# Automated Medical Image Analysis for Disease Diagnosis Using Deep Learning

Kanish Shandilya[1], Devansh Dhingra[2], Komal Kashyap[3], Pooja Singh[4]

*[123] Department of Computer Science Engineering, Mahatma Gandhi Mission's College of Engineering & Technology, Noida, NCR, India*

*[4]Guide, Department of Computer Science and Engineering, MGM's College of Engineering and Technology, Noida, India*

**Abstract- This paper presents an automated medical image analysis system for disease diagnosis using deep learning techniques. The proposed system consists of three independent modules designed to perform pneumonia detection, cardiac disease classification, and atrium segmentation from medical images. Convolutional Neural Networks (CNNs) are employed for classification tasks, while a U-Net architecture is utilized for segmentation.**

**The pneumonia detection model is trained on chest X-ray images and achieves an accuracy of approximately 95–97%. Similarly, the cardiac classification model, trained on processed heart image data, achieves an accuracy in the range of 93–96%. For atrium segmentation, the U-Net model is trained on medical imaging data in NIfTI format, and the performance is evaluated qualitatively through visual comparison of predicted and ground truth masks.**

**The system is implemented using Python and PyTorch, incorporating standard preprocessing, training, and evaluation techniques. The results demonstrate that deep learning models can effectively assist in medical image analysis tasks by providing reliable and consistent outputs across different tasks.**

**Keywords— Deep Learning, Medical Image Analysis, Convolutional Neural Network, U-Net, Disease Diagnosis, Image Segmentation**

## I. INTRODUCTION

Medical image analysis has become an essential component of modern healthcare systems, enabling early detection and diagnosis of various diseases. Traditional methods rely heavily on manual interpretation by medical professionals, which can be time-consuming and subject to human error. With the advancement of deep learning techniques, automated systems have shown significant potential in improving diagnostic accuracy and efficiency [5].

Convolutional Neural Networks (CNNs) have been widely used for image classification tasks, including disease detection from chest X-ray images [2], [3]. These models are capable of learning hierarchical features directly from image data, making them highly effective for medical imaging applications. Similarly, U-Net architectures have proven to be highly suitable for medical image segmentation tasks due to their encoder-decoder structure and ability to capture both spatial and contextual information [1].

The increasing availability of medical imaging datasets and computational resources has further accelerated the adoption of deep learning in healthcare. Automated systems can assist clinicians by providing faster and more consistent analysis of medical images.

This work presents an automated medical image analysis system that integrates multiple diagnostic tasks into a single framework. The system includes pneumonia detection, cardiac disease classification, and atrium segmentation. Unlike many existing works that focus on a single task, this study proposes a modular framework integrating multiple medical image analysis tasks, including classification and segmentation, within a unified system. This approach enhances practical usability and demonstrates how deep learning models can be combined for multi-disease diagnostic support.

## II. LITERATURE REVIEW

Deep learning techniques have significantly improved the performance of automated medical image analysis systems in recent years. CNN-based models have demonstrated strong results in detecting pneumonia from chest X-ray images by learning discriminative features directly from data [3], [5]. These approaches have achieved high accuracy and are widely used in medical imaging research.

In the domain of cardiac disease analysis, CNN architectures have been successfully applied to classify medical images into disease and non-disease categories. These models are capable of handling structured datasets and have shown reliable performance across different studies [5].

For segmentation tasks, the U-Net architecture has become a standard model due to its encoder-decoder structure with skip connections, enabling precise localization of anatomical structures [1]. Variants such as V-Net have also been proposed for volumetric medical image segmentation tasks [6].

Recent studies have also explored combining classification and segmentation approaches to improve diagnostic performance. However, most existing systems focus on individual tasks rather than integrating multiple tasks into a single framework. This work addresses this gap by implementing a modular system capable of performing both classification and segmentation tasks within a unified structure.

## III. METHODOLOGY

The proposed system is a modular deep learning-based framework designed to perform multiple medical image analysis tasks, including classification and segmentation. The system consists of three independent modules: pneumonia detection, cardiac disease classification, and atrium segmentation.

### A. Pneumonia Detection

The pneumonia detection module is based on a Convolutional Neural Network (CNN) designed for binary classification of chest X-ray images. The input images are pre-processed by converting them to grayscale, resizing them to 128 × 128 pixels, and applying normalization.

The model is trained using Binary Cross Entropy loss and optimized using the Adam optimizer with a learning rate of 0.001. Training is performed for 20 epochs with a batch size of 32. A threshold of 0.5 is used to classify the outputs into pneumonia and normal categories. The CNN model consists of multiple layers followed by pooling and fully connected layers for classification.

### B. Cardiac Disease Classification

The cardiac classification module utilizes a CNN-based architecture for binary classification. The dataset consists of pre-processed heart image data stored in NumPy array format. The input images are converted into single-channel format before being fed into the model. The model is trained using Binary Cross Entropy with Logits loss along with class weighting to handle imbalance in the dataset. The Adam optimizer with a learning rate of 0.001 is used for training. The model is trained for 25 epochs with a batch size of 32. The CNN model consists of multiple convolutional layers followed by pooling and fully connected layers for classification.

### C. Atrium Segmentation

The atrium segmentation module is implemented using the U-Net architecture, which is widely used for medical image segmentation. The model follows an encoder-decoder structure with skip connections to preserve spatial information.The dataset consists of medical images in NIfTI (.nii.gz) format along with corresponding segmentation masks. The model is trained using a combination of Binary Cross Entropy and Dice-based loss for 25 epochs with a batch size of 4. The U-Net architecture follows an encoder-decoder structure with skip connections to preserve spatial information and improve segmentation accuracy.

### D. System Workflow

The overall workflow of the system is illustrated in Fig. 1. The system allows users to select a specific diagnostic task, provide the corresponding medical image as input, and obtain predictions from the selected model.

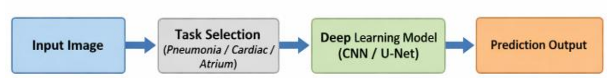


Fig. 1. Overall System Workflow

### E. Dataset Description

The datasets used in this study consist of publicly available medical imaging datasets. The pneumonia detection model is trained on chest X-ray images, while the cardiac classification dataset consists of pre-processed heart image data. The atrium segmentation dataset includes medical images in NIfTI format along with corresponding ground truth masks. The datasets are divided into training and testing sets in an approximate ratio of 80:20.

## IV. IMPLEMENTATION

The system is implemented using Python and the PyTorch framework, along with supporting libraries such as NumPy and Matplotlib. Each module is developed independently and follows a modular design.

The pneumonia detection module processes chest X-ray images using standard preprocessing techniques before feeding them into the CNN model. The cardiac classification module operates on NumPy-based datasets, converting them into tensors for training. The atrium segmentation module processes NIfTI-format medical images along with corresponding masks.

All models are trained using mini-batch gradient descent on a CPU-based environment. Training progress is monitored using loss values across epochs.

## V. RESULTS AND DISCUSSION

The performance of the proposed system is evaluated separately for each module using appropriate metrics and observations.

### A. Pneumonia Detection

The pneumonia detection model achieves an accuracy in the range of approximately 95–97%. The training loss decreases consistently, indicating effective learning. Sample predictions demonstrate that the model is capable of correctly classifying images into pneumonia and normal categories (Fig. 2).

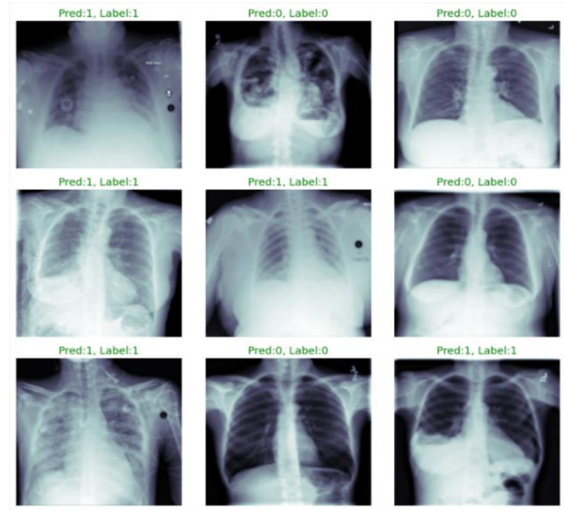


Fig. 2. Pneumonia Detection Output

### B. Cardiac Disease Classification

The cardiac classification model achieves an accuracy of approximately 93–96%. The use of weighted loss functions helps address class imbalance. The model produces reliable predictions for both classes (Fig. 3). These results highlight the effectiveness of deep learning techniques in handling medical image classification tasks with high reliability.In addition, the model benefits from the use of convolutional feature extraction, which allows it to identify subtle patterns that may not be easily visible through manual inspection. The training process also ensures that the model generalizes well across different input samples, reducing the chances of overfitting. The modular design of the system allows this classification component to operate independently, making it easier to integrate with other diagnostic modules. Furthermore, the approach can be extended to larger datasets and more complex architectures in future implementations. These characteristics make the model suitable for practical applications in automated medical diagnosis systems.

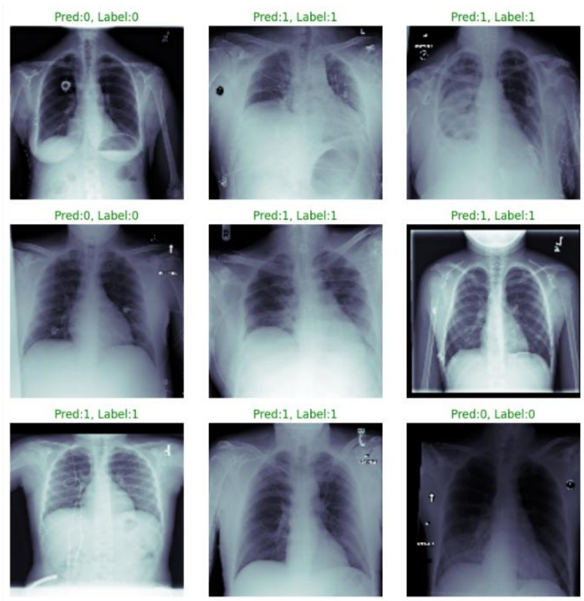

Fig. 3. Cardiac Classification Output

C. Atrium Segmentation

The atrium segmentation model is evaluated qualitatively using visual comparison between predicted masks and ground truth masks. The U-Net architecture successfully captures the structure of the atrium region. The predicted segmentation outputs closely resemble the actual masks (Fig. 4).

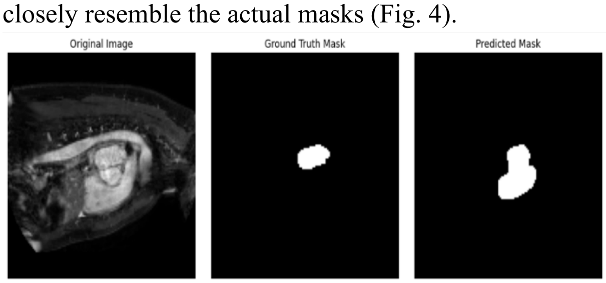

Fig. 4. Atrium Segmentation Output (Predicted vs Ground Truth)

D. Discussion

The results demonstrate that deep learning models can effectively perform both classification and segmentation tasks in medical imaging. The modular structure of the system allows flexibility in handling different types of data. However, performance may be influenced by dataset size and training limitations. Despite promising results, the system has certain limitations. The performance depends on dataset size and quality, and the models are trained on limited data. Future work can focus on larger datasets, real-time deployment, and integration with clinical systems.

E. Evaluation Metrics

In addition to accuracy, the performance of classification models is evaluated using Precision, Recall, and F1-score. These metrics provide better insight into model performance, especially in imbalanced datasets.

For segmentation tasks, Dice Coefficient is used to measure the overlap between predicted masks and ground truth masks.

TABLE I: Model Performance

| Model | Task | Metric |
|---|---|---|
| CNN | Pneumonia Detection | 95–97% Accuracy |
| CNN | Cardiac Classification | 93–96% Accuracy |
| U-Net | Atrium Segmentation | Qualitative |

## VI. CONCLUSION

This paper presents an automated medical image analysis system for disease diagnosis using deep learning techniques. The system integrates classification and segmentation tasks within a unified framework using CNN and U-Net architectures.

The results demonstrate that the models achieve reliable performance across different tasks. This work focuses on practical implementation and highlights the effectiveness of deep learning techniques in medical image analysis. The system can be further improved by incorporating larger datasets and more advanced models. The proposed system demonstrates the practical applicability of deep learning techniques in real-world medical image analysis scenarios. Future improvements may include the use of advanced architectures such as ResNet and EfficientNet, incorporation of larger datasets, and deployment of the system as a real-time clinical decision support tool.

## REFERENCES

[1] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," MICCAI, 2015, pp. 234–241.

[2] Alex Krizhevsky, I. Sutskever, and Geoffrey Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," NeurIPS, 2012.

[3] D. Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," arXiv:1711.05225, 2017.
[4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," CVPR, 2016.
[5] G. Litjens et al., "A Survey on Deep Learning in Medical Image Analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
[6] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation," 3DV, 2016.
[7] J. Deng et al., "ImageNet: A Large-Scale Hierarchical Image Database," CVPR, 2009.
[8] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," arXiv:1409.1556, 2014.
[9] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, MIT Press, 2016.
[10] M. Abadi et al., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems," 2016.