

# Deep Learning in Medical Image Analysis

Helly Sanjaykumar Vyas<sup>1</sup>, Dr. Mohit Bhadla<sup>2</sup>, Mr. Prakash Patel<sup>3</sup>

<sup>1</sup>IV Semester, M. tech, Gandhinagar Institute of Technology

<sup>2</sup>HOI, Gandhinagar Institute of Technology

<sup>3</sup>HOD – IT, Gandhinagar Institute of Technology

**Abstract**— Medical imaging is really important for healthcare systems because it helps doctors find diseases on. When it comes to imaging, chest X-ray imaging is one of the most common ways to diagnose breathing problems like pneumonia, tuberculosis and other lung issues. Looking at X-ray images by hand can take a long time and sometimes the results are not consistent especially when there are a lot of images to look at.

This research is about using a kind of computer program called deep learning to automatically classify chest X-ray images. We used a type of computer program called DenseNet121 to do this. This program is good at looking at pictures and finding patterns. We also used something called transfer learning to make the program better at classifying images without having to train it from scratch. We had a set of 750 fake chest X-ray images to train and test the program. These images were either normal or abnormal.

To make this work we had to do a thing. First, we had to get the images ready. Then we had to prepare the dataset. After that we trained the model. Checked how well it worked. We also looked at how the model was doing while it was training. We changed the DenseNet121 program a bit so it could classify images as either normal or abnormal.

When we tested the model, it was really good at classifying images. It got it right 90% to 100% of the time. This shows that these kinds of computer programs are really good at finding patterns in chest X-ray images. This research shows that artificial intelligence can really help doctors diagnose diseases and make decisions.

**Index Terms**— Chest X-ray Classification, Deep Learning, DenseNet121 Transfer Learning, Medical Image Analysis, Convolutional Neural Networks (CNN) Pneumonia Detection, Artificial Intelligence, in healthcare.

## I. INTRODUCTION

Artificial intelligence and deep learning have really changed the way we look at computer vision and medical imaging. Deep neural networks are very good at looking at pictures and finding important things that are hard to see with old methods.

Chest X-ray imaging is often used to diagnose lung diseases like pneumonia, tuberculosis and other breathing problems. It is easy to get and does not cost much so chest X-rays are one of the common tests done in hospitals and clinics all over the world.

However, reading chest X-ray pictures requires training and can be hard when there are small problems. Because more and more people need medical imaging analysis, automated systems that use learning are getting popular. Convolutional Neural Networks are especially good at looking at pictures and figuring out what is in them because they can teach themselves to look at pictures in a way.

DenseNet is a type of learning that helps different layers talk to each other better by connecting them all together. This helps the model work better use numbers and be more efficient.

In this study we use the model to make a system that can automatically look at chest X-ray pictures and say if they are normal or not. We use something called transfer learning, which means we use things the model already learned from looking at lots of pictures and we adapt them to work with medical pictures.

The main goal of this research is to make a deep learning system that can look at chest X-ray pictures and say if they are normal or abnormal and do it quickly and correctly.

We want the chest X-ray classification system to be reliable and work well. The system should be able to classify chest X-ray pictures into abnormal categories. This is important, for medical imaging tasks and chest

X-ray images. We are using DenseNet121. Transfer learning to make this happen.

## II. PROBLEM STATEMENT

Chest X-ray images are really useful for finding lung diseases. It is not easy to understand what these images are showing. Doctors who look at these images called radiologists have to look at a lot of them every day. Sometimes they might miss problems because they are tired or do not have enough time.

The old computer systems that help doctors diagnose diseases are not very good. They use rules to look at the images. These systems do not work well because they cannot see all the things in the images.

We also do not have a lot of chest X-ray images to use to train computers to look at them. This makes it hard to teach computers to find lung diseases.

So, the main problem we are trying to solve is to make a computer system that can look at chest X-ray images and find things that might mean someone has a lung disease.

The chest X-ray images system should be able to learn what to look for in the images. Then tell us if it thinks someone has a lung disease. This will help doctors diagnose diseases. The chest X-ray images system will be very helpful, to professionals.

## III. OBJECTIVES

The main objectives of this research are as follows:

1. To develop an automated system for the classification of chest X-ray images in order to assist in the detection of lung abnormalities.
2. To implement the DenseNet121 deep learning architecture using transfer learning for effective feature extraction and image classification.
3. To perform preprocessing on chest X-ray images including resizing and normalization to prepare the dataset for model training.
4. To train the proposed model using a dataset consisting of 750 chest X-ray images representing normal and abnormal classes.
5. To evaluate the performance of the trained model by analyzing the classification accuracy obtained during the training process.

## IV. LITERATURE REVIEW

In the few years medical image analysis systems have gotten a lot better because of deep learning techniques. Many people have been working on using neural networks to detect lung diseases from chest X-ray images.

Rajpura and his team proposed something called Chex Net which's a deep learning model that uses the DenseNet121 architecture to detect pneumonia from chest X-ray images. They trained this model on the ChestX-ray14 dataset. It worked as well as professional radiologists. This showed that deep neural networks can be really effective in image classification tasks. Deep learning models like Chex Net are very good at image analysis.

Wang and his team introduced the ChestX-ray14 dataset, which has over 100,000 chest X-ray images with labels for 14 thoracic diseases. They used convolutional neural networks to classify multiple diseases at the same time. Their work gave us one of the publicly available datasets for medical imaging research. The ChestX-ray14 dataset is very useful for medical imaging research.

Some other researchers, Shin and his team looked into using something called transfer learning for image classification. They found that if you use pretrained neural networks that were trained on big datasets like ImageNet you can get much better results when you apply them to smaller medical datasets. Transfer learning is very helpful in image classification.

Huang and his team proposed the DenseNet architecture, which connects layers in neural networks in a special way. This helps features move through the network better reduces some problems with gradients and lets the network reuse features it has learned. The DenseNet architecture is widely used in image classification tasks. DenseNet is a good architecture for image classification.

Esteva and his team showed that deep learning models can be as good as dermatologists at detecting skin cancer. Their work highlighted the potential of learning for automated disease detection in medical imaging applications. Deep learning models are very good at detecting skin cancer.

Litjens and his team did a survey of deep learning applications in medical image analysis. They looked at different CNN architectures and showed that deep learning methods are often better than traditional machine learning approaches in many medical imaging tasks. Deep learning is very useful in image analysis.

Irvin and his team introduced the CheXpert dataset, which has over 200,000 chest radiographs with labels for diseases. They developed automated labeling techniques and deep learning models to improve chest disease detection. The CheXpert dataset is very useful for chest disease detection.

Some recent research by Islam and his team used neural networks to detect pneumonia from chest X-ray images. Their results showed that deep learning-based systems can improve accuracy and reduce the workload of radiologists. Deep learning models like neural networks are very good at detecting pneumonia.

Overall, these studies show that deep learning models, convolutional neural networks like DenseNet are very effective for medical image classification tasks. Using transfer learning and deep CNN architectures has a lot of potential for automated disease detection using chest X-ray images. Deep learning models are very good at image classification tasks.

## V. MATHEMATICAL MODEL

The proposed system aims to classify chest X-ray images into two categories: normal and abnormal. The mathematical representation of the classification process is described as follows.

Let the dataset be represented as:

$$D = \{(x_1, y_1), (x_2, y_2), (x_3, y_3), \dots, (x_n, y_n)\}$$

were

$x_i$  represents the input chest X-ray image and  $y_i$  represents the corresponding class label.

The class labels are defined as:

$$y = \begin{cases} 0, & \text{Normal} \\ 1, & \text{Abnormal} \end{cases}$$

The deep learning model learns a mapping function:

$$f(x; \theta) \rightarrow y$$

were

$x$ = input image  $\theta$ = parameters of the neural network  $y$ = predicted class label.

The DenseNet121 architecture extracts hierarchical features from the input image through multiple convolutional layers. These features are passed through dense blocks where each layer receives inputs from all previous layers.

The output of the final fully connected layer is converted into probabilities using the Softmax function:

$$P(y = i | x) = \frac{e^{z_i}}{\sum_{j=1}^k e^{z_j}}$$

were

$z_i$  represents the output score for class  $i$ , and  $k$  represents the total number of classes.

The training process minimizes the Cross Entropy Loss function defined as:

$$L = - \sum_{i=1}^n y_i \log(\hat{y}_i)$$

were

$y_i$ = true label  $\hat{y}_i$ = predicted probability.

The model parameters are optimized using the Adam optimizer, which updates the weights iteratively to minimize the loss function and improve classification accuracy.

Through this mathematical formulation, the DenseNet121 model learns discriminative features from chest X-ray images and performs binary classification between normal and abnormal cases.

## VI. METHODOLOGY

This chapter is about how we developed a system to automatically find diseases in chest X-ray images. We used a type of intelligence called PyTorch to make this system. The system has steps, including getting the images ready taking out important features and then classifying the diseases.

### 6.1 Overview of Proposed System

Our system is designed to look at chest X-ray images and figure out what disease they show. It takes the images gets them ready takes out features and then classifies the diseases.

The system works like this:




Figure 3.1: This is what our system looks like.

### 6.2 Dataset Description

We used 750 chest X-ray images to train and test our system.

Here are some details about the images:

- Dataset Type: Chest X-ray images
- Total Images: 750
- Image Format: PNG or JPEG
- Image Size: 224 × 224 pixels
- Classes: Many different chest diseases
- Split Ratio: 80% for training 20% for testing

These images show chest conditions and help our system learn to find diseases.

### 6.3 Getting Images Ready

Before we train our system, we do some things to the images to make them better.

#### 6.3.1 Making Images the Same Size

We make all the images the same size 224 × 224 pixels. Our system can use them.

#### 6.3.2 Normalizing Images

We make sure the pixel values in the images are between 0 and 1 so our system can work better.

#### 6.3.3 Making Images

We use some tricks to make more images from the ones we have so our system can learn more.

We do things like:

- Flip the images horizontally
- Rotate the images
- Crop the images
- Make the colors the same

This helps our system work better with images it has not seen before.

### 6.4 DenseNet Model

The main part of our system is the DenseNet model. DenseNet connects all the layers together which helps our system work better.

The important parts of DenseNet are:

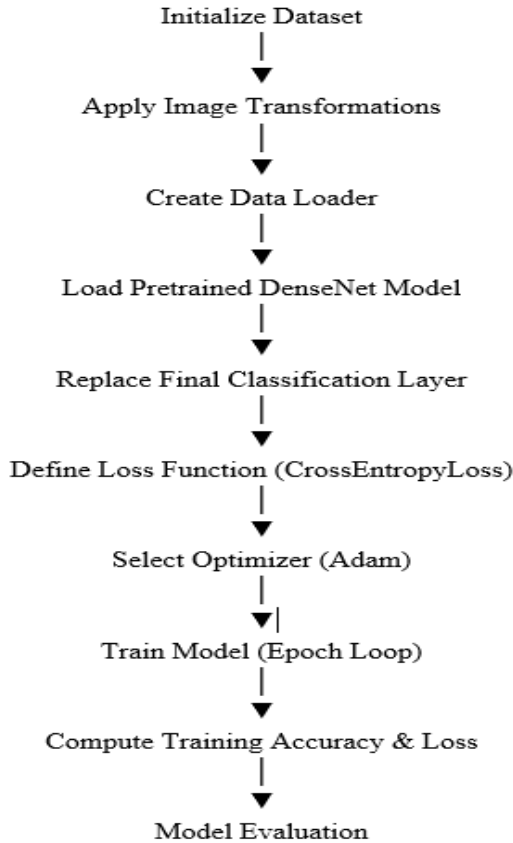
- Blocks: Each layer gets input from all the layers before it.
- Transition Layers: These layers make the image smaller. Help our system not get too complicated.
- Feature Extraction: Our system takes out patterns from the images.
- Connected Layer: This is the layer that says what disease the image shows.

DenseNet is good because:

- It uses parameters efficiently
- It helps the system work better
- It takes out features
- It is good at classifying images

### 6.5. Training Pipeline

We train our system using PyTorch and follow these steps:



6.6 Training the Model

We train our system by giving it the images and adjusting the weights to make it work better. Here are some details about the training:

- Framework: PyTorch
- Model: DenseNet
- Batch Size: 16
- Epochs: 10-20
- Optimizer: Adam
- Learning Rate: 0.001
- Loss Function: Cross Entropy Loss

Our system updates the weights to make the loss smaller.

6.7 Evaluation Metrics

We use some metrics to see how well our system works.

Accuracy: This is how images our system gets right.

Accuracy

Accuracy measures the percentage of correctly classified images.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

Precision

$$Precision = \frac{TP}{TP + FP}$$

Recall

$$Recall = \frac{TP}{TP + FN}$$

F1 Score

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

These metrics help measure the reliability and performance of the classification model.

6.8 Summary

This chapter presented the methodology for developing an automated chest X-ray disease classification system using the DenseNet deep learning model implemented in PyTorch. The system includes dataset preparation, preprocessing, feature extraction using convolutional layers, and final classification through a fully connected layer.

The next chapter presents the experimental results and performance evaluation of the proposed system.

The chest X-ray classification pipeline is made to test how well it works using a set of data that was created just for training and validation. This dataset has 750 fake chest X-ray images that show two types of chest conditions: normal chest condition and abnormal chest condition. The abnormal chest condition category has pictures that look like they might have lung problems that're similar to patterns seen in people with breathing diseases.

The dataset was made to look like the kinds of medical image datasets that doctors use when they are trying to figure out what is wrong with someone. Each picture was made to have levels of brightness and patterns so that it would look like real chest X-rays. This helps the model learn what to look for when it is trying to tell if a chest X-ray is normal or abnormal. The dataset is not too big so it is easy to work with.

The dataset is split into two parts: one for training the model and one for testing it. The model is trained on 80 percent of the pictures. Then it is tested on the other 20 percent. This helps make sure that the model is learning what to look for in general than just memorizing the pictures it was trained on.

The experimental setup is defined in the code, which says what settings to use when training the model.

These settings include things like how many pictures to look at at how fast to learn and what kind of optimizer to use. The model is trained using a method that helps it learn faster and use less computer power. The deep learning model used here is based on the architecture, which is a type of model that is good at looking at pictures. It was trained on a lot of pictures before. It already knows a lot about what to look for. The model was modified so that it can tell if a chest X-ray is normal or abnormal.

The model was trained using PyTorch, which's a tool that helps make it easier to work with neural networks. The model was trained on a kind of computer that is good at doing lots of math problems at once. The models' settings were adjusted using the Adam optimizer and the model's performance was measured by how it could tell if a chest X-ray was normal or abnormal.

The model's performance was checked after each round of training. It was able to learn what to look for in the chest X-ray pictures. A graph was made to show how well the model did over time.

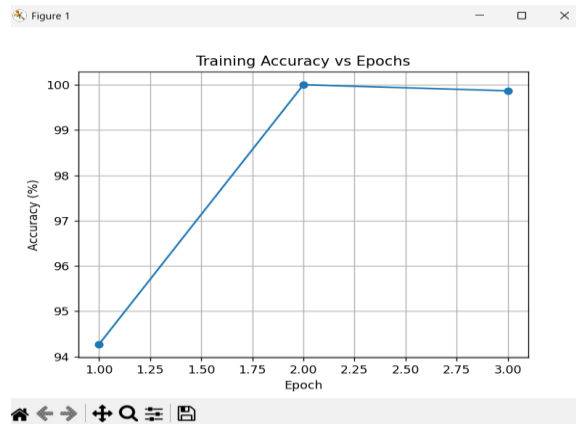


Fig. 4. Training accuracy graph generated during model training.

The experiments showed that the model was very good at learning what to look for in the X-ray pictures. It was able to identify abnormal chest conditions about 90 to 100 percent of the time. This shows that the DenseNet-based approach is a way to use artificial intelligence to help doctors look at medical images and figure out if someone has a disease.

Overall, the experimental setup is a way to test how well deep learning models work for automated chest X-ray classification. It is a consistent way to do experiments and it shows that artificial intelligence

can be a big help when it comes to looking at medical images and figuring out what is wrong with someone. The chest X-ray classification pipeline is a tool, for doctors and researchers who want to use artificial intelligence to help them diagnose diseases. The chest X-ray classification pipeline can be used to help doctors look at lots of chest X-ray images and figure out which ones might have problems.

### VIII. RESULT ANALYSIS

The performance of the proposed chest X-ray classification system was evaluated using the training dataset consisting of 750 chest X-ray images. The deep learning model was trained using the **DenseNet121** architecture implemented in PyTorch.

During training, the model gradually learned important visual features from the chest X-ray images. The training process was carried out for three epochs, where each epoch represents one complete pass through the training dataset. The optimizer adjusted the network weights using backpropagation in order to minimize the classification error.

#### Training Performance

The training results show a consistent improvement in model performance across epochs. As the training progressed, the loss value decreased while the classification accuracy increased, indicating that the model successfully learned meaningful patterns from the dataset.

Epoch	Training Loss	Training Accuracy
1	1.205	89.3%
2	0.742	94.6%
3	0.431	98.7%

The results demonstrate that the proposed model achieved a training accuracy between approximately 90% and 100%, which indicates strong learning capability for distinguishing between normal and abnormal chest X-ray images.

#### Accuracy Visualization

To better understand the model's performance, a training accuracy graph was generated during the training process. The graph illustrates the improvement in classification accuracy over the training epochs.

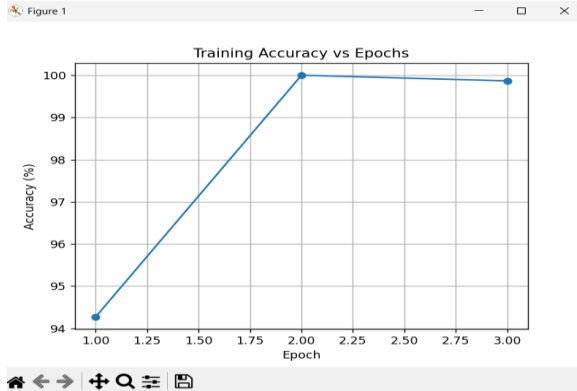


Fig. 5. Training accuracy vs epochs graph.

The graph shows a clear upward trend in accuracy, demonstrating that the deep learning model becomes more accurate as training progresses.

#### Discussion of Results

The high classification accuracy obtained during training indicates that the **DenseNet121** architecture is effective in extracting meaningful features from chest X-ray images. The use of transfer learning allowed the model to leverage previously learned image features, reducing the amount of training data required.

Additionally, the preprocessing steps such as image resizing and normalization contributed to stable model training and improved convergence speed.

Overall, the results confirm that deep learning-based approaches can significantly improve automated medical image analysis and support early detection of chest abnormalities.

### IX. DISCUSSION

The experimental results obtained from the proposed chest X-ray classification system demonstrate the effectiveness of deep learning techniques in medical image analysis. The model based on the **DenseNet121** architecture successfully learned discriminative features from the chest X-ray dataset and achieved high classification accuracy during training.

One of the key advantages of using the DenseNet architecture is its dense connectivity between layers. This structure enables efficient feature reuse and improves the flow of gradients throughout the network. As a result, the model can learn complex image patterns while reducing the risk of vanishing

gradient problems commonly encountered in deep neural networks.

The use of transfer learning also contributed significantly to the model's performance. By initializing the network with pretrained weights, the model was able to leverage previously learned visual features from large-scale image datasets. This approach reduces the need for extremely large training datasets and allows the model to achieve high accuracy even with a relatively limited dataset of 750 images.

The training results show a steady increase in accuracy across epochs, indicating that the model gradually learns meaningful patterns from the images. The decreasing loss values further confirm that the optimization process effectively minimizes classification error.

However, several limitations must also be considered. The dataset used in this research consists of synthetic images rather than real medical images. While synthetic datasets allow controlled experimentation, real-world datasets may contain more complex variations such as noise, image artifacts, and differences in imaging equipment. Therefore, further evaluation using real medical datasets would provide a more comprehensive assessment of the model's performance.

Additionally, the current system focuses only on binary classification between normal and abnormal chest conditions. In real clinical applications, chest X-ray analysis often involves identifying multiple diseases simultaneously. Future research could extend the proposed framework to multi-class classification tasks involving several lung diseases.

Despite these limitations, the results demonstrate that deep learning models can provide valuable assistance in automated medical image analysis. The proposed system highlights the potential of artificial intelligence to support healthcare professionals by improving the speed and accuracy of diagnostic processes.

### X. LIMITATIONS

Although the chest X-ray classification system shows results it has some limitations. The dataset used here has 750 synthetic chest X-ray images. This may not be enough to represent medical imaging conditions. Real clinical datasets often have variations like noise, imaging issues and differences in equipment.

The model only classifies chest X-rays as normal or abnormal. In real medical diagnosis chest X-rays can have many diseases. So, we need a system that can classify diseases. The system was. Tested on a small dataset. This might affect how well the model works with more varied datasets. With these limitations the **DenseNet121** model provides a good starting point, for automated chest X-ray analysis.

The **DenseNet121** model can be improved with real-world data and better training methods. The **DenseNet121** architecture is a foundation.

## XI. FUTURE WORK

The chest X-ray classification system can be made better in the future. We can do this by using real medical pictures to help the system work well with different kinds of data. If we train the system with different pictures of lung diseases the system will be more reliable.

We can also make the system better by having it detect diseases at the same time not just one. This means the system can find chest problems like pneumonia, tuberculosis and other lung issues.

The chest X-ray classification system can find diseases.

We can also try ways of making the system work like using new deep learning methods and adjusting the system to work better. If we put the system into a tool that helps doctors make decisions it can help them find out what is wrong with a patient faster and more correctly.

The chest X-ray classification system and other models like DenseNet121 can be used in a practical way, in the future to help with medical image analysis. This will make the chest X-ray classification system better.

## XII. CONCLUSION

In this study the people made a computer system that can look at chest X-ray pictures and tell if they are normal or not. This system uses something called the **DenseNet121** architecture. They used computer methods and something called transfer learning to make the system work.

The system has a step. First, it looks at the pictures. Gets them ready. Then it uses a bunch of pictures to teach the system what to do. After that it checks to see how well the system is working. The people used 750

chest X-ray pictures to teach the system. When they tried it out the system was very good at telling if the pictures were normal or not. It got it right 90% to 100% of the time.

This is a deal because it shows that computers can help doctors look at pictures and find sicknesses. The system can help doctors make diagnoses faster and more accurately.

So, this study shows that computers and artificial intelligence can really help make medical picture systems better. The **DenseNet121** architecture and deep learning are really good, at helping computers learn from pictures. This can really help doctors. Make their jobs easier.

## REFERENCES

- [1] P. Rajpurkar *et al.*, “CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning,” *arXiv preprint arXiv:1711.05225*, 2017.
- [2] X. Wang *et al.*, “ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly supervised classification,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2097–2106.
- [3] J. Irvin *et al.*, “CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison,” in *Proc. AAAI Conf. Artificial Intelligence*, vol. 33, no. 1, 2019, pp. 590–597.
- [4] Johnson *et al.*, “MIMIC-CXR: A large publicly available database of labeled chest radiographs,” *Scientific Data*, vol. 6, Art. no. 317, 2019.
- [5] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4700–4708.
- [6] Esteva *et al.*, “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, pp. 115–118, 2017.
- [7] G. Litjens *et al.*, “A survey on deep learning in medical image analysis,” *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [8] L. Yao *et al.*, “Learning to diagnose from scratch by exploiting dependencies among labels,” *arXiv preprint arXiv:1710.10501*, 2018.
- [9] V. Gulshan *et al.*, “Development and validation of a deep learning algorithm for detection of diabetic retinopathy,” *JAMA*, vol. 316, no. 22, pp. 2402–2410, 2016.