

SELF-SUPERVISED VISION TRANSFORMERS FOR EARLY DISEASE DETECTION IN LOW-RESOURCE MEDICAL IMAGING ENVIRONMENTS FRAMEWORK

Anindya Bose¹, Denesh Sooriamoorthy^{2*}, and Kumaresu Murugasu³

¹ Microelectronics and VLSI Technology, Maulana Abul Kalam Azad University of Technology, West Bengal, India;

² School of Engineering, Asia Pacific University of Technology and Innovation, Kuala Lumpur, Malaysia;

³ School of Global Hospitality & Tourism, Asia Pacific University of Technology & Innovation (APU), Kuala Lumpur,

* **Corresponding author:** denesh.sooriamoorthy@apu.edu.my.

ABSTRACT: This study proposes a Self-Supervised Vision Transformer (ViT) framework for early disease detection in low-resource medical imaging environments. The research addresses major challenges in healthcare systems, including limited labeled medical imaging datasets, inadequate computational resources, and a shortage of medical experts. Traditional Convolutional Neural Network (CNN)-based models generally require large annotated datasets and extensive computational power, limiting their applicability in resource-constrained healthcare settings. To overcome these limitations, the proposed framework integrates self-supervised learning techniques with a Vision Transformer architecture to improve medical image classification performance using limited labeled data. The framework uses chest X-ray images for pneumonia detection and employs preprocessing methods, including normalization, augmentation, resizing, and denoising, to improve image quality and consistency. Self-supervised learning techniques, including contrastive learning and masked image modeling, are used to learn meaningful image representations from unlabeled data. The Vision Transformer architecture captures both local and global contextual information through self-attention mechanisms, enhancing feature extraction and disease representation. Experimental evaluation is conducted using metrics such as accuracy, precision, recall, F1-score, confusion matrix, and ROC-AUC, with a comparative analysis against the ResNet50 CNN model. The results demonstrate that the proposed Vision Transformer framework achieves strong disease-sensitive classification performance and excellent recall for pneumonia detection, while reducing dependency on large annotated datasets. Although the CNN model achieved higher overall classification accuracy, the Vision Transformer demonstrated superior contextual learning and disease-discrimination capabilities. The study highlights the potential of self-supervised Vision Transformers for scalable, efficient, and cost-effective AI-assisted medical image analysis in low-resource healthcare environments.

Keywords: self-supervised learning; vision transformers; medical imaging; early disease detection; deep learning; chest x-ray classification; artificial intelligence in healthcare; pneumonia detection; healthcare informatics.

I. INTRODUCTION

The healthcare sector has undergone profound changes with the advent of Artificial Intelligence (AI), especially in medical imaging and diagnosis. AI-powered medical imaging systems can help healthcare professionals identify diseases at early stages, aid diagnosis, minimize human error, and facilitate rapid treatment decisions (Khalifa & Albadawy, 2024). With the advent of technologies such as deep learning, computer vision is increasingly used to analyze medical images, including X-rays, Magnetic Resonance Imaging (MRI), Computed Tomography (CT), and ultrasound scans. Timely disease diagnosis can affect

patient survival and treatment outcomes in diseases such as cancer, pneumonia, cardiovascular disease, and, in particular, neurological disease (Egbuna et al., 2025).



(Source: Khalifa and Albadawy, 2024).

FIGURE 1. The four AI domains and eight functions in diagnostic imaging.

In the field of medical image classification, traditional deep-learning models, especially Convolutional Neural Networks (CNNs), have proven effective. But these models are generally large and require a labeled dataset and significant computational power, which may be hard to access in low-resource health care settings. In recent years, the power of Vision Transformers (ViTs) has risen to a level that makes them a strong alternative, thanks to their ability to leverage long-range contextual information via self-attention. Moreover, self-supervised learning methods can learn meaningful image representations from unlabeled data, reducing the need for large annotated datasets (Sallam & Shnan, 2025). Hence, the combination of self-supervised learning and Vision Transformers is an exciting prospect for achieving effective and reliable early disease detection in resource-limited healthcare settings.

In low-resource healthcare settings, early disease diagnosis remains a significant challenge due to a lack of high-quality medical image datasets, computational resources, and medical experts. Medical imaging models based on traditional convolutional neural networks (CNNs) have been shown to be well-suited for disease classification (Öksüz et al., 2025). But these models are very dependent on supervised learning and large annotated imaging datasets. These constraints limit their usefulness in medical imaging systems with limited resources. Moreover, CNN models might have difficulty learning relationships in high-resolution medical images, which could impact the early detection of diseases with a high degree of accuracy (Salehi et al., 2023).

The challenges outlined above can be addressed by recent developments in Vision Transformers and self-supervised learning, which enable the learning of meaningful image representations from unlabeled data and enhance contextual feature extraction. However, few studies have specifically addressed the use of self-supervised Vision Transformer models for early disease diagnosis in low-resource health care environments. In this work, we propose to address this gap by developing and evaluating an efficient self-supervised Vision Transformer framework for medical image analysis.

The purpose of this work is to propose an early disease detection system using a self-supervised Vision Transformer (ViT) model trained on small labeled medical datasets.

The following are the research hypotheses:

- H1: Self-supervised Vision Transformers show great promise in early disease detection in low-resource medical imaging settings.
- H2: The proposed self-supervised Vision Transformer framework outperforms the conventional CNN-based medical imaging models.

- H3: Self-supervised learning alleviates the need for huge labeled medical imaging datasets while retaining classification performance.
- H4: The Vision Transformer-based feature extraction significantly contributes to accurate disease classification in medical images by providing a better context of the image data.

The study is significant because it addresses important problems, such as early disease detection in low-resource healthcare settings, where labeled medical imaging data and sophisticated computational resources are limited. The research presents a self-supervised Vision Transformer architecture. It can enhance the performance of AI-driven medical image analysis without relying on extensive annotated datasets (Rane, 2023). The research can help healthcare systems determine more accurately and efficiently whether patients have a particular disease, when resources for medical imaging are scarce. The research has a strong impact on the fields of self-supervised learning and Vision Transformers, as well as on medical image classification in healthcare informatics and AI. Comparing the proposed framework with conventional CNN models also provides insights into the performance of transformer architectures in medical image applications. From a practical point of view, the results could facilitate better diagnosis or interpretation, reduce manual workload, and aid quicker decision-making in clinical practice, especially in developing and resource-limited countries.

The study focuses on developing and testing a self-supervised Vision Transformer model for early detection of diseases from medical imaging data in resource-limited healthcare settings. The research mainly focuses on the use of self-supervised learning techniques to reduce reliance on large labeled medical datasets while maintaining high disease classification performance. In the study, medical images include X-ray, MRI, CT, and ultrasound images, which are frequently used to diagnose diseases. The research ranges from data preprocessing, feature extraction, implementation of Vision Transformers, training, fine-tuning with limited labeled data, to performance evaluation using classification metrics such as accuracy, precision, recall, F1-score, and ROC-AUC. A comparative analysis between the proposed frameworks and traditional CNN-based medical imaging models is also included in the study. However, the research is limited to image-based disease classification tasks and does not include real-time clinical deployment, multimodal integration of healthcare data, or large-scale implementation in hospital systems.


II. LITERATURE REVIEW

1. OVERVIEW OF MEDICAL IMAGING SYSTEMS

Medical imaging systems play a significant role in the modern healthcare sector, helping to diagnose, monitor, and plan treatment for diseases. Medical imaging technologies include X-ray, Magnetic Resonance Imaging (MRI), Computed Tomography (CT), and ultrasound. X-ray imaging is commonly used to expedite the diagnosis of bone fractures, chest infections, and dental abnormalities because it is inexpensive (Cleveland Clinic, 2022). MRI is particularly useful for diagnosing neurological and musculoskeletal conditions and has a high ability to visualize soft tissues (Spinnato et al., 2022). Using CT scanning, imaging of the human body's internal organs is performed using a large number of X-rays, which helps enable more precise diagnosis of tumors, injuries, and cardiovascular diseases (Rayan et al., 2025). A form of imaging that uses sound waves to make a real-time view of the body's organs and tissues, ultrasound is often used in prenatal monitoring and cardiac examinations. The volume of medical imaging data produced by these systems is large and continues to grow, placing a high demand on artificial intelligence to analyze this data for precise and efficient disease diagnosis.

2. EARLY DISEASE DETECTION USING ARTIFICIAL INTELLIGENCE

Artificial Intelligence (AI) has become an important technology for improving early disease detection in healthcare systems. With the help of AI technology such as machine learning, deep learning, and computer vision, clinical and medical images and information can be analyzed automatically with high accuracy and efficiency (Elyan et al., 2022). Early disease diagnosis is key because it enables detection in the initial stages, which can lead to successful management and better survival rates for patients.



(Source: Elyan et al., 2022).

FIGURE 2. Schematic diagram a of a CNN model with an arbitrary architecture.

Medical image analysis is also emerging as a field for detecting diseases such as cancer, pneumonia, tuberculosis, cardiovascular diseases, and neurological diseases using AI-based systems. These systems help radiologists by reducing diagnostic errors and manual work (Kothinti, 2024). Also, AI models can detect subtle patterns in images that are not easily visible to human observers. The application of AI in medical imaging is therefore helping to improve diagnostic accuracy, aid quicker clinical decision-making, and deliver improved healthcare.

3. DEEP LEARNING IN MEDICAL IMAGING

The use of deep learning has greatly advanced the field of medical image analysis, facilitating the automatic extraction of features and the precise classification of medical conditions. Deep learning models learn complex visual patterns from medical images without feature engineering. CNNs are one of the most prevalent deep learning architectures in medical imaging owing to their excellent performance in image recognition and classification. X-ray, MRI, CT, and ultrasound images have been successfully used to detect diseases with CNN-based models.



(Source: Natarajan et al., 2024).

FIGURE 3. Example of a chest X-ray.

Some applications include cancer detection, pneumonia diagnosis, brain tumor classification, and diabetic retinopathy screening (Natarajan et al., 2024). In healthcare imaging systems, deep learning methods can be employed for segmentation, object detection, and anomaly detection. But numerous deep learning models need a significant amount of labeled data and computing power to train properly. All these constraints pose difficulties when implementing a dependable AI-based diagnostic system in low-resource healthcare settings.


4. VISION TRANSFORMERS (ViT) IN HEALTHCARE

Recently, Vision Transformers (ViTs) have been developed as powerful deep learning models for medical image analysis and healthcare applications. Vision Transformers differ from CNNs in that they rely on a self-attention mechanism to process image patches and capture long-range context (Wang et al., 2025). This

function facilitates a better understanding of global features in images and complex disease patterns. In the medical field, healthcare applications of Vision Transformers have been growing for medical imaging tasks such as disease classification, tumor detection, organ segmentation, and abnormality detection in X-ray, MRI, CT, and ultrasound images. ViT-based models have been shown to be effective at detecting subtle visual abnormalities that are difficult to detect with traditional CNN architectures (Singh et al., 2024). In addition, Vision Transformers have the potential to enhance diagnostic accuracy and generalization performance when combined with large-scale pretraining strategies. Therefore, the use of ViTs in developing state-of-the-art AI-based health imaging applications and early detection models is emerging as a promising direction for future research.

5. SELF-SUPERVISED LEARNING TECHNIQUES

Self-Supervised Learning (SSL) is a new machine learning paradigm in which networks learn meaningful feature representations from unlabeled data without requiring extensive manual annotation (Abdulrazzaq et al., 2024). Self-supervised learning is also particularly beneficial in medical imaging, where labeled data is scarce, costly, and time-intensive to acquire. This strategy automatically generates supervisory signals from the data itself, enabling models to acquire useful visual patterns that can be fine-tuned on a small labeled dataset.



(Source: Abdulrazzaq et al., 2024).

FIGURE 4. The major workflow related to SSL.

Popular self-supervised learning methods are contrastive learning and masked image modeling. The contrastive learning method encourages models to learn differences between image representations to enhance feature-extraction performance. A masked image modeling method is used to mask parts of the image and train the model to fill in the missing regions. Self-supervised learning is particularly promising in the context of low-resource healthcare imaging, where valuable medical images are scarce, as these techniques enhance representation learning, minimize reliance on labeled medical images and boost classification performance.

6. LOW-RESOURCE MEDICAL IMAGING CHALLENGES

There are several obstacles to successfully implementing AI-based disease diagnosis systems in low-resource medical imaging settings. The biggest problem is the lack of large-scale, labeled medical imaging datasets for training deep learning models (Piffer et al., 2024). Often, medical image annotation relies on

expert radiologists, making the process costly and time-consuming. In addition, in resource-limited areas, healthcare centers often have inadequate image quality due to outdated equipment, substandard infrastructure, and non-standard imaging practices. The adoption of more sophisticated AI models is also hindered by computational resources, including memory, processing power, and storage. A further difficulty stems from the lack of healthcare professionals and the technical skills required to operate AI-based diagnostic systems. These problems make traditional supervised deep learning models less reliable and scalable.


7. EXISTING DISEASE DETECTION MODELS

Existing disease detection models in medical imaging are primarily based on CNNs, more recently, transformer-based architectures. CNN-based models such as ResNet, VGGNet, DenseNet, et and EfficientNet have been extensively employed for disease classification, tumor detection, image segmentation, and abnormality identification because of their ability to extract image features effectively (Mienye et al., 2025). These models have been shown to be accurate in analyzing X-ray, MRI, CT, and ultrasound images. However, CNNs are primarily geared towards local spatial characteristics and tend to perform well only with large amounts of labeled data. Recently, ViTs have demonstrated the effectiveness of their self-attention mechanism in exploiting contextual understanding. CNNs cannot represent global image relationships as well as transformer-based frameworks can.

8. RESEARCH GAP IDENTIFICATION

Current medical imaging research mainly focuses on supervised CNN-based models, which require large amounts of labeled data and significant computational power. The lack of adequate annotated medical imaging data and advanced infrastructure in low-resource healthcare settings presents a challenge to this. While there has been some success in image analysis with ViT, there has been little research on combining self-supervised learning with ViT-based architectures for early disease detection. Additionally, existing work lacks attention to computational efficiency, contextual feature learning, and real-world applications in resource-limited medical imaging settings with limited labeled datasets.

9. CONCEPTUAL FRAMEWORK DEVELOPMENT



(Source: Self-developed).

FIGURE 5. Conceptual framework.

The combination of self-supervised learning and a vision transformer architecture can assist in accurate disease detection and early diagnosis using the limited chest X-ray dataset available in low-resource medical imaging environments.

III. PROPOSED FRAMEWORK OR METHODOLOGY

1. MEDICAL IMAGE DATASET COLLECTION

A medical image database is generated from chest X-ray images for training, validation, and testing of the proposed framework. The datasets contain both labeled and unlabeled sets of X-ray images from public medical imaging repositories and healthcare research databases. The unlabeled datasets are mostly used for fine-tuning and disease classification. The gathered X-ray images include abnormal conditions such as pneumonia, tuberculosis, and COVID-19, which help in accurate early detection of these diseases and in evaluating the model.

2. DATA PREPROCESSING

Medical image data preprocessing is important in model training to enhance the quality and consistency of medical images. To improve the performance of the proposed self-supervised Vision Transformer framework, the pre-processing techniques of normalization, augmentation, resizing, and denoising are used in this study. First, all medical images are resized to the same resolution. The image is then normalized to make the intensities of the same image from different imaging devices more consistent. Data augmentation is used to diversify the data and reduce overfitting during training, including rotation, flipping, cropping, and brightness adjustment. Also, irrelevant artifacts are removed, and the image is clarified using denoising methods. Such pre-processing steps facilitate the extraction of relevant visual information, optimize learning capabilities, and boost the accuracy of disease diagnosis in resource-constrained medical imaging settings.

3. SELF-SUPERVISED LEARNING FRAMEWORK

The proposed framework aims to learn meaningful image representations from large quantities of unlabeled medical images using self-supervised learning. Self-supervised learning differs from conventional supervised learning approaches that require large amounts of labeled data, as it allows the model to generate its own supervisory signals from image data. The framework is trained efficiently using various techniques, including contrastive learning and masked image modeling. The model is capable of recognizing similarities and differences between image representations and detecting contextual patterns in medical scans that are not apparent. In the pretraining stage, unlabeled data from X-ray, MRI, CT, and ultrasound images are loaded to strengthen feature extraction capabilities. Pretraining is followed by fine-tuning with a small labeled dataset for disease classification tasks. This method reduces reliance on costly annotation and improves the model's generalization. The self-supervised framework effectively improves feature learning, classification accuracy, and early disease detection in low-resource health care settings.

4. VISION TRANSFORMER ARCHITECTURE

This study uses the Vision Transformer architecture to enhance medical image analysis through high-level self-attention mechanisms. In contrast to traditional Convolutional Neural Networks, which consider only local spatial features, the Vision Transformer (ViT) processes both local and global contextual information in medical images. First, each medical image is segmented into patches, and embedding vectors are then generated from these patches for use in the model. Image regions are then encoded with positional encoding to maintain spatial relationships. The transformer encoder is composed of multiple layers of self-attention to model interactions among image patches and extract key disease patterns. It allows the framework to detect complex structures and abnormalities in X-ray, MRI, CT, and ultrasound images. In addition, with the help of self-supervised learning, the Vision Transformer can efficiently learn features even when only a small number of labeled examples are available. Thus, the proposed architecture enhances the performance of the classification model, context understanding, and early disease detection in resource-limited health care settings.

5. FEATURE EXTRACTION AND REPRESENTATION

In medical image analysis, feature extraction and representation play an essential role in discovering patterns associated with the disease. In this paper, we present a method to extract deep features and

semantics from chest X-ray images based on a Vision Transformer (ViT). The Vision Transformer methodology involves breaking down the image into smaller patches and using the self-attention mechanism to learn contextual and specific features, providing global context for the image (Aburass et al., 2025). The chest X-ray images are downsized to 224×224 pixels and then embedded into the transformer encoder layers. The extracted features are characteristic of abnormalities such as lung opacities and pneumonia-related patterns. The deep feature representation improves the framework's ability to capture hidden disease features and supports medical image analysis.

6. FINE-TUNING WITH LIMITED LABELED DATA

The proposed system provides a framework for addressing the problem of labeled medical datasets using transfer learning and fine-tuning techniques. We fine-tuned a pre-trained Vision Transformer model for chest X-ray classification using a small labeled dataset. The trained layers contain the learned visual representations, and the final classification layer is fine-tuned on the target disease classes (Huang et al., 2023). The fine-tuning process is performed using labeled chest X-ray images of pneumonia and normal cases. This method simplifies the training process, reduces computational resource requirements, and increases the model's adaptability for medical imaging tasks. The fine-tuning also reduces reliance on large annotated datasets, thereby enabling the system to be applicable in low-resource healthcare environments (Tang et al., 2021).

7. EARLY DISEASE DETECTION AND CLASSIFICATION

The framework proposed in this paper can automatically detect and classify disease in chest X-ray images. The Vision Transformer will take the extracted image features and classify the input image as either the normal or pneumonia class. The classification process is performed by the final fully connected layer within the transformer architecture (Zhang et al., 2024). Early detection of disease is significant because it can improve treatment and minimize health risks. The framework provides binary classification and prediction output as probability scores. The automated detection process increases diagnostic efficiency and supports the development of AI-based healthcare systems.

8. MODEL OPTIMIZATION

Optimization techniques for models' performance and computational efficiency are applied in this paper. The Adam optimizer is used to optimize the model parameters during training. The cross-entropy loss function is used to measure the error of predictions in classification problems. Other functions of learning rate optimization include ensuring training stability and accelerating convergence (Chen et al., 2021). GPU acceleration on the Google Colab platform decreases training time and computational complexity. Besides, image pre-processing methods like normalization and image resizing ensure robustness and the capability of feature learning for the models. These optimizations improve the overall performance of this framework.

9. EXPLAINABILITY AND VALIDATION

Explainability and validation are crucial to ensuring the reliability of AI-powered medical imaging systems. The proposed framework is tested and evaluated with accuracy, precision, recall, F1-score, and the confusion matrix analysis. The above metrics are used to evaluate disease classification performance (Jiang et al., 2022). The framework is validated by implementing it on other test sets and evaluating its ability to generalize to them. For comparison, the Vision Transformer is also compared with the ResNet50 CNN. In addition, a confusion matrix visualization is used to analyze correctly and incorrectly classified cases, thereby enhancing the transparency and reliability of medical diagnosis systems.

IV. PERFORMANCE EVALUATION

Evaluating the performance of medical image classification systems is critical to determining their effectiveness and clinical applicability. We evaluate this framework using a Self-Supervised Vision Transformer and common classification metrics, including Accuracy, Precision, Recall, F1-score, and ROC-AUC (Huang et al., 2023). These evaluation metrics provide a holistic picture of the proposed model's ability to detect the disease. The experimental evaluation is performed on chest X-ray images with pneumonia and

normal classes. To test the framework's ability to generalize, the set is split into training and testing sets. Moreover, the proposed Vision Transformer model is compared with the ResNet50 CNN to assess the transformer model's performance and effectiveness in low-resource medical imaging settings.

1. ACCURACY

One of the most commonly used evaluation metrics in medical image classification tasks is accuracy. Accuracy measures the framework's performance in terms of the overall correctness of disease classification predictions. Accuracy is defined as the proportion of correctly classified samples to the total number of samples in the testing dataset. The mathematical formula for accuracy is:

Where:

- TP = True Positives
- TN = True Negatives
- FP = False Positives
- FN = False Negatives

The Vision Transformer model in this study achieved a test accuracy of about 70.5%, compared to the ResNet50 CNN model, which achieved about 85.4%. From the results obtained, it was observed that the CNN model achieved better overall performance in classifying the selected chest X-ray dataset. Although the Vision Transformer had lower overall accuracy, its ability to identify pneumonia cases was excellent. Medical diagnosis accuracy may not be representative of the overall performance of the medical diagnosis system, as classification accuracy on imbalanced datasets can be biased. Hence, many other evaluation metrics, such as Precision, Recall, and F1-score, are required for a complete performance evaluation.

2. PRECISION

Precision measures the percentage of positive disease cases that are correctly predicted among all predicted positive cases. It quantifies how reliable the disease predictions made by the framework are. A high-precision model has fewer false-positive predictions. Precision is mathematically defined as:

Precision is important in medical imaging systems because it helps avoid additional medical tests and reduce healthcare costs. The pneumoniaa classification modelVision Transformerr precision score is ~ 0.68. The ResNet50 CNN model has a higher precision score. The results show that the ResNet50 model achieved a better balance in classification performance, while the Vision Transformer had a higher false-positive rate for normal chest X-ray images. However, despite this, the Vision Transformer achieved high accuracy in pneumonia detection and good sensitivity in disease detection.

3. RECALL


Recall evaluates the system's ability to detect disease. Recall evaluates the performance of the model in classifying positive disease samples in the dataset. Recall is especially important in healthcare applications, where missing a case of disease may have severe clinical consequences. Recall formula:

The pneumonia score was about 1.00, indicating that the framework detected nearly all pneumonia cases in the test set. This is further evidence of how fragile transformer architectures are for diseases. The ResNet50 CNN model also performed well, achieving a recall above 0.99 in detecting pneumonia. However, the Vision Transformer slightly outperformed the CNN model in terms of disease sensitivity. High recall is desired in medical diagnosis systems to help prevent dangerous false-negative predictions in disease detection.

4. F1-SCORE

The F1-score gives a balanced view of Precision and Recall. It is especially beneficial when the dataset is imbalanced across classes, as it accounts for false-positive and false-negative predictions. The F1-score is mathematically represented as:

The Vision Transformer model achieved a weighted F1 score of about 0.64, while the ResNet50 CNN model achieved about 0.85. The CNN model's performance is more balanced between the normal and pneumonia classes, as indicated by the higher F1 score.



(Source: Python).

FIGURE 4.1: Confusion matrix of the Vision Transformer (ViT) model.

Because many normal images were misclassified as pneumonia cases, the Vision Transformer achieved a lower F1-score. However, in the context of disease detection, the Vision Transformer showed high performance in detecting disease-positive samples, which is useful in healthcare systems that emphasize early detection.

5. ROC-AUC ANALYSIS

The classification capability of the proposed framework is assessed using ROC-AUC analysis, which compares the True Positive Rate (TPR) with the False Positive Rate (FPR). The Receiver Operating Characteristic (ROC) curve shows the trade-off between sensitivity and specificity of a binary classifier across different decision thresholds. The Area Under the Curve (AUC) is a quantitative measure of classification performance. AUC (Area under the curve) measures a model's ability to discriminate between disease and normal states.





FIGURE 3. ROC curve of the Vision Transformer (ViT) model.

(Source: Python)




(Source: Python).

FIGURE 4. ROC curve of the ResNet50 CNN model.

The proposed Vision Transformer framework achieved an AUC of ~0.961, and the CNN model ResNet50 achieved an AUC of ~0.950. The overall classification accuracy was better for the ResNet50 model, while the Vision Transformer performed a bit better in the disease discrimination task as measured by ROC-AUC. The ROC curve results confirmed that both models demonstrated effective disease classification performance in chest X-ray image analysis (Aburass et al., 2025).

6. COMPARATIVE ANALYSIS WITH EXISTING MODELS

The effectiveness of the proposed Vision Transformer framework is compared with the traditional CNN-based model by performing a comparative analysis. To compare the proposed framework with the ResNet50 CNN model, the evaluation metrics such as Accuracy, Precision, Recall, F1 Score, and ROC AUC analysis are used. The results of the experimental evaluation showed that the ResNet50 CNN model achieved the highest overall classification accuracy and was more balanced in prediction performance on the chest X-ray dataset.



(Source: Python).

FIGURE 5. Confusion matrix of the ResNet50 CNN model.

The sensitivity of the Vision Transformer, however, was better than the other models, with almost full recall and a slightly better ROC-AUC score, for the detection of pneumonia. The analysis shows that Vision Transformers excel in both disease representation and disease-sensitive classification tasks. CNN models, however, might perform better on smaller medical datasets, as they are generally more data-efficient than transformer-based models. In general, both frameworks demonstrated a high degree of automated disease classification in a low-resource medical image analysis context.

V. RESULTS AND DISCUSSION

The experimental results and discussion of the proposed Self-Supervised Vision Transformer framework for early disease diagnosis in low-resource medical imaging environments are presented in this chapter. The framework was applied to chest X-ray images with pneumonia and normal classes. The performance metrics, such as Accuracy, Precision, Recall, F1-score, Confusion Matrix analysis, and ROC-AUC evaluation, were used in the experimental evaluation. Furthermore, a comparative analysis between the proposed Vision Transformer (ViT) and the ResNet50 CNN model was conducted to evaluate their classification performance and disease discrimination capabilities.

1. EXPERIMENTAL RESULTS

The suggested framework was implemented in the experimental setting with Google Colab with GPU acceleration. During preprocessing of the chest X-ray dataset, image resizing, normalization, and tensor conversion were performed before training the model. Both the Vision Transformer model and the ResNet50 CNN model were fine-tuned with small amounts of labeled chest X-ray images, with the Vision Transformer

model using pretrained weights. We fine-tuned a Vision Transformer model with pretrained weights on a small set of labeled chest X-ray images and used a ResNet50 CNN as a baseline for comparison. The ResNet50 CNN model has a test accuracy of around 85.4%, whereas the Vision Transformer model achieves a test accuracy of around 70.5%. The ROC-AUC results show that the Vision Transformer achieved an AUC of ~ 0.961 , while ResNet50 achieved an AUC of ~ 0.950 . This means that the transformer-based architecture was able to discriminate among the diseases to a greater extent, but with lower classification accuracy. Moreover, the confusion matrix analysis showed that the Vision Transformer model classified almost all patients with pneumonia correctly, with no false negatives. However, the ResNet50 model was more efficient, with higher classification performance for both classes (normal and pneumonia) and fewer false predictions.

2. DISEASE DETECTION PERFORMANCE

The proposed model was assessed for its ability to detect diseases using Precision, Recall, F1-score, and ROC-AUC. The Vision Transformer model performed very well in detecting pneumonia in patients' chest X-rays, as it classified almost all images in the test set as either Positive (pneumonia) or Negative (Healthy). The model's performance in disease-sensitive classification was almost 100%. Pneumonia recall score. A high recall level is important because false-negative predictions may delay diagnosis and pose health risks. The model correctly recognized disease-related abnormalities in most chest X-rays, but it had more false-positive predictions on normal chest X-rays. The CNN model ResNet50 achieved the highest overall classification accuracy and the highest accuracy for classifying images as pneumonia or normal. The CNN model outperformed the Vision Transformer in terms of weighted F1-score and also the number of false positives. However, this self-attention mechanism within the Vision Transformer helped with effective contextual learning and representation of diseases from chest X-ray images.

3. IMPACT OF SELF-SUPERVISED LEARNING

The proposed framework was very successful in fields such as medical imaging, where labeled data is limited, and the use of self-supervised learning contributed significantly to its success. The model leverages learned visual knowledge and integrates it into the framework to improve feature representation in medical images using pretrained Vision Transformer representations. By fine-tuning the pretrained transformer model on a small, labeled chest X-ray dataset, reliance on large annotated medical datasets was reduced, and computational efficiency was improved. Additionally, self-supervised learning improved contextual representation learning, enabling it to capture global semantic relationships and image patterns related to diseases. Self-supervised transformer learning has recently proved that high-quality medical image representation can be obtained with only limited annotations. The results of the experiments performed for this research corroborate this development; the system was able to detect pneumonia-related abnormalities while using only a small labeled training set. The results show that self-supervised Vision Transformers have the potential to be efficient and scalable for disease detection in low-resource healthcare settings.

4. DISCUSSION OF FINDINGS

The results of this study show that both Vision Transformer and CNN-based architectures can be used for the automated diagnosis of diseases in chest X-ray images. However, differences in classification behavior and disease representation capability were observed between the two approaches. The Vision Transformer's contextual feature learning and self-attention mechanisms enabled it to effectively detect pneumonia. The model has high recall and better ROC-AUC performance, suggesting strong disease-discrimination ability. The results are consistent with recent progress in transformer-based medical imaging systems, which have demonstrated improved global feature representation and semantic learning performance through self-attention architectures. On the contrary, the ResNet50 CNN model had higher classification accuracy and better balance between the normal and pneumonia classes. CNN architectures remain efficient and effective when datasets are limited in size (in terms of data quantity), thanks to the efficient, local spatial feature-capturing process of convolutional operations. The findings also reveal a significant challenge for Vision Transformers in scenarios with limited data. Generally, if the datasets are large enough, transformer models generalize well, while CNN models perform well for classification with relatively small amounts of data.

However, the Vision Transformer exhibited excellent disease-sensitive classification ability and is very useful for applications related to early disease diagnosis in healthcare systems.

5. CLINICAL IMPLICATIONS

The proposed framework has significant practical applications in health care systems, including low-resource medical imaging environments. Automated disease detection with Vision Transformers has the potential to assist radiologists in making diagnostic predictions and alleviating the workload of manual interpretation. The high recall and ROC-AUC of the Vision Transformer demonstrate its strong ability to detect pneumonia-positive samples. This is clinically significant, as early disease diagnosis can improve treatment outcomes and reduce time to diagnosis. The framework can also apply to the large-scale screening of chest X-ray images in hospitals and remote health care centers. Additionally, the use of self-supervised learning and pre-trained transformer models enables greater adaptability within healthcare settings, which do not always have extensive annotated medical datasets. Recent developments in lightweight transformer architectures and the deployment of AI in the cloud further bolster the promise of real-time clinical use of AI-driven disease detection systems in resource-limited healthcare settings.

VI. CONCLUSION AND FUTURE WORK

The following chapter highlights the primary conclusions, contributions, and areas where this proposed Self-Supervised Vision Transformer approach has limitations and requires further study in the field of early disease detection in low-resource medical imaging settings. The aim of the study was to improve automated disease classification using chest X-ray images and to evaluate the efficiency of transformer-based architectures compared with CNN models.

1. SUMMARY OF FINDINGS

The study developed and successfully applied an automated framework for detecting and classifying diseases in chest X-ray images using a Self-Supervised Vision Transformer. The framework was based on transfer learning, feature extraction, and fine-tuning techniques for classifying normal and pneumonia images in a low-resource healthcare environment. Experimental testing showed that the Vision Transformer model provided robust, disease-sensitive classification capabilities, especially for detecting pneumonia. The framework was good at identifying disease-positive cases and at representing contextual features, with good recall and ROC-AUC scores. The overall classification accuracy and balanced prediction performance on the smaller datasets were higher for the CNN architectures than for the ResNet50 CNN model. The results confirm that transformer-based architectures are highly effective for contextual feature learning and disease discrimination tasks. The study also demonstrated that it is possible to reduce reliance on large, labeled medical datasets while preserving disease-detection performance by using self-supervised learning and pretrained transformer models.

2. CONTRIBUTIONS OF THE STUDY

This study offers both theoretical and practical contributions to AI-assisted medical imaging research. From a theoretical perspective, our work contributes to the ongoing development and application of self-supervised learning and Vision Transformers to the healthcare image analysis. The results demonstrate the potential of context-aware representation learning using transformers in disease diagnosis and classification. The proposed framework offers a medical imaging solution that is practical for medical institutions with limited computational infrastructure and annotated datasets. The deployment of the framework on Google Colab and publicly available chest X-ray datasets illustrates the potential of creating cost-efficient AI-aided diagnostic systems. Moreover, the comparison with CNN-based models provides insights into the advantages and disadvantages of Vision Transformer and CNN architectures in medical image classification problems.

3. LIMITATIONS

Despite the limitations identified in this study, the results still showed potential. The framework was first tested on a relatively small chest X-ray dataset with only two classes: normal and pneumonia. This limitation

could limit the generalization of models in more complex medical imaging settings. Second, training and optimizing transformer-based architectures are very resource-intensive. GPU acceleration was used in Google Colab, but the size of the datasets and complexity of the transformer models were limited by hardware and memory constraints. Furthermore, the study primarily focused on binary disease classification and did not evaluate multimodal imaging data or multiple disease categories. The lack of advanced explainability tools, such as Grad-CAM and attention heatmaps, also reduces the model's interpretability for clinical decision-making applications.

4. FUTURE RESEARCH DIRECTIONS

Future work might improve the proposed framework by incorporating multimodal medical imaging information, for example, from computed tomography (CT), magnetic resonance imaging (MRI), and ultrasound, to achieve a more holistic understanding of the disease. The multimodal learning training method can improve feature representation and disease classification performance. Moreover, there is potential for further research into lightweight transformer models, hybrid CNN-transformer models, and improvements in self-supervised learning to make more efficient use of computational resources and improve performance in resource-limited healthcare environments. Also, explainable AI approaches such as Grad-CAM and attention visualization techniques can be used to increase transparency of clinical diagnosis systems. Real-time disease detection systems in hospitals and health centers for medical diagnosis and large-scale disease screening can be a valuable addition to future AI applications.

Author Contributions

The author conducted the conceptualization, methodology, data analysis, investigation, writing, review, editing, and final approval of the manuscript.

Funding

This research received no external funding.

Data Availability

The dataset will be available from the author upon reasonable request.

Conflicts of Interest

The author declares no conflict of interest.

REFERENCES

1. Abdulrazzaq, M. M., Ramaha, N. T., Hameed, A. A., Salman, M., Yon, D. K., Fitriyani, N. L., ... & Lee, S. W. (2024). Consequential advancements of self-supervised learning (SSL) in deep learning contexts. *Mathematics*, 12(5), 758. <https://doi.org/10.3390/math12050758>
2. Aburass, S., Dorgham, O., Al Shaqsi, J., Abu Rumman, M., & Al-Kadi, O. (2025). Vision Transformers in Medical Imaging: A Comprehensive Review of Advancements and Applications Across Multiple Diseases. *Journal of Imaging Informatics in Medicine*, 38(6), 3928–3971. <https://doi.org/10.1007/s10278-025-01481-y>
3. Chen, X., Xie, S., & He, K. (2021). An Empirical Study of Training Self-Supervised Vision Transformers (Version 4). arXiv. <https://doi.org/10.48550/ARXIV.2104.02057>
4. Cleveland Clinic (2022). *X-Ray: What It Is, Types, Preparation and Risks*. [online] Cleveland Clinic. Available at: <https://my.clevelandclinic.org/health/diagnostics/21818-x-ray>.
5. Egbuna, I., Okei, N. C., Ajoku, E., Edobor, O., Issah, P., Olatokun, T., & Akinode, A. O. (2025). Advancing early disease detection with AI: innovations in medical imaging, EHR analytics, and wearable technologies. *International Journal of Life Science Research Archive*, 9, 120-140. <https://doi.org/10.53771/ijlsra.2025.9.1.0049>
6. Elyan, E., Vuttipittayamongkol, P., Johnston, P., Martin, K., McPherson, K., Moreno-García, C. F., ... & Sarker, M. M. K. (2022). Computer vision and machine learning for medical image analysis: recent advances, challenges, and way forward. *Artificial Intelligence Surgery*, 2(1), 24-45. <http://dx.doi.org/10.20517/ais.2021.15>
7. Huang, S.-C., Pareek, A., Jensen, M., Lungren, M. P., Yeung, S., & Chaudhari, A. S. (2023). Self-supervised learning for medical image classification: a systematic review and implementation guidelines. *Npj Digital Medicine*, 6(1). <https://doi.org/10.1038/s41746-023-00811-0>

8. Jiang, J., Tyagi, N., Tringale, K., Crane, C., & Veeraraghavan, H. (2022). Self-supervised 3D Anatomy Segmentation Using Self-distilled Masked Image Transformer (SMIT). In *Lecture Notes in Computer Science* (pp. 556–566). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-16440-8_53
9. Khalifa, M. and Albadawy, M. (2024). AI in diagnostic imaging: Revolutionizing accuracy and efficiency. *Computer Methods and Programs in Biomedicine Update*, 5(100146), pp.100146–100146. doi: <https://doi.org/10.1016/j.cmpbup.2024.100146>.
10. Kothinti, R. R. (2024). The role of deep learning in radiology and medical imaging: improving diagnostic accuracy. *International Journal of Academic Research and Development*, 9, 128-143. https://www.academia.edu/download/121565934/IJNRD2409420_Paper_2024.pdf
11. Singh, S., Kumar, M., Kumar, A., Verma, B. K., Abhishek, K., & Selvarajan, S. (2024). Efficient pneumonia detection using Vision Transformers on chest X-rays. *Scientific Reports*, 14, 2487.
12. Mienye, I. D., Swart, T. G., Obaido, G., Jordan, M., & Ilono, P. (2025). Deep convolutional neural networks in medical image analysis: A review. *Information*, 16(3), 195. <https://doi.org/10.3390/info16030195>
13. Natarajan, K., Muthusamy, S., Sha, M. S., Sadasivuni, K. K., Sekaran, S., Charles Gnanakkan, C. A. R., & A. Elngar, A. (2024). A novel method for the detection and classification of multiple diseases using transfer learning-based deep learning techniques with improved performance. *Neural Computing and Applications*, 36(30), 18979-18997. <https://doi.org/10.1007/s00521-024-09900-x>
14. Öksüz, C., Urhan, O., & Güllü, M. K. (2024). An integrated convolutional neural network with attention guidance for improved performance of medical image classification. *Neural Computing and Applications*, 36(4), 2067-2099. <https://doi.org/10.1007/s00521-023-09164-x>
15. Piffer, S., Ubaldi, L., Tangaro, S., Retico, A., & Talamonti, C. (2024). Tackling the small data problem in medical image classification with artificial intelligence: a systematic review. *Progress in Biomedical Engineering*, 6(3), 032001. <https://doi.org/10.1088/2516-1091/ad525b>
16. Rane, N. (2023). Transformers for medical image analysis: Applications, challenges, and future scope. *Challenges and Future Scope* (November 2, 2023). <https://dx.doi.org/10.2139/ssrn.4622241>
17. Rayan, A. M., Adam, A., Al-Arabi, G., & Ahmed, M. R. (2025). The applications of X-ray technology in medical imaging: advances, challenges, and future perspectives (A review). *Journal of Sustainable Food, Water, Energy and Environment*, 1(2), 39- 61. https://journals.ekb.eg/article_454848_13d4a86fdddbc57bd9db3b245a746e11.pdf
18. Salehi, A.W., Khan, S., Gupta, G., Alabdullah, B.I., Almjally, A., Alsolai, H., Siddiqui, T. and Mellit, A., 2023. A study of CNN and transfer learning in medical imaging: Advantages, challenges, future scope. *Sustainability*, 15(7), p.5930. <https://doi.org/10.3390/su15075930>
19. Sallam, M., & Shnan, M. A. (2025). Enhancing semantic image retrieval using self-supervised learning: A label-efficient approach. *Babylonian Journal of Machine Learning*, 2025, 42- 60. <https://doi.org/10.58496/BJML/2025/004>
20. Spinnato, P., Patel, D. B., Di Carlo, M., Bartoloni, A., Cevolani, L., Matcuk, G. R., & Crombé, A. (2022). Imaging of musculoskeletal soft-tissue infections in clinical practice: a comprehensive updated review. *Microorganisms*, 10(12), 2329. <https://doi.org/10.3390/microorganisms10122329>
21. Tang, Y., Yang, D., Li, W., Roth, H., Landman, B., Xu, D., Nath, V., & Hatamizadeh, A. (2021). Self-Supervised Pre-Training of Swin Transformers for 3D Medical Image Analysis (Version 2). arXiv. <https://doi.org/10.48550/ARXIV.2111.14791>
22. Wang, Y., Deng, Y., Zheng, Y., Chattopadhyay, P., & Wang, L. (2025). Vision transformers for image classification: A comparative survey. *Technologies*, 13(1), 32. <https://doi.org/10.3390/technologies13010032>
23. Zhang, J., Li, F., Zhang, X., Wang, H., & Hei, X. (2024). Automatic Medical Image Segmentation with Vision Transformer. *Applied Sciences*, 14(7), 2741. <https://doi.org/10.3390/app14072741>