray (Pneumonia) data set, which included 5,863 anterior-posterior chest radiographs from children between the ages of one and five years [13]. The dataset is categorized into two groups: *Normal* and *Pneumonia*. The photos were first gathered as part of standard clinical care and underwent quality control assessment; expert doctors graded the diagnoses to verify that the labels were correct [13].

For robust and reproducible evaluation, we deviated from the original dataset splits. All images were first combined into a single set. We used stratified partitioning, allocating 20% of the data for a reserved test set, 10% for validation, and 70% for training. This procedure ensures there is no data leakage from the original test set into our training or validation sets.

To conform to the input dimensions of our network, all photos had their sizes reduced to 224×224 pixels. To enhance model generalization, real-time data was included to the training set. The augmentations comprised random rotations of up to 20 degrees, random horizontal flips, and color jittering with brightness and contrast factors of 0.3. For the validation and test sets, only resizing and normalization were applied. All input photos were normalized using the mean and standard deviation obtained from the ImageNet dataset, in accordance with conventional procedures for models utilizing pre-trained weights.

B. Model Architecture: PneumoniaXAttnNet

Our model architecture, PneumoniaXAttnNet, comprises of a pre-trained feature extractor, a novel attention module, and a custom classification head, which are detailed below.

1) *Fine-Tuned Xception Backbone*: The foundation of our model is the Xception architecture [6], which we use as a feature extractor. Weights pre-trained on the ImageNet dataset were used to initialize the network to make use of the rich, general-purpose features learned from a large-scale image corpus. We utilized a selective fine-tuning method to adjust the model for the particular area of medical imaging. The weights of the initial and intermediate layers were fixed, maintaining their capacity to identify fundamental patterns such as edges and textures. We unfroze only the final two separable convolution blocks (`block10` and `block11`) and the initial stem layers. This enables the model to modify its advanced feature representations to the nuanced patterns indicative of pneumonia on radiographs. The original fully-connected classification layer was removed to accommodate our custom modules.

2) *Channel-Spatial Attention Module*: To improve the model’s capacity to concentrate on diagnostically important regions, we incorporated a tailored attention mechanism positioned directly after the Xception backbone. Drawing inspiration from the Convolutional Block Attention Module

(CBAM), this component refines the feature maps by applying attention sequentially across two distinct dimensions. The channel attention component first determines *what* is important by using global average pooling and a small multi-layer perceptron (MLP) to compute channel-wise weights. These weights recalibrate the feature map to emphasize informative channels. Subsequently, the spatial attention unit identifies *where* the model should concentrate by employing a 7×7 convolutional layer to produce a two-dimensional spatial attention map. This map is applied to the feature representation, attenuating irrelevant background information while enhancing the prominent pathological regions.

3) *Classification Head*: The attended feature maps are reduced to a fixed-size feature vector via global average pooling. This vector is then fed into a custom classification head designed for robust prediction. The head is composed of two dense blocks, each containing a linear layer, a `BatchNorm1d` layer for training stability, a Swish-like activation function (SiLU), and a `Dropout` layer (with rates of 0.5 and 0.3, respectively) for regularization. The final linear layer outputs the raw logits for the binary classification.

C. Loss Function

For training, the network was optimized using the standard cross-entropy loss, which is widely adopted for classification problems. Considering a batch of M samples, the loss function $\mathcal{J}$ can be expressed as:

$$\mathcal{J} = -\frac{1}{M} \sum_{j=1}^M \sum_{k=1}^K t_{j,k} \log(\hat{p}_{j,k}) \quad (1)$$

In this formulation, K represents the total number of categories (2 in our case), $t_{j,k}$ is a binary indicator (equal to 1 if instance j belongs to class k , and 0 otherwise), and $\hat{p}_{j,k}$ denotes the predicted probability, obtained through the softmax function applied to the network outputs, that instance j is assigned to class k . The model weights are iteratively updated by minimizing this loss with the AdamW optimizer.

D. Explainable AI for Model Interpretation

In a clinical context, the transparency of an automated system is paramount. To interpret our model’s predictions, we employed several gradient-based explainability techniques to generate saliency maps. We predominantly employed Gradient-weighted Class Activation Mapping (Grad-CAM) [7], which generates a class-discriminative localization map by examining the gradients that propagate into the final convolutional layer. To conduct a more detailed study, we calculated saliency maps utilizing basic input gradients, a fundamental technique that directly depicts the gradient of the class score in respect to the input pixels. To mitigate the noise inherent in these raw gradient maps, we applied SmoothGrad, which averages the gradients over multiple noise-perturbed copies of the input image. The resulting visualizations allow for a qualitative assessment of whether the model identifies clinically relevant indicators of pneumonia, such as lobar consolidations, rather than learning from confounding artifacts.




Fig. 1: The Proposed PneumoniaXAttnNet Framework

IV. RESULTS AND DISCUSSION

This section provides a thorough assessment of our proposed PneumoniaXAttnNet. We detail the experimental setup, provide an extensive quantitative research of the model’s effectiveness on the reserved test set, analyze our results against previous methods, and conduct a qualitative analysis using the explainability techniques described in our methodology.

A. Experimental Setup

The PyTorch deep learning framework was used to implement it on a computer with an NVIDIA Tesla P100 GPU. Training lasted up to 20 epochs, and the AdamW optimizer was set up with a weight decay factor of 1×10^{-3} and a starting learning rate of 1×10^{-4} . During training, the learning rate was modified using a cosine annealing scheduler to further enhance convergence. To mitigate overfitting and select the best-performing model, we applied a mechanism called early stopping with 5 epochs of patience, monitoring the validation loss. The total training time to convergence was approximately 10.7 minutes (644 seconds). Subsequent inference on the entire held-out test set, which comprised 1,172 images, was completed in 12.3 seconds, highlighting the model’s computational efficiency.

B. Quantitative Performance Analysis

We assessed the model’s performance utilizing a range of conventional measures, emphasizing its training behavior, classification accuracy on the test set, and discriminative capability.

1) *Training and Validation Dynamics*: Fig. 2 presents the learning curves obtained across the training epochs. The training loss decreases steadily in a smooth manner, while the validation loss reflects a similar pattern, suggesting that the model achieved good generalization rather than simply memorizing the training set. Similarly, the validation and training accuracy curves rise in tandem, both achieving a plateau above 97%. The close proximity of the validation and training curves suggests that our data augmentation and regularization strategies were effective in preventing significant overfitting.



Fig. 2: Model training and validation curves. The close alignment between accuracy (left) and loss (right) curves indicates stable convergence and effective generalization.

2) *Test Set Performance Metrics*: Table I provides the classification report that summarizes the model’s results on the independent test set. The proposed PneumoniaXAttnNet achieved an impressive overall accuracy of 97.35%. Specifically, for the clinically significant *Pneumonia* category, the model reached a precision of 98.00%, a recall of 98.30%, and an F1-score of 98.20%. The high recall is particularly important in a diagnostic setting, as it signifies the model’s capacity for accurate identification the vast majority of true pneumonia cases, thereby minimizing false negatives. The high precision means that the model is very likely to be accurate when it predicts pneumonia.

TABLE I: Performance metrics of PneumoniaXAttnNet on the held-out test set.

Class	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
NORMAL	97.35	95.71	94.83	95.27
PNEUMONIA	97.35	97.99	98.34	98.16
Macro Avg	-	96.85	96.59	96.72
Weighted Avg	97.35	97.35	97.35	97.35

3) *Error Analysis via Confusion Matrix*: A more detailed view of the classification performance is presented in the confusion matrix shown in Fig. 3. Among the 843 actual pneumonia cases in the test set, the model correctly predicted 829 as positive (True Positives), while 14 cases were incorrectly labeled as normal (False Negatives). For the 329 true normal instances, 312 were accurately classified (True Negatives), and only 17 were mistakenly identified as pneumonia (False

Positives). This low number of misclassifications, particularly the small number of false negatives, further validates the model’s reliability for this diagnostic task.




Fig. 3: Confusion matrix of the proposed model on the test set, detailing correct and incorrect predictions for each class.

4) *Discriminative Power Analysis*: To assess the model’s ability to distinguish between the two categories in all possible classification thresholds, we constructed the Receiver Operating Characteristic (ROC) curve, illustrated in Fig. 4. The curve shows a low false positive rate and a high true positive rate as it increases dramatically toward the top-left corner. The computed Area Under the Curve (AUC) was an impressive 0.9927, which is extremely near to the optimal value of 1.0. This high AUC value demonstrates the remarkable discriminative power of the model.




Fig. 4: The ROC curve for the test set shows an AUC of 0.9927, highlighting the model’s strong discrimination between classes.

C. Ablation Study

To determine the impact of the suggested Channel-Spatial Attention (CSA) module, we performed an ablation study. We compared the performance of our full PneumoniaXAttnNet model against a baseline counterpart, which uses the identical fine-tuned Xception backbone and classification head but omits

the attention module. Both models were trained and evaluated under the same experimental conditions to ensure a direct and fair comparison.

The results, summarized in Table II, confirm the significant contribution of the attention mechanism. While the baseline XceptionNet performs well, the introduction of the CSA module boosts performance across all key metrics. Notably, accuracy improves from 95.65% to 97.35%, and the Macro F1-Score rises from 94.48% to 96.72%. This enhancement indicates that by directing the model’s focus toward the most important features, the attention module contributes to more reliable and accurate classification, supporting its integration into our architecture.

TABLE II: Ablation study results comparing the baseline Xception against the proposed PneumoniaXAttnNet.

Model	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)
Xception	95.65	95.67	93.45	94.48
PneumoniaXAttnNet	97.35	96.85	96.59	96.72

D. Qualitative Analysis and Model Interpretability

Fig. 5 provides visual explanations for the model’s predictions on representative test images using Grad-CAM [7], simple gradients, and SmoothGrad. For the *Pneumonia* case (bottom row), the Grad-CAM visualization correctly highlights the region of lung opacity in the upper right lobe, a critical clinical marker. This indicates that the model is effectively identifying clinically relevant features. For the correctly classified *Normal* case (top row), the model’s attention is more diffuse and centered on expected anatomical structures like the cardiac silhouette and hilum, without focusing on any specific anomalous region. These visualizations provide strong evidence that the model’s decision-making process is grounded in meaningful visual patterns.




Fig. 5: The figure shows the original image alongside heatmaps generated by Grad-CAM, Simple Gradients, and SmoothGrad for a correctly classified ‘Normal’ case (top row) and a ‘Pneumonia’ case (bottom row).

E. Discussion

The experimental findings indicate that the proposed PneumoniaXAttnNet serves as a highly effective approach for detecting pneumonia in pediatric chest X-rays. In this section,

we analyze the results, explore their clinical implications, and highlight the study's limitations along with potential directions for future research.

The quantitative success, particularly the high recall (98.3%) and AUC (0.9927), is significant. We attribute this strong performance primarily to our architectural design. By using a pre-trained Xception network, we leveraged a robust set of features learned from a massive dataset. However, the key innovation is the subsequent channel-spatial attention module. This module forces the model to explicitly learn and weight the importance of different feature channels ('what' to look for) and spatial locations ('where' to look), effectively guiding its focus towards diagnostically relevant regions and away from confounding background information. This targeted feature refinement is likely what gives our model an edge over more generic CNN architectures.

From a clinical perspective, the implications of this work are promising. The model's high accuracy and, more importantly, its high recall for pneumonia, suggest its potential as a reliable second reader or a screening tool for radiologists, especially in high-workload environments. Its rapid inference time makes it practical for real-time applications. Furthermore, the interpretability provided by the XAI visualizations is crucial for building clinical trust. Clinicians can view not only the model's prediction but also an accompanying visual explanation, enabling them to confirm whether the model's attention corresponds with their own clinical assessment.

Despite these encouraging findings, we recognize certain limitations. Firstly, the model was trained and tested on a single, though sizeable, publicly available dataset of pediatric patients. Its performance on adult patients or on images from different medical centers with different imaging equipment remains to be validated. Second, the current model performs a binary classification. The model is unable to differentiate among various types of pneumonia (such as bacterial versus viral), which represents a clinically significant distinction.

Future work will focus on three primary directions. The first is to assess the model's generalizability using external, multi-center datasets that also include adult patients. The second is to expand the framework to a multi-class classification task capable of distinguishing between bacterial and viral pneumonia. Finally, we plan to investigate combining our image-based model with additional clinical information from electronic health records (EHR) to develop a comprehensive, multi-modal diagnostic system.

V. CONCLUSION

This study introduces PneumoniaXAttnNet, a novel deep learning framework designed to detect pneumonia in pediatric chest X-ray images. By integrating a channel-spatial attention module with a fine-tuned Xception backbone, our model learns to focus on the most salient diagnostic regions, significantly improving classification performance. Evaluation on a large public dataset demonstrates the effectiveness of our approach, achieving a high accuracy of 97.35% and an outstanding AUC of 0.9927. Crucially, we paired this high accuracy with robust

interpretability. Through the use of XAI techniques like Grad-CAM, we confirmed that the model's predictions rely on clinically meaningful features, such as areas of lung opacity, thereby fostering the trust necessary for clinical adoption. The model's computational efficiency further supports its potential for real-world deployment as a reliable diagnostic aid for radiologists. The primary limitation of this study is its reliance on a single pediatric dataset. Consequently, future efforts will aim to evaluate the model on diverse external datasets, including adult populations, and to extend its capability to distinguish between bacterial and viral pneumonia.

REFERENCES

- [1] C. Troeger *et al.*, "Estimates of the global, regional, and national morbidity, mortality, and aetiologies of lower respiratory tract infections in 195 countries: a systematic analysis for the global burden of disease study 2015," *The Lancet Infectious Diseases*, vol. 17, no. 11, pp. 1133–1161, 2017.
- [2] W. Khan, N. Zaki, and L. Ali, "Intelligent pneumonia identification from chest x-rays: A systematic literature review," *IEEE Access*, vol. 9, pp. 51 747–51 771, 2021.
- [3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [4] M. R. Minar and J. Naher, "Recent advances in deep learning: An overview," *arXiv preprint arXiv:1807.08169*, 2018.
- [5] A. Torfi, R. A. Shirvani, Y. Keneshloo *et al.*, "Natural language processing advancements by deep learning: A survey," *arXiv preprint arXiv:2003.01200*, 2020.
- [6] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1251–1258.
- [7] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [8] T. Ching, D. S. Himmelstein, B. K. Beaulieu-Jones *et al.*, "Opportunities and obstacles for deep learning in biology and medicine," *Journal of the Royal Society Interface*, vol. 15, no. 141, p. 20170387, 2018.
- [9] C. Gu and M. Lee, "Deep transfer learning using real-world image features for medical image classification, with a case study on pneumonia x-ray images," *Bioengineering*, vol. 11, no. 4, p. 406, 2024.
- [10] W. Akbar, A. Soomro, A. Hussain, T. Hussain, F. Ali, M. I. U. Haq, R. W. Attar, A. Alhomoud, A. A. Alzubi, and R. Alsagri, "Pneumonia detection: A comprehensive study of diverse neural network architectures using chest x-rays," *International Journal of Applied Mathematics and Computer Science*, vol. 34, no. 4, pp. 679–699, 2024.
- [11] M. Widiarto, P. N. Andono, C. Supriyanto, M. A. Soeleman, and A. Wijaya, "Ensemble cnn models for accurate pneumonia classification in chest x-ray images," in *2024 International Conference on Computer Engineering, Network, and Intelligent Multimedia (CENIM)*. IEEE, 2024, pp. 1–6.
- [12] S. Phine, "Pneumonia classification using deep learning vgg19 model," in *2023 IEEE Conference on Computer Applications (ICCA)*. IEEE, 2023, pp. 67–71.
- [13] D. S. Kermay, M. Goldbaum, W. Cai, C. C. Valentim, H. Liang, S. L. Baxter, A. McKeown, G. Yang, X. Wu, F. Yan *et al.*, "Identifying medical diagnoses and treatable diseases by image-based deep learning," *cell*, vol. 172, no. 5, pp. 1122–1131, 2018.