

Detection Of Chest Diseases Using EfficientNet-B0

Varun V¹, Adithya M², Nagatejas P³, Nethra Moorthi Hosahalli⁴, Emmanuel Didymus Sebastian⁵ & Dr. Gowthul Alam⁶

^{1 2 3 4 5 6}Department of Computer Science and Engineering, JAIN (Deemed-to-be-University)

Abstract - Chest X-rays are vital diagnostic tools in identifying thoracic diseases, but interpreting them manually can be time-consuming and prone to variability. With the advancement of deep learning, automated multi-label classification has become increasingly feasible for aiding clinical diagnosis. In this research, we investigate the performance of EfficientNetB0, a lightweight convolutional neural network, for classifying 14 chest diseases using the ChestMNIST dataset. The dataset comprises grayscale X-ray images, and the model is trained from scratch without using pre-trained weights to better adapt to the grayscale nature of the input. This research showcase that that EfficientNetB0 can effectively handle multi-label classification in chest radiography, even without leveraging transfer learning, making it a promising candidate for real-world clinical applications. The model provides a high accuracy and promising Area Under the Curve (AUC) value.

Key Words: ChestMNIST, AUC, EfficientNetB0, Chest Disease, Transfer Learning

1.INTRODUCTION

All Contemporary healthcare heavily relies on medical imaging techniques to peer inside the human body and precisely pinpoint health problems. Chest X-rays are exceptionally common within this field, frequently employed to assess heart and lung health because they are relatively inexpensive, readily accessible, and quick to perform, proving indispensable for both emergency situations and routine check-ups. Despite their utility, interpreting these images requires specialized skill, as diagnosticians must identify often faint or ambiguous signs that could point to numerous conditions, leading to potential inconsistencies and mistakes between readers. Implementing automated analysis systems for chest conditions thus offers a significant opportunity to boost the speed and reliability of diagnoses, especially benefiting medical facilities that handle large numbers of patients or operate with limited specialist availability[13].

The evolution of artificial intelligence, especially deep learning techniques, has significantly transformed the interpretation of medical images. Convolutional Neural Networks (CNNs), a specific type of AI demonstrating strong capabilities in recognizing visual patterns, have found widespread application in medicine, assisting in diagnosing conditions like diabetic eye damage, classifying skin lesions, and detecting diseases in the chest. When analyzing chest X-rays, these CNNs have shown they can identify multiple different health issues within a single image, a capability known as multi-label classification. This is particularly valuable because various chest ailments frequently exist together or produce similar visual evidence on X-rays. Therefore, focusing on building and training CNNs adept at this multi-label identification task can greatly improve the efficacy of automated diagnostic processes.

For this project, we selected EfficientNetB0, an advanced Convolutional Neural Network known for striking an effective equilibrium between high performance and computational frugality by simultaneously adjusting its depth, width, and image resolution scaling. Unlike more computationally demanding networks such as ResNet or DenseNet, EfficientNetB0 achieves similar accuracy levels with a markedly reduced number of parameters, positioning it as an excellent choice for applications where computing resources are restricted. Significantly, our approach involved training the model entirely from the beginning, without initializing it with knowledge gained from the common ImageNet dataset. This decision was made because our input data, the grayscale medical images of ChestMNIST, differs fundamentally from ImageNet's collection of primarily natural, color photographs; applying pre-trained knowledge across such disparate data types often impedes the model's learning process on the target task.

To improve how well our model performs on new data and prevent it from merely memorizing the training examples, we incorporated a mix of data augmentation approaches and regularization controls, such as dropout and stopping the training process early if validation performance plateaued. We partitioned the dataset

carefully, using 70% for training the model, 10% for validating its progress during training, and holding out 20% for final testing, ensuring that the proportion of different conditions was similar across these splits. During the model's development, we employed binary cross-entropy to measure error and tracked performance on the validation set to identify and save the best-performing model version. Our final assessment considered not just accuracy and loss figures but also Area Under the Curve (AUC) scores, which are particularly useful for evaluating multi-label classification tasks where some conditions might be less common. The model ultimately showed strong and consistent results, reaching 95.03% accuracy on the training data, 94.93% on the validation data, and 94.82% on the unseen test data, indicating reliable performance and good generalization.

To gain deeper insight beyond overall accuracy, we performed a more detailed assessment using Receiver Operating Characteristic (ROC) curves generated independently for each disease category, along with their calculated Area Under the Curve (AUC) scores. This method provides a finer-grained view of the model's capabilities by illustrating its trade-offs between correctly detecting each specific condition (sensitivity) and correctly identifying its absence (specificity). While the AUC results differed among the individual disease types, we computed a macro-average AUC of 0.76, which considers each class equally and suggests a solid level of diagnostic performance when averaged across all the conditions studied.

A key difficulty in classifying medical images involves managing datasets where certain conditions are significantly less common than others, a situation known as class imbalance. While we explored the use of specific weights for different classes to address this, our initial attempts led to inconsistent training patterns and reduced the model's effectiveness on new data. Consequently, we chose not to apply explicit class weights during training, relying instead on our data augmentation and regularization methods to foster a more balanced learning environment implicitly. This proved to be a successful decision, enabling the model to reach commendable AUC performance levels for both rare and common disease categories. Additionally, avoiding class weights simplified the overall training process, making the model's reproduction and implementation more straightforward[14].

A. Dataset

A compilation of chest radiographic images, known as ChestMNIST, serves as a significant tool for the progression of studies aimed at identifying multiple thoracic illnesses simultaneously. The ChestMNIST is based on the NIH-ChestXray14 dataset[2]. This collection, extracted from X-ray scans, includes 112,120 individual images, organized into training, validation, and testing subsets, containing 78,468, 11,219, and 22,433 images, respectively. Each image is annotated with the presence or absence of up to 14 distinct pathological conditions. The use of binary labels for each condition facilitates multi-label classifications, enabling the prediction of whether each specific illness is present or absent within a given image[1].

2. Literature Survey

I. Allaouzi et al research introduces a unique method for automatically identifying multiple thoracic diseases in chest X-ray images by intelligently merging Convolutional Neural Networks for image understanding with problem transformation strategies for handling multiple labels. The study rigorously tested the DenseNet-121 model for feature extraction on the challenging ChestX-ray14 and CheXpert datasets. Among different classification approaches like Binary Relevance, Label Powerset, and Classifier Chains, the Binary Relevance technique demonstrated superior performance, achieving an average AUC of 0.882 on ChestX-ray14 and 0.812 on CheXpert, surpassing existing state-of-the-art methods. This innovative approach holds significant potential for enhancing the accuracy and reliability of Computer-Aided Diagnosis systems, thereby providing valuable support to radiologists in medical image interpretation, particularly in scenarios with imbalanced data[6].

S. Ms. Kavitha et al research explores a Convolutional Neural Network model designed to categorize chest conditions evident in X-ray images. The image collection used contains 1,120 X-rays distributed across six diagnostic categories. To enhance the model's diagnostic capabilities, several image preparation strategies were implemented, with Contrast Limited Adaptive Histogram Equalization proving most effective, resulting in a top accuracy of 91.22%. The model training with this preprocessing took 170.816 seconds. The study contrasts the diagnostic precision and training duration achieved with different preprocessing steps, highlighting the capability of CNNs in identifying chest diseases. Subsequent work will focus on further increasing the accuracy and improving the model's efficiency[4].

I. F. Jassam et al study introduces a deep learning system employing a Dense-Net architecture for identifying pneumonia in pediatric chest X-ray images. Utilizing a collection of X-rays from young patients (1-5 years old) at Guangzhou Women and Children's Medical Centre, the model integrates Dense-Net with supplementary convolutional layers and a global average pooling layer to boost its diagnostic capabilities. Strategies such as overlapping pooling and Decoupled Weight Decay Regularization are applied to reduce overfitting. The model achieves a 90% accuracy rate, along with an F1-score of 93% and an AUC of 88.4%, offering a valuable and rapid tool for radiologists in the initial detection of pneumonia and improving the workflow for chest X-ray assessments. This approach contributes to the progress of automated medical image analysis[9].

J. F. V. Oraño et al research investigates the use of a convolutional neural network to categorize lung conditions based on chest X-ray images. The model was trained using 8,125 images representing six distinct diseases: atelectasis, effusion, infiltration, mass, nodule, and pneumothorax. The achieved performance metrics were a training accuracy of 98.46%, a validation accuracy of 83.99%, and a total accuracy of 82.53%. The CNN structure comprises four convolutional layers followed by max pooling and fully connected layers employing ReLU and softmax activation functions. This model was incorporated into computer and mobile applications to support radiologists in their diagnostic process for lung diseases. While the outcomes are encouraging, the study suggests further enhancements to improve the model's accuracy and its ability to generalize to new data[3].

A. Djoudi et al introduces a novel approach for identifying lung diseases in chest X-ray images. By integrating deep CNNs with Genetic Algorithms, the model refines the selection of important image features and optimizes crucial settings to enhance its diagnostic capabilities. Evaluated on the COVID-19[8] Radiography dataset, this method achieved accuracy rates of 77.31%, 78.23%, and 78.87% across various data subsets, demonstrating improvements of 0.51%, 1.56%, and 1.7% over conventional techniques. The study acknowledges the inherent difficulties in analysing chest X-rays for disease detection, such as the visual similarities between different conditions and the constraints of available datasets. The findings indicate that this GCNN method can improve the precision of diagnoses and holds considerable potential for advancing medical image analysis and disease identification[7].

A. Chatchaiwatkul et al research explores a deep learning technique employing the VGG16 architecture to identify and categorize lung illnesses, specifically COVID-19[8], Pneumonia, and Pneumothorax, using chest X-ray images. The model underwent training on extensive, carefully curated datasets obtained from Kaggle and was assessed using key performance indicators such as accuracy, precision, recall, and F1-score. The achieved outcomes were notable, with classification accuracy ranging from 93% to 100% across different scenarios. The study underscores the efficacy of convolutional neural networks in the realm of medical imaging and proposes the development of an online platform to offer complimentary access to this diagnostic capability for wider application[11].

Y. Huang et al research presents a novel Category-wise Residual Attention Learning (CRAL) structure designed for the complex task of multi-label chest X-ray classification to better identify various thorax conditions. CRAL employs Residual Attention Mechanisms to diminish the impact of unimportant image details while amplifying those specific to each disease category, combining modules for feature representation and attention learning. Evaluated on the Chest X-ray14 dataset, the framework achieved a strong average AUC of 0.816, surpassing the performance of previously established models. The underlying network architectures utilized were ResNet-50 and DenseNet-121, and two different attention mechanisms (att1 and att2) were investigated. The findings suggest that att1 yields marginally superior results when paired with ResNet-50, whereas both attention mechanisms show similar effectiveness with DenseNet-121. This study emphasizes the capability of CRAL to enhance the accuracy of identifying multiple diseases in chest X-rays[5].

M. Irtaza et al research investigates the use of deep learning for classifying multiple lung conditions, comparing the effectiveness of various models including MobileNet, DenseNet, VGG-16, EfficientNet, Xception, and InceptionV3 on the NIH Chest X-ray dataset. Among these, MobileNet demonstrated the strongest performance, reaching 93.4% binary accuracy, 57% recall, a 0.553 F1-score, and an AUC of 81. Notably, the F1-score of MobileNet was improved by 5% through the incorporation of synthetic data generated using deep convolutional GANs. Although the GAN-augmented data initially showed lower individual performance, the study underscores their promise in tackling the issue of unevenly distributed disease categories. The research puts forward a transfer-learning model that integrates

MobileNetV1 with a three-layer classifier and employs geometric image augmentation to improve classification outcomes. The study recommends continued efforts to enhance the model's ability to generalize to new data and to optimize its computational demands[10].

3. Proposed Work

Our research adhered to a methodical sequence, starting with data gathering, followed by preparation steps, model construction, the training phase, and finally, performance evaluation. We utilized the ChestMNIST dataset, part of the larger MedMNIST collection known for offering standardized medical imaging resources suitable for benchmarking. We specifically configured the data loading to yield high-resolution 224×224 grayscale X-ray images; this approach aimed to retain more intricate visual information vital for spotting subtle disease indicators. The ChestMNIST data includes annotations for 14 different conditions using a multi-label system, meaning each radiograph can be associated with one or several pathologies. For dividing the data into training, validation, and test sets, we adopted the splits already defined within the ChestMNIST resource itself to ensure uniformity and facilitate dependable comparisons with other research efforts.

Our data preparation process involved adapting the images for use with deep learning models, primarily through normalization and reshaping. To promote stable and efficient training, we scaled the image pixel intensities to a standard range between 0 and 1 by dividing each value by 255. As the input images were grayscale, we modified their structure by adding a channel dimension, transforming them into the (224, 224, 1) format expected by the TensorFlow library. The disease labels, initially represented as binary vectors indicating which conditions were present, were converted to the float32 numerical type for better computational performance and smooth integration with TensorFlow functions. In line with earlier experiments showing that class weighting caused erratic training and poorer performance on new data, we intentionally did not employ any class balancing or weighting techniques in this phase.

Our deep learning framework was constructed around EfficientNetB0, a convolutional neural network recognized for its computational efficiency derived from a balanced scaling of its dimensions and input resolution. We deliberately started the training process from random initial values (weights=None), bypassing the use of standard ImageNet pre-trained parameters because they are optimized for color photographs rather than the

grayscale medical imagery used here. The architecture sequentially included the main EfficientNetB0 network, followed by a Global Average Pooling operation, then a densely connected layer with 128 nodes utilizing ReLU activation, and finally, an output layer consisting of 14 dense nodes with sigmoid activation functions to appropriately manage the multi-label classification nature of the task. Throughout training, the weights within the base EfficientNetB0 model were allowed to update, enabling the entire system to learn features specific to the medical domain directly from the data, beginning from its uninitialized state.

The neural network's setup involved employing the Adam method for refining its internal adjustments, utilizing a rate of 0.001, a selection rooted in its capacity to dynamically adjust learning and its stable behavior observed in initial trials. To quantify the disparity between the network's outputs and the desired outcomes, the binary cross entropy method was employed, a typical choice for scenarios where multiple independent classifications are made. Throughout the learning process, two performance indicators were monitored: the precision of individual label assignments and the area under the receiver operating characteristic curve, which assesses the model's ability to distinguish between categories at various decision points, evaluated separately for each category. The learning process spanned 25 complete cycles, with data processed in segments of 32 instances.

To bolster reliability and avoid excessive adaptation to the training examples, the learning procedure incorporated a pair of monitoring mechanisms. A mechanism was in place to automatically cease the training process if the error on a separate validation set showed no reduction over five successive complete passes through the data. Additionally, the specific set of learned parameters that yielded the smallest error on this validation set was preserved. These tactics were implemented to guarantee the resulting model could effectively handle new, unseen data without mirroring the training data's peculiarities too closely. The parameter configuration that performed optimally on the validation data was subsequently assessed using a distinct dataset reserved for final verification. This assessment involved quantifying the overall error, the area under the receiver operating characteristic curve, and the correctness of predictions on this unseen data, thereby offering a comprehensive understanding of the model's capacity to generalize to real-world scenarios.

The Fig-1 represents the flow diagram of the proposed method. A 224x224 grayscale chest X-ray initiates the process, entering an initially untrained EfficientNetB0 core. This core progressively extracts visual traits through layered filtering and normalization. Spatial data is then condensed into a linear form, processed by a 128-unit ReLU-activated layer for complex pattern recognition, with random connection deactivation to prevent overfitting. The final stage yields 14 sigmoid-activated outputs, each providing an independent likelihood for a specific disease, enabling simultaneous multi-label diagnosis.




Fig -1: Framework of proposed method

Following the learning phase, the model's proficiency was depicted through graphical representations of its learning progression and an examination of its ability to discriminate between categories. Diagrams illustrating the error and the two-performance metrics (individual label correctness and area under the ROC curve) on both the training and a separate validation dataset across each learning cycle were produced to analyze patterns and consistency in learning. For each of the 14 diagnostic categories, curves illustrating the balance between correctly identifying positive cases and incorrectly identifying negative cases as positive were created, enabling a visual evaluation of the model's discernment for each condition. A summary curve, representing the average of the rates of correctly identified positive cases across all categories, was also generated to provide an overall, balanced view of the model's effectiveness. This overall summary measure indicated an area under the

curve of 0.76. On a final, unseen evaluation dataset, the overall area under the curve was measured at 0.835, and the overall correctness of predictions was 0.9482.

The entire investigative process was executed using the Python programming language, with TensorFlow and Keras serving as the foundational tools for constructing and training the deep learning model. Visual representations were generated using the Matplotlib library, and records of the model's learning progress were maintained to ensure the experiments could be replicated and further scrutinized. By employing a design known for its efficiency and scalability, specifically EfficientNetB0, in conjunction with resilient training techniques and detailed visual information, this approach establishes a workable and resource-conscious procedure for categorizing multiple chest conditions from medical images.

4. Results and Discussions

The model was trained using the ChestMNIST image collection, adapted to 224x224 grayscale format for the EfficientNetB0 framework. Throughout this phase, the discrepancy between predictions and actual labels was quantified, and the learning was automatically halted if improvement ceased, with the top-performing model configuration being saved to avoid over-specialization. The concluding evaluation demonstrated a Area Under the Receiver Operating Characteristic Curve of 0.839, signifying a robust ability to distinguish between various disease categories. Employing a sigmoid activation in the final layer facilitated the effective identification of multiple concurrent conditions within a single X-ray.

Table -1: Comparison of AUC with Other Benchmarks

	Model	AUC
Z. Li et al [15]	ResNet – 50	0.755
Y. Guan et al [5]	DenseNet - 121	0.816
Proposed Method	EfficientNet - B0	0.835

The Table-1, is a comparison of the proposed method with state-of-the-art models, showing that the EfficientNet-B0 model achieves the highest AUC of 0.839, outperforming ResNet-50 (0.755) and DenseNet-121 (0.816).




Fig -2: ROC curve of each class

The visual representation as in Fig-2, of the model's diagnostic capability for each of the 14 categories offered deeper understanding of its functioning. The aggregated curve, averaging performance across all categories, provided a unified perspective on the model's overall effectiveness and affirmed its capacity to generalize across the spectrum of conditions.




Fig -3: Loss analysis of proposed method




Fig -4: AUC analysis of proposed method




Fig -5: Accuracy analysis of proposed method

The performance plots provide a clear overview of the model's training dynamics over 20 epochs. In the Loss Analysis graph as in Fig-3, both training and validation losses decrease steadily, indicating progressive learning and convergence. The validation loss stabilizes near the test loss, suggesting that the model maintains its generalization without overfitting. In the AUC Analysis as in Fig-4, the training AUC improves consistently, while the validation AUC shows a slight increase before leveling off, closely matching the test AUC value of 0.835, highlighting strong and consistent performance on unseen data. The Accuracy Analysis as in Fig-5 reflects a similar trend, with validation accuracy hovering around 94.9% and the training accuracy gradually rising to surpass the test accuracy by the final epochs. These plots collectively confirm that the model is well-trained, generalizes effectively, and avoids significant overfitting or underfitting throughout the training process.

A significant architectural choice involved employing EfficientNetB0 starting with no prior knowledge, enabling the model to develop feature recognition tailored to the specific characteristics of grayscale chest X-rays, rather than depending on pre-existing knowledge derived from standard color images. This strategy illustrated that streamlined convolutional designs, when carefully adapted, can still independently acquire distinguishing features within the realm of medical imaging. Furthermore, enlarging the initial 28x28 pixel ChestMNIST images to 224x224 was crucial for integration with the EfficientNet structure and aided in retaining more spatial information throughout the learning process.

Despite the encouraging outcomes, certain constraints warrant consideration. The image collection, while structurally consistent, might contain inaccuracies in its labels due to less rigorous annotation methods common in medical imaging. Additionally, the lack of an independent dataset for verification limits the assessment of the

model's applicability in actual healthcare environments. Future research could explore combining multiple models, incorporating mechanisms to focus on relevant image regions, or employing visualization techniques to understand the model's reasoning. Nonetheless, the present findings provide a firm starting point for the simultaneous categorization of chest conditions using contemporary, computationally economical convolutional neural networks.

5. CONCLUSIONS

This investigation introduces a methodology based on deep neural networks for simultaneously categorizing multiple chest ailments, utilizing the ChestMNIST image collection and the EfficientNetB0 framework. By tailoring the model to process grayscale images and implementing a specific learning procedure, the system effectively learned to forecast the occurrence of as many as 14 distinct chest conditions. The model exhibited robust performance, achieving an area under the ROC curve of 0.839, a prediction accuracy of 94.82%, and a prediction error of 0.1632 on unseen data, showcasing its ability to acquire relevant feature representations without relying on pre-existing knowledge. These outcomes affirm the potential of EfficientNetB0 to tackle intricate medical image analysis tasks even when trained without prior external knowledge.

Examining the performance curves for each of the 14 categories revealed that the model excelled at pinpointing conditions like enlarged heart and fluid accumulation, whereas others, such as small abnormal growths[12] and tissue density changes, presented greater difficulty, possibly due to less distinct visual characteristics or resemblances between categories. The overall performance curve, averaged across all categories, further confirmed the model's consistent reliability. The strategy of adjusting the input image size to 224x224 and implementing appropriate scaling and regularization techniques proved vital for the model's ability to generalize. In conclusion, this study establishes a hopeful and resource-effective procedure for identifying diseases in chest X-rays, providing a solid groundwork for subsequent medical investigation and practical application.

6. FUTURE ENHANCEMENTS

Although the proposed model demonstrates strong performance, there is substantial scope for enhancement in future iterations. Applying strategies like class rebalancing or experimenting with loss functions such as focal loss may improve prediction consistency across all

disease categories, particularly those that are less frequent or more complex to identify. Enhancing model transparency through visualization methods like Grad-CAM or integrated gradients could also increase interpretability, helping healthcare professionals better understand and trust the model's decisions.

Future enhancements might also explore transfer learning by initializing with pretrained models trained on large-scale chest X-ray datasets like CheXpert or NIH ChestX-ray14, which could improve both training speed and model generalization. Incorporating ensemble strategies, attention layers, or even transformer architectures could further strengthen the model's accuracy and resilience. Validating the approach on external datasets from real-world clinical environments and diverse patient populations would provide deeper insights into its practical utility. Ultimately, integrating multiple data sources such as clinical notes or symptom data alongside imaging could lead to more comprehensive and accurate diagnostic tools.

REFERENCES

1. J. Yang, R. Shi, D. Wei, et al., "MedMNIST v2 - A large-scale lightweight benchmark for 2D and 3D biomedical image classification", doi: 10.1038/s41597-022-01721-8.
2. X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri and R. M. Summers, "ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases," 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 2017, pp. 3462-3471, doi: 10.1109/CVPR.2017.369.
3. J. F. V. Oraño, J. F. V. Oraño-Maaghop and E. A. Maravillas, "CXR-based Lung Disease Classification Using Convolutional Neural Network," 2022 IEEE 14th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM), Boracay Island, Philippines, 2022, pp. 1-5.
4. S. Ms. Kavitha, S. Thaarani, A. P. Singh and G. Santhosh, "Chest Disease Classification Using Convolutional Neural Networks," 2023 International Conference on Computer Communication and Informatics (ICCCI), Coimbatore, India, 2023, pp. 1-4, doi: 10.1109/ICCCI56745.2023.10128179.
5. Y. Guan and Y. Huang, "Multi-label chest X-ray image classification via category-wise residual attention learning," Pattern Recognition Letters, vol.

- 116, pp. 58-65, 2018, doi: 10.1016/j.patrec.2018.10.027.
6. I. Allaouzi and M. Ben Ahmed, "A Novel Approach for Multi-Label Chest X-Ray Classification of Common Thorax Diseases," in IEEE Access, vol. 7, pp. 64279-64288, 2019, doi: 10.1109/ACCESS.2019.2916849.
7. A. Djoudi, A. A. Zaid, Y. Guellouma and H. Cherroun, "Genetic Convolutional Neural Networks Approach for Disease Detection From Chest X-Ray Images," 2023 5th International Conference on Pattern Analysis and Intelligent Systems (PAIS), Sétif, Algeria, 2023, pp. 1-8, doi: 10.1109/PAIS60821.2023.10322027.
8. V. G. Sreena, P. L. Deepa and D. N. Ponraj, "Comparative Analysis of Deep Learning Models for Covid-19 Detection from Chest X-rays," 2022 International Conference on Disruptive Technologies for Multi-Disciplinary Research and Applications (CENTCON), Bengaluru, India, 2022, pp. 113-118.
9. I. F. Jassam, S. M. Elkaffas and A. A. El-Zoghabi, "Chest X-Ray Pneumonia Detection by Dense-Net," 2021 31st International Conference on Computer Theory and Applications (ICCTA), Alexandria, Egypt, 2021, pp. 176-179, doi: 10.1109/ICCTA54562.2021.9916637.
10. M. Irtaza, A. Ali, M. Gulzar and A. Wali, "Multi-Label Classification of Lung Diseases Using Deep Learning," in IEEE Access, vol. 12, pp. 124062-124080, 2024, doi: 10.1109/ACCESS.2024.3454537.
11. A. Chatchaiwatkul, P. Phonsuphee, Y. Mangalmurti and N. Wattanapongsakorn, "Lung Disease Detection and Classification with Deep Learning Approach," 2021 36th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC), Jeju, Korea (South), 2021, pp. 1-4, doi: 10.1109/ITC-CSCC52171.2021.9501445.
12. J. Rocha, S. C. Pereira, J. Pedrosa, A. Campilho and A. M. Mendonça, "Attention-driven Spatial Transformer Network for Abnormality Detection in Chest X-Ray Images," 2022 IEEE 35th International Symposium on Computer-Based Medical Systems (CBMS), Shenzhen, China, 2022, pp. 252-257, doi: 10.1109/CBMS55023.2022.00051.
13. W. Fan, X. Guo, L. Teng and Y. Wu, "Research on Abnormal Target Detection Method in Chest Radiograph Based on YOLO v5 Algorithm," 2021 IEEE International Conference on Computer Science, Electronic Information Engineering and Intelligent Control Technology (CEI), Fuzhou, China, 2021, pp. 125-128, doi: 10.1109/CEI52496.2021.9574450.
14. Y. Li and X. He, "COVID-19 Detection in Chest Radiograph Based on YOLO v5," 2021 IEEE International Conference on Computer Science, Electronic Information Engineering and Intelligent Control Technology (CEI), Fuzhou, China, 2021, pp. 344-347, doi: 10.1109/CEI52496.2021.9574463.
15. Z. Li et al., "Thoracic Disease Identification and Localization with Limited Supervision," 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 2018, pp. 8290-8299, doi: 10.1109/CVPR.2018.00865.