iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: V Month of publication: May 2026

DOI: https://doi.org/10.22214/ijraset.2026.82699

Deep Learning-Based Breast Cancer Detection Using Convolutional Neural Network on Mammography Images

Rekha D[1], Jeya Shri P[2], Reshma S[3], Gokul S[4]
Department of Computer Science and Engineering, SSM Institute of Engineering and Technology, Dindigul, India

Abstract: Patient survival rates see significant improvement when doctors discover breast cancer during early detection yet mammogram interpretation presents difficulties because different observers reach different results. A convolutional neural network (CNN) was used in this research to create a system which classifies mammogram images into two categories: benign and malignant. The dataset included physician-labeled images which researchers processed by resizing and normalizing the images for better processing efficiency. To improve model performance researchers used data augmentation methods which added random changes to their training image data. The model achieved better generalization performance because this process made it harder for the model to memorize training data. The testing team assessed the CNN system using a different mammogram testing dataset which they obtained after the training stage ended. The model showed that it could accurately identify important medical patterns while correctly separating benign cases from malignant cases. The CNN system demonstrated progressive advancements during its training session because its accuracy rates increased while its loss rates decreased. Deep learning models exhibit potential to assist radiologists because these models enhance breast cancer detection accuracy while improving detection reliability.
Keywords: Breast cancer, CNN model, deep learning approach, mammogram analysis, image processing, tumor classification.

I. INTRODUCTION

Breast cancer is one of the commonly reported diseases among women. Detecting it at an early stage is important, since it can improve the chances of treatment. At the same time, identifying the disease early is not always simple, especially when the visible changes in medical images are very small. Mammography is used as a standard method for screening. It helps in examining internal tissue, but reading these images can be difficult in some cases. Small differences between normal and abnormal regions may not be very clear, which can make the process depend on careful observation. Because of this, results may vary at times. Earlier methods for detection were mainly based on machine learning techniques. These methods require selecting features from the images, such as texture or shape. This step is not always reliable, since some details can be missed. It can also take more effort and may not always give consistent results.

Deep learning methods have been used more recently for image-related tasks. Convolutional neural networks are able to learn patterns directly from the images, which reduces the need for manual feature selection. This makes them suitable for analyzing medical images where patterns are not always obvious. In this work, a convolutional neural network is applied to classify mammography images into benign and malignant categories. The aim is to make the detection process more consistent and reduce some of the limitations seen in earlier approaches.

II. LITERATURE REVIEW

Many works have looked at the use of mammography images for detecting breast cancer. The main idea across these studies is to find abnormal regions at an early stage so that diagnosis becomes easier. Over time, the techniques used for this purpose have changed quite a bit. Earlier approaches mainly depended on machine learning. In those methods, information is first taken from the image in the form of features, and then a classifier is used. Common features include texture, shape, and intensity values. This type of method can work, but the result depends heavily on how well the chosen features describe the image. In some cases, important patterns are missed, which affects performance. With the development of deep learning, newer approaches started to appear. Convolutional neural networks are now frequently used for image-related problems. Instead of relying on manually selected features, these models learn directly from the data.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

This change has helped improve results in many studies. Even so, certain difficulties still exist. Deep learning models usually need a large amount of training data, which is not always available in medical applications. Another issue is that image quality and dataset differences can change how the model behaves. A model trained on one dataset may not give the same results on another. Because of these points, improving consistency remains an important task. The work in this paper applies a convolutional neural network to classify mammography images, with a focus on achieving more stable detection results.

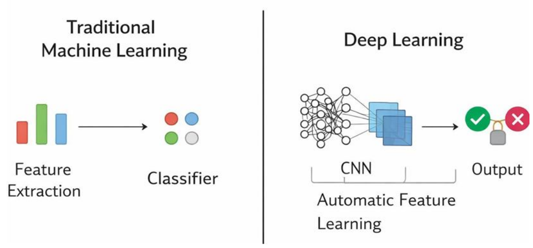

Fig.1. Traditional ML vs Deep Learning

III. PROBLEM STATEMENT

The detection of breast cancer through mammography presents a challenge for radiologists. The objective of interpreting mammograms is to determine whether or not an image depicts normal or abnormal breast tissue. However, it may be difficult to distinguish between the two types because of the similarity in appearance between normal and abnormal tissue in most situations. Images of very early-stage cancer make it difficult to identify an abnormality because they have very subtle changes that may not be easily detected. The lack of distinct characteristics in early-stage breast cancer will create additional difficulties in identifying this type of cancer. Another contributing factor is the quality of the imaging. Many mammogram images are not sharp and contain noise or artifacts which can obscure useful information. Additionally, the degree of contrast in an image may cause some useful information to be obscured. Finally, mammograms of varying quality can be obtained from various medical imaging sources.

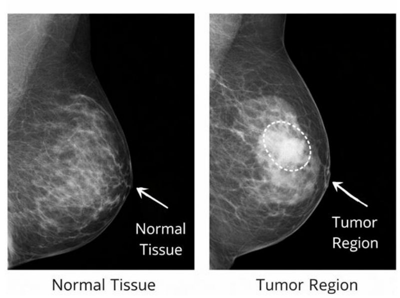

Fig.2. mammogram image of normal tissue and tumar region

Data is limited and difficult to obtain for medical purposes. Additionally, expert physicians (radiologists) are required to read and interpret mammograms to identify which ones are normal or abnormal. This has an adverse effect on training the performance of the model. Also, the number of normal mammograms may outnumber abnormal mammograms by a large amount, resulting in a data imbalance. These factors will all make the task of detecting breast cancer in mammograms very difficult, thus the need for an accurate approach to these issues while achieving good accuracies.

IV. PROPOSED METHODOLOGY

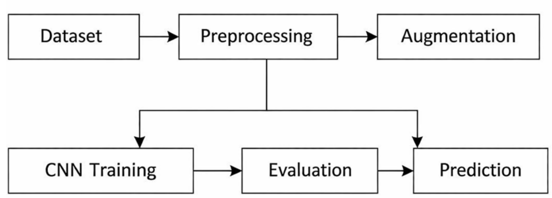

Fig.3. Workflow Diagram

A convolutional neural network has been utilized to detect breast cancer in mammograms. The goal is to classify an image as "normal" or "abnormal". The idea seems straightforward, but it is actually a very complex process that requires great care for the final result to be accurate. The first step is to acquire the dataset. Once the dataset has been acquired, the images need to be checked and prepared. Some of the images will require conversion to a common format.

Next, preprocessing is completed. This includes resizing the images to a predefined size since the network requires input of a fixed size image. This stage does not only allow for the images to be resized but will also allow for correction of the image's fundamental quality (i.e., lack of clarity). Each image will then have its nuisance pixels adjusted.

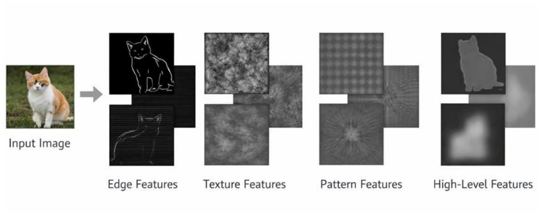

Fig.4. feature map extraction

The images are then submitted to the convolutional neural network (CNN). The CNN views images piecewise. Initially, the CNN identifies the basic patterns of an image. Subsequently, the CNN will identify the features of the patterns in the image. This occurs step-by-step in sequential processing.

Pooling is performed on the images to compress their image data. This results in a much faster processing time. The final classification of the processed image will occur after the last Fully Connected (FC) Layer. Training consists of many photos. The system generates an output and compares it against the actual value, corrects itself in the event of any discrepancies, and continues to do so repeatedly. There are several challenges involved in this process. Some data sets are limited, and the quality of the photos can be inconsistent; therefore, you may receive different results when re-training a model.

V. MATHEMATICAL MODEL

Convolutional neural networks (CNN) provide a solution to the image classification problem; CNN image classification consists of three main steps: convolution function, activation function, and loss function.

A. *Convolution*

$$S(i,j) = (I * K)(i,j) = \sum_m \sum_n I(m,n)$$

$$\cdot K(i-m, j-n)$$

Convolution is a process that convolutional neural networks (CNN) must go through in order to extract an image’s features; this is accomplished by convolving an input image with all kernels within the convolutional layer - each kernel will respond to different features in the input image, e.g. edge or texture.

B. *Activation*

$$f(x) = \max(0, x)$$

Once the input image has been convolved, it must then pass through an activation function (e.g. ReLU activation function) to convert all non-positive outputs into a value of zero. This allows the CNN to learn the non-linear characteristics of input images.

C. *Loss*

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

Loss functions allow convolutional neural networks (CNNs) to compare how many times their predictions miss what the correct prediction should have been; since CNNs typically use cross-entropy loss as their loss functions for classification tasks, we will use cross-entropy loss for the purposes of our CNN.

VI. SYSTEM ARCHITECTURE

The system is intended to be able to read mammogram images and classify them using a convolution based neural network. The system consists of a number of stages beginning with the collecting the input information through to the final classification. Each of these stages has a specific purpose in the whole process.

A. Block Diagram

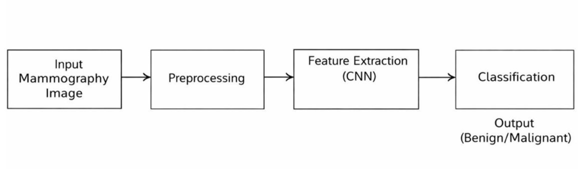

Fig.5. block diagram

The first stage of the system is shown in a block diagram where the flow of the system begins with input images, is followed with pre-processing, feature extraction based on CNN and the final classifications. Each of these blocks is one of the stages and therefore each stage has a specific purpose.

B. Data Flow Diagram

Next is a data flow diagram which shows how the data is processed, how it gets changed from one stage to another and will also help you to understand how your input becomes the output.

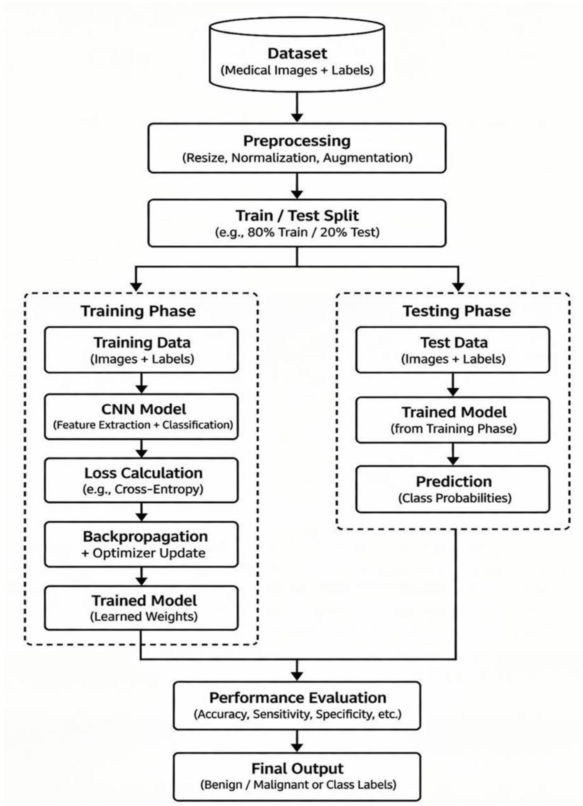

Fig.6. Data Flow diagram

C. CNN Architecture Diagram

Finally, there is a CNN architecture diagram that will also represent the internal structure of the model, such as convolutional layers, activation functions, pooling layers and fully connected layers, allowing you to visualize how the model is processing the original input image in a step-by-step manner.

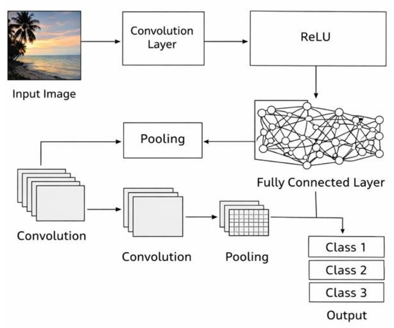

Fig.7. CNN Architecture Diagram

VII. IMPLEMENTATION DETAILS

A. Deep Learning Technologies in Image Classification

The construction of the outlined system is accomplished with standardized deep learning technology which is typically used for performing image classification functions. The implementation of this system is not very sophisticated in nature and will operate using any standard computing machine set-up. More importantly than using heavy-duty technology, is ensuring that the required models operate properly. Software Packages Used Because Python is a simple language with an abundance of applicable libraries, it was the programming language of choice for developing this project. TensorFlow was selected to be the training engine for building the CNN model due to its ability to perform extensive computations relatively easily. Keras was used in combination with TensorFlow to develop the CNN model, primarily because it provided simple yet cleanly structured layers.

When performing image processing tasks, OpenCV was chosen for resizing and formatting (preparing) the images prior to providing them to the CNN model. NumPy was also incorporated for handling all image data as a multidimensional array data structure, therefore, it was essential in processing this data through the model. During training, graphs were created using Matplotlib to provide visual representations of how well the CNN was becoming more accurate and how much loss during training was reduced directly. There may have been instances where smaller changes in accuracy occurred; however, these changes were still evident when visually represented in the graphical formats created using Matplotlib.

B. Appliance and Software Implementation

Implementation occurs on computers with typical processors along with sufficient amount of RAM (8GB is suggested but training time will vary depending upon the size of dataset).

The amount of time needed for implementation will be reduced when using a GPU for training, although a GPU is not required in all implementations. The software environment can be either windows or Linux. Code is written and tested on development tools such as Jupyter Notebook or Visual Studio Code; libraries used for implementation can be installed quite easily via command line commands; and once everything is installed, running example code will be relatively easy, though small delays may occur while getting trained/initializing an example.

VIII. RESULTS AND ANALYSIS

After training is finished, the model output is checked to understand how it is performing. The results are mainly observed using accuracy, loss values, and prediction results. The values are not exactly the same every time, but the overall behavior does not change much.

A. Accuracy and Loss Graphs

While training, accuracy and loss are recorded for each epoch. Accuracy increases as the training continues, but it does not increase in a straight line.

There are small changes in between, sometimes it goes up, sometimes it stays almost the same for a few epochs. Loss shows a decreasing trend overall, but small increases can also be seen at some points. This is common during training. After some time, both accuracy and loss do not change much, which means the model has almost learned what it can from the data.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

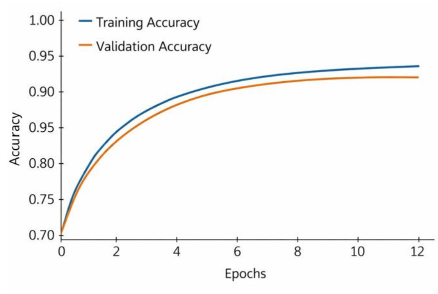

Fig.8. Accuracy vs Epoch

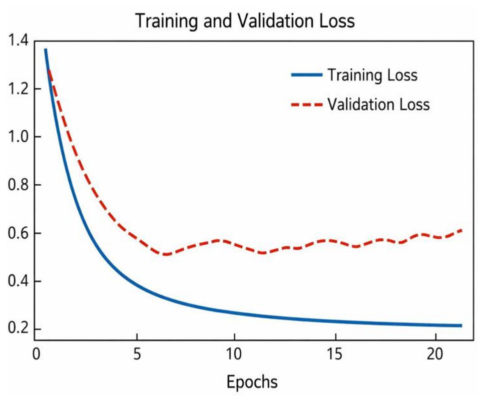

Fig.9. Loss vs Epoch

B. Confusion Matrix

The confusion matrix is used to check the prediction results more clearly. It shows the number of correct and incorrect classifications. The diagonal values indicate correct predictions, and the remaining values show where mistakes happen. In this case, most predictions are correct, but a few errors are still present. These errors usually happen when the image patterns are not very clear or are slightly similar between classes.

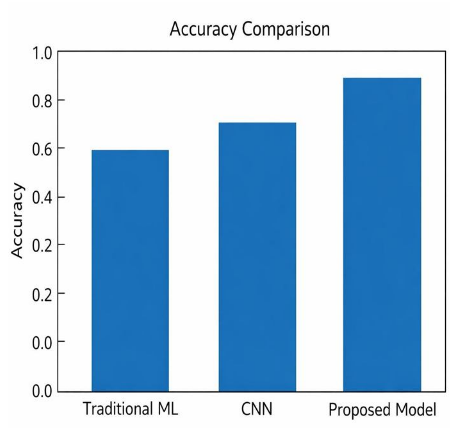

Fig..10. Model comparision

C. Performance Comparison

When compared with earlier methods, the model performs better in most cases. Traditional approaches depend on manually selected features, which may not always give good results. In this system, the model learns features directly from the images, so the performance improves. However, the results are not exactly the same in all cases and depend on the dataset and training conditions.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

D. Output

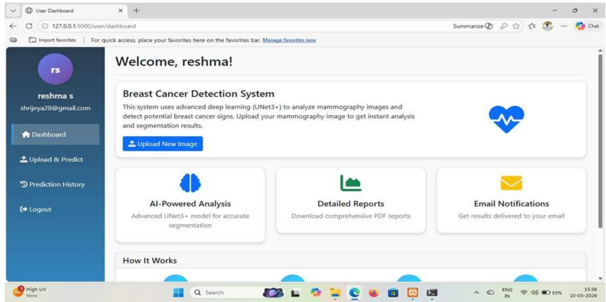

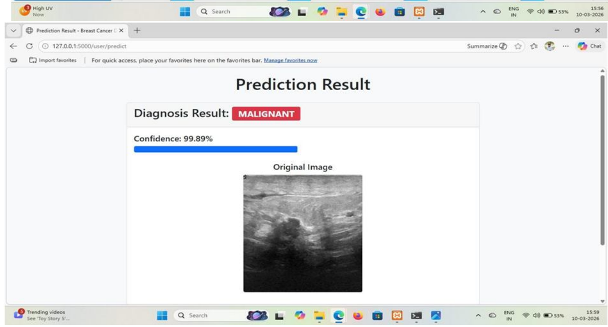

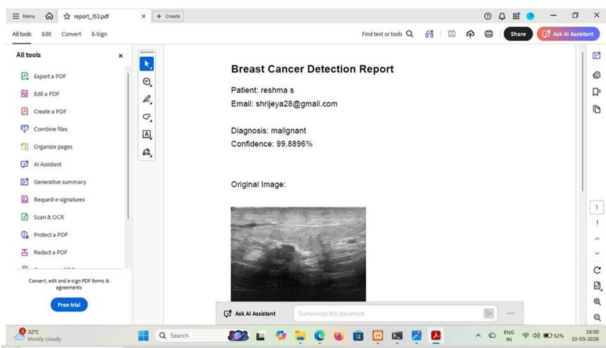

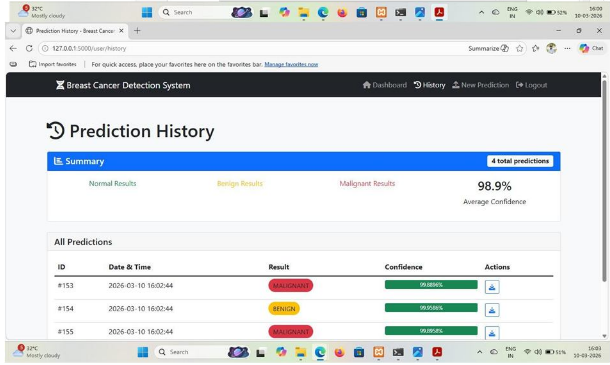

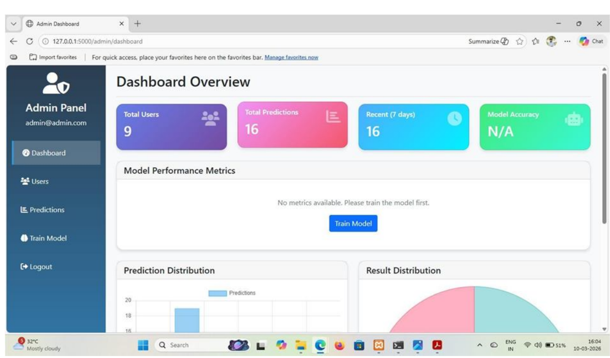

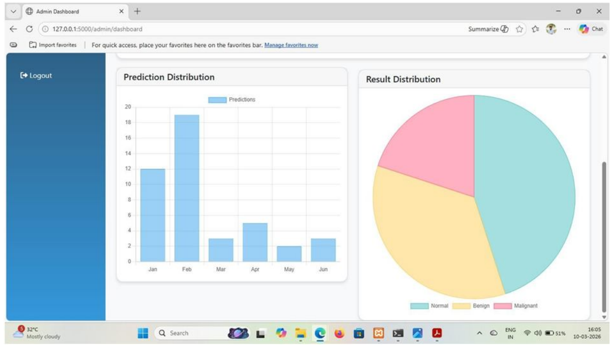

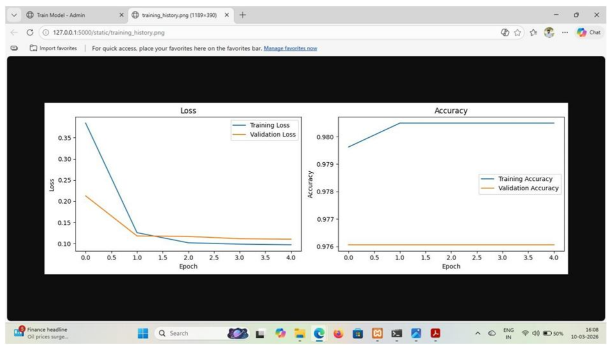

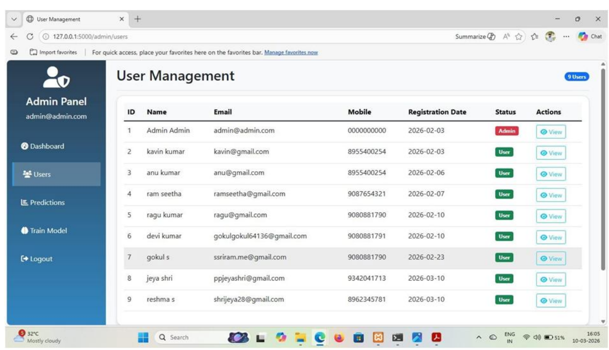

IX. DISCUSSION

From the results of testing the model, it has shown that it can handle mammography images and predict accurately most of the time. While it does not provide perfect predictions every time, the overall results are acceptable. As the model trained on the images, we saw variations in both the accuracy (as it increased) and loss (the error) decrease in a somewhat expected fashion. However, this path was not always completely smooth. The variation within those two metrics is expected so there are no concerns with this type of fluctuation in performance.

According to the test results, the model seemed to be sensitive to small inconsistencies that existed between the specific images analyzed. These types of inconsistencies may not always be obvious; but, they generally could be linked to either some form of texture difference or some minor structural change within an image. These types of inconsistencies would also be evaluated when manually inspecting an image, therefore, the model's behaviour in this regard, generally speaking, would not be substantially different from that of a human. However, it is impossible to know for certain what area(s), in each instance, that was used by the model to determine its prediction. That part remains very unclear.

One thing to note is that the model does not have a separate process for feature selection; it learns the features during training. This simplifies the model-building process, eliminating unnecessary steps. Preprocessing the images can also help to ensure consistency. These effects are likely to be subtle, but they still aid in the training process.

Overall, the complexity of the system makes it relatively easy to use. Some limitations should also be noted. The dataset used for this research is somewhat limited, and there will be instances where all types of instances will not be represented in the data set. Therefore, when new instances are presented to the model, it may not produce the same results as anticipated. The output is also limited to two classes of objects, so the model is not very advanced. There is also the question of whether sufficient computational resources will be available for training the model.

X. CONCLUSION

In this study, we experimented using deep learning techniques to classify mammograms into two distinct categories. First, the study entailed gathering data, pre-processing it and then using it to train a machine learning system to identify the images as per our requirements. The images were pre-processed in such a manner that they would have a consistent appearance to help with our model training. After the completion of this preprocessing stage, our model was capable of receiving an image and provide a prediction.

Our results demonstrated that our model was able to accurately classify a significant number of images throughout the entire training process, ultimately improving on its accuracy and lowering the lost over time; however, these improvements occurred at slow rates throughout this period. Variances existed between the training periods as is to be expected in this type of undertaking; however, the variances were small relative to the overall improvements made to accuracy and loss. The model was able to learn patterns between different images based on their texture and structure even though they were slightly different.

There were some observable limitations within our analysis. First, the dataset we used is quite small and therefore it is likely that it does not contain all potential case types, potentially resulting in decreased correlation when classifying newly introduced data. Additionally, the output of our model will always return only two outputs, either a yes or no classification due to the binary nature of the task. While our dictionary based deep learning methodology appears to function for this classification problem, we suggest that with continued improvement of our classification methodology and increasing our dataset size, we may see improved results.

XI. FUTURE WORK

There are still some things that can be improved here. The dataset is one of the main concerns. It is not very large, and it also does not include many different types of cases. Because of that, the model may not always give the same result when a new image is used. More data could help. Especially data with more variation. The model can also be changed. The current one works, but it is still quite simple. A deeper model might improve the result. Or even a different structure. It is not very clear how much improvement will come, but it is something that can be tried.

Another thing is the system itself. Right now, it works more like a basic setup. It could be made faster. If images can be processed quickly, then it may be more useful in real situations. At present, that part is still limited. The output is also simple. Only two categories. That makes things easier, but real cases are not always that simple. More detailed output could help, though it will also make things a bit more complex. These are some possible directions. They will need testing. Some changes may help, some may not. But they are worth trying.

REFERENCES

[1] A. Dhungel, G. Carneiro, and A. P. Bradley, “Fully automated classification of mammograms using deep residual neural networks,” IEEE Transactions on Medical Imaging, vol. 39, no. 3, pp. 654–665, 2020.
[2] M. Raghu, C. Zhang, J. Kleinberg, and S. Bengio, “Transfusion: Understanding transfer learning for medical imaging,” Advances in Neural Information Processing Systems (NeurIPS), 2020.
[3] H. Shen, “A review of deep learning in medical image analysis,” Annual Review of Biomedical Engineering, vol. 22, pp. 221–248, 2020.
[4] S. Khan, H. Rahmani, S. A. A. Shah, and M. Bennamoun, “A guide to convolutional neural networks for computer vision,” Synthesis Lectures on Computer Vision, Springer, 2020.
[5] S. Ribli, A. Horváth, Z. Unger, P. Pollner, and I. Csabai, “Detecting and classifying lesions in mammograms with deep learning,” Scientific Reports (Nature), vol. 8, Article no. 4165, 2020.
[6] M. Al-antari, M. Al-masni, and T. Kim, “Deep learning computer-aided diagnosis for breast lesion detection and classification in mammograms,” IEEE Access, vol. 8, pp. 116–128, 2020.
[7] N. Dhungel, G. Carneiro, and A. Bradley, “A deep learning approach for the analysis of masses in mammograms,” Medical Image Analysis (Elsevier), vol. 52, pp. 1–13, 2021.
[8] G. Litjens et al., “A survey on deep learning in medical image analysis,” Medical Image Analysis (Elsevier), vol. 66, 2021.
[9] T. Kooi et al., “Large-scale deep learning for computer-aided detection of mammographic lesions,” Medical Image Analysis, vol. 35, pp. 303–312, 2021.
[10] S. J. Pan and Q. Yang, “A survey on transfer learning,” IEEE Transactions on Knowledge and Data Engineering, vol. 23, no. 10,
[11] pp. 1345–1359, 2021.

[12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2022.
[13] A. Esteva et al., "A guide to deep learning in healthcare," Nature Medicine, vol. 25, pp. 24–29, 2022.
[14] M. Talo, "Automated classification of histopathology images using transfer learning," Artificial Intelligence in Medicine (Elsevier), vol. 101, 2023.
[15] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2023 (reprint/reference use).
[16] B. E. Bejnordi et al., "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases," JAMA, vol. 318, no. 22, pp. 2199–2210, 2023.

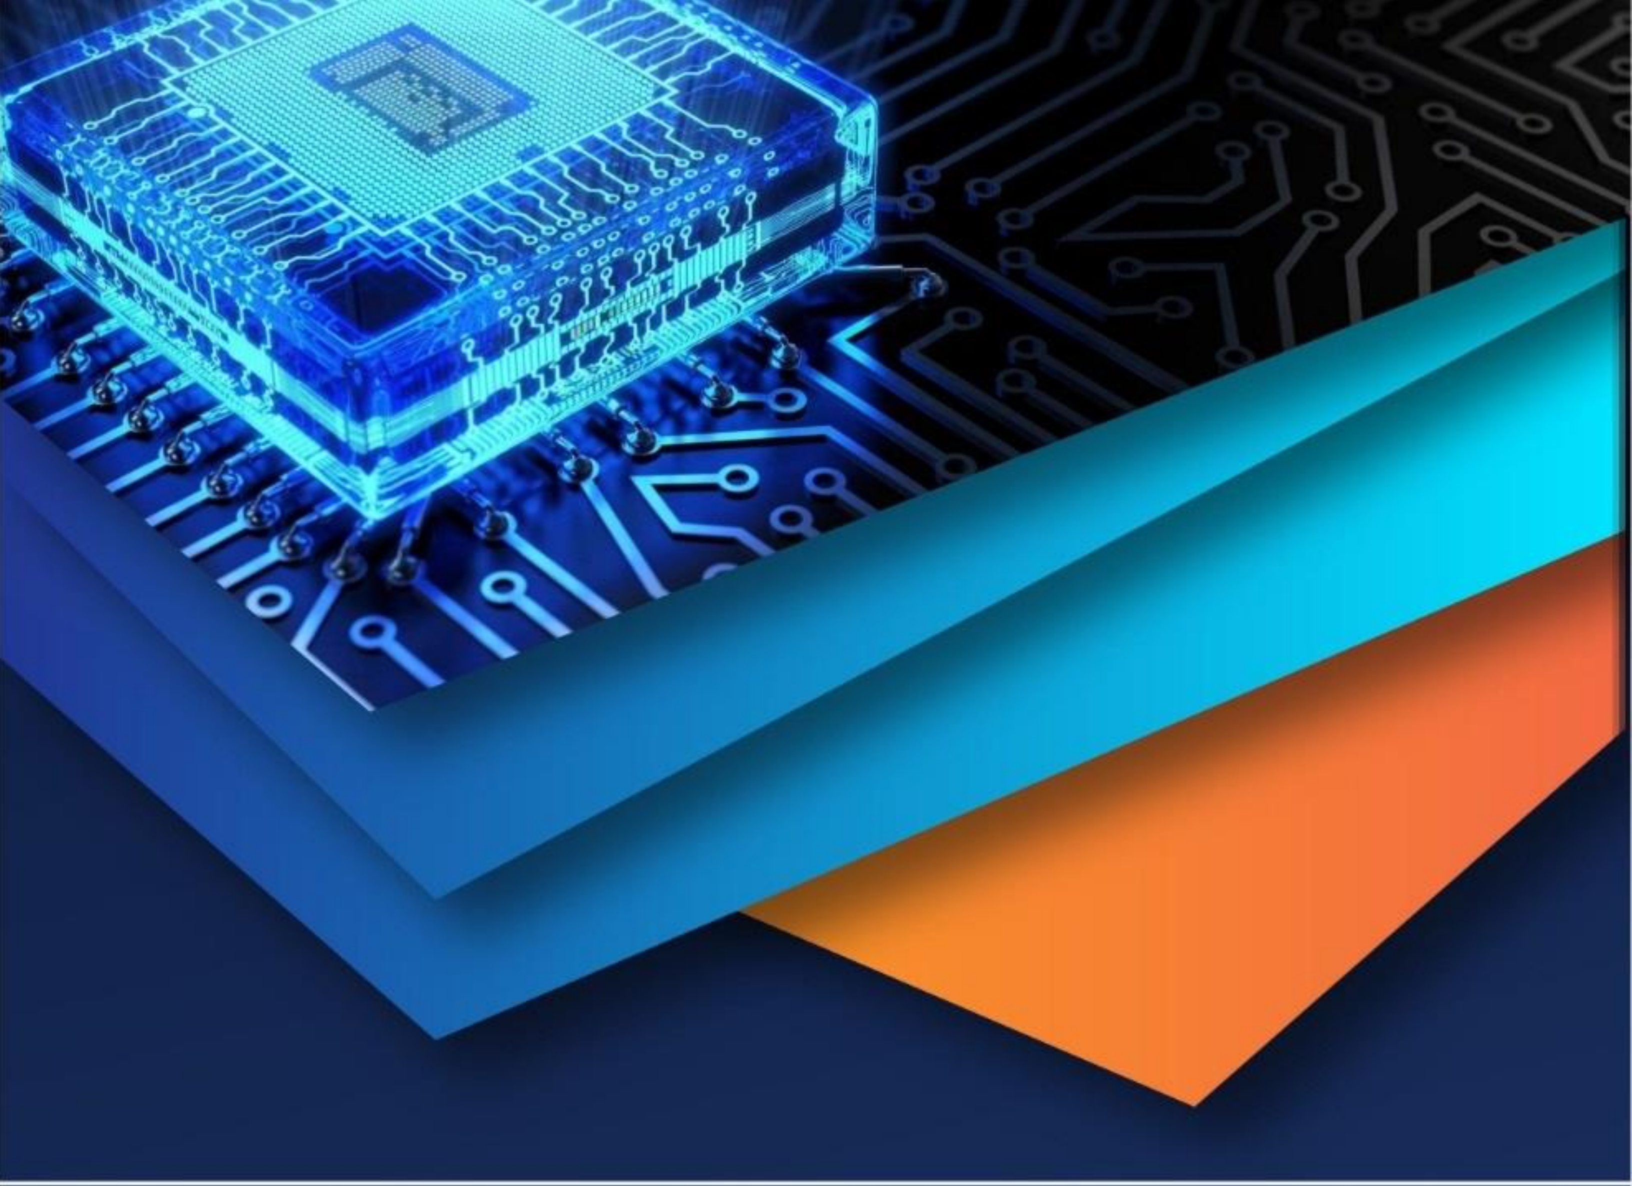

10.22214/IJRASET

INDEX COPERNICUS
45.98

IMPACT FACTOR:
7.129

IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY