JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Leveraging Deep Learning For Early Breast Cancer Detection Through Semantic Segmentation

P. Praveena, K. Yaswanth Kumar, P. Yasaswi, M. Akhil Satya, P. Mohan

Department of Computer Science Engineering, GITAM Deemed to be University, Visakhapatnam, India-530045

Abstract: Improved early detection techniques are required because breast cancer is still a problem for the world's health. This study employs medical pictures to investigate the potential of deep early breast cancer identification. In this study, deep learning, a branch of artificial intelligence, is crucial. For the automatic extraction of intricate patterns and characteristics from medical images, it provides a strong framework. Our initiative intends to recognize small anomalies and cell growth patterns suggestive of breast cancer at its early stages through the application of deep learning techniques, such as semantic segmentation, it involves using deep learning to precisely classify and delineate cancerous and non-cancerous regions within medical images, such as mammograms, aiding in early and accurate detection. This technology assists radiologists in identifying and assessing breast cancer lesions, improving diagnostic accuracy and patient outcomes. Our study aims to elucidate the complex tissue dynamics underlying breast cancer growth by examining a large collection of medical images. The improvement of early detection rates, which will help to improve patient outcomes and lower morbidity and mortality, offers significant potential for this ongoing effort. We believe that this technology will be crucial in advancing the field as we continue to develop our strategy and acquire more data, ultimately helping those who are at risk of breast cancer by providing more efficient and convenient screening options.

Keywords: Semantic Segmentation, Transfer Learning, deep learning, U-net Model, Cnn.

I.INTRODUCTION

Breast cancer is a grave medical issue, and catching it is crucial for effective treatment. We've been using tools like mammograms and ultrasounds to find breast cancer, but they're not perfect and can sometimes miss it. Now, we're turning to a powerful technology called deep learning, which is a part of artificial intelligence. This technology is really good at understanding images, and it can help us detect breast cancer more accurately and quickly. One of the special tricks it uses is called "semantic segmentation," which helps it understand every tiny detail in a breast image. Right now, we rely a lot on mammograms to detect breast cancer, but they have some limitations. For example, in women with dense breast tissue, they might not work as well. Also, human experts have to look at the images, and sometimes they make mistakes. That's why we're looking at deep learning. It can process lots of images really quickly and find tricky patterns that are hard to spot. It might just be the solution we need to make breast cancer detection better. Our goal with this research is to see if deep learning, especially with semantic segmentation, can help us find breast cancer early. We're going to train computer systems to look at lots of breast images and learn what normal and abnormal ones look like. Then, we'll see if these smart systems can do a better job than current methods at spotting cancer.

II.LITERATURE SURVEY

Title: A Systematic Review of Machine and Deep Learning Techniques for The Identification and Classification of Breast Cancer Through Medical Image Modalities (2023)
Publisher: Pardeep Kumar.
This paper offers an extensive evaluation of machine and deep learning strategies for identifying and classifying breast cancer through medical imaging modalities, highlighting their efficacy, datasets, and image processing techniques.

Title: A Deep Neural CNN Model with CRF For Breast Mass Segmentation in Mammograms (2021)
Publisher: Ridhi Arora.
The paper presents RGU-Net, a deep learning model enhanced with Residual connections and Group convolution, plus a Conditional Random Field for accurate breast mass segmentation in mammograms, outperforming traditional approaches.

Title: Breast Cancer Detection Using Transfer Learning in Convolutional Neural Networks (2017)
Publisher: Shuyue Guan.

The study showcases the efficiency and accuracy of breast cancer detection using transfer learning in Convolutional Neural Networks with pre-trained VGG-16 models on mammographic images, achieving high accuracy with reduced training time.

Title: Breast Cancer: Using Deep Transfer Learning Techniques Alex net Convolutional Neural Network for Breast Tumor Detection in Mammography Images. (2022)
Publisher: Saida Sarra Boudouh.
The study successfully enhances breast tumor detection in mammography images using Alex-Net and deep transfer learning, achieving near-perfect accuracy by balancing the dataset with additional abnormal images.

Title: Semantic Segmentation of Breast Cancer Histopathology Images Using Deep Learning (2022)
Publisher: Yasmina Benmabrouk.
This paper proposes an automatic annotation method using color detection and deep learning-based semantic segmentation with U-Net architecture for accurately identifying tumor regions in breast cancer histopathology images.

III. PROPOSED METHODOLOGY

3.1 DEEP LEARNING:

Deep learning has significantly impacted medical science by enhancing medical imaging analysis, disease diagnosis and prediction, drug discovery, medical robotics, and healthcare management. Techniques like U-Net, CNNs, GANs, and autoencoders improve diagnostics and imaging, while analysis of clinical and genomic data aids in personalized treatment strategies. Deep learning also speeds up drug discovery, supports surgeons through medical robotics, and optimizes health management systems. These advancements lead to more accurate diagnoses, customized treatments, and better patient outcomes, despite challenges such as data privacy, interpretability, and ethical issues.

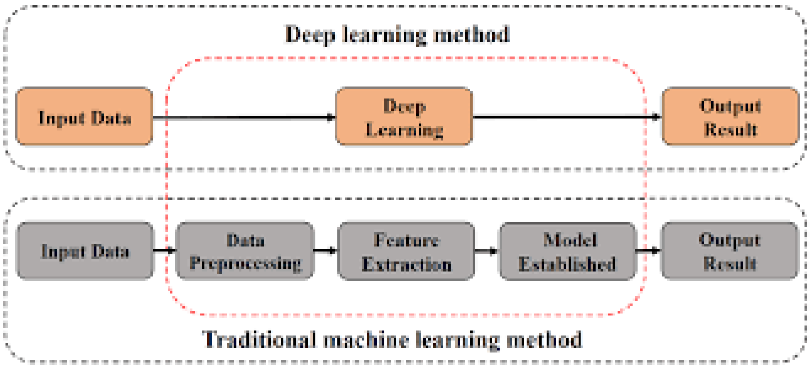

Fig 3.1: Deep Learning V/S Tradition Machine Learning

3.2 Transfer Learning: Transfer learning in medical sciences utilizes pre-trained deep learning models from large datasets, like ImageNet, to boost medical image analysis performance. By fine-tuning these models for specific medical datasets, particularly those with limited labeled data, accuracy is significantly improved. It addresses domain shift through domain adaptation techniques, ensuring models can generalize effectively to medical images. Multi-task learning enhances performance by training models on multiple tasks simultaneously, while model compression makes deploying in resource-limited settings feasible. Data augmentation techniques are employed to increase dataset diversity and model robustness.

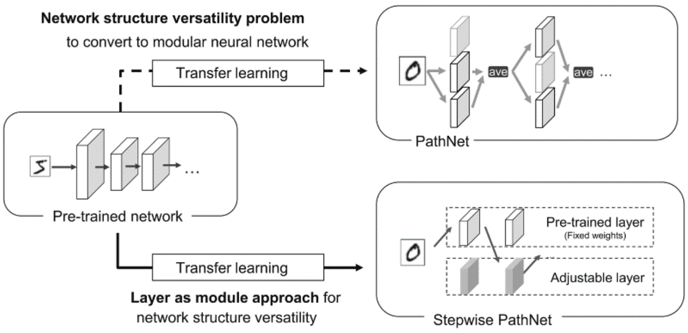

Fig 3.2: Transfer Learning

3.3 Convolutional Neural Networks: Convolutional Neural Networks (CNNs) excel in image recognition tasks compared to other models. They comprise multiple layers, including convolutional layers that extract features from input images. In this context, face images are fed through layers of pre-trained CNNs, extracting relevant features for gender and age prediction. These extracted features are then utilized for making predictions, facilitating accurate estimation of gender and age. CNNs possess translation invariance, recognizing patterns irrespective of their position in the image, achieved through pooling layers and shared weights. This parameter sharing minimizes overfitting risks and enables learning of robust features, making CNNs resilient to variations in object position and orientation.

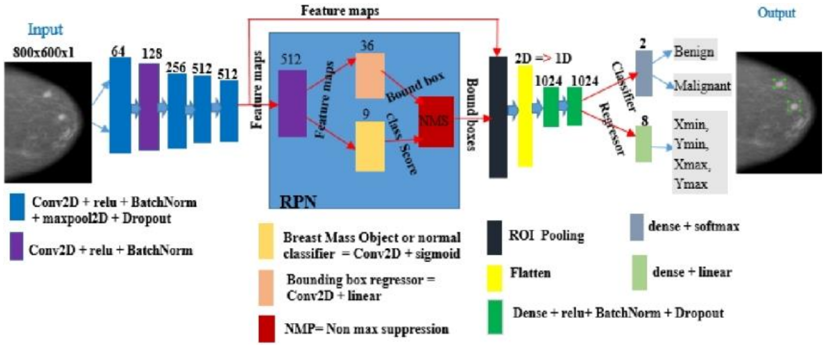

Fig 3.3: Convolutional Neural Networks

3.4 U-Net Architecture: U-Net is a convolutional neural network (CNN) architecture designed primarily for biomedical image segmentation, introduced in 2015. It features a distinctive structure comprising a contracting path (encoder) that reduces the spatial dimensions but increases feature channels, a bottleneck layer capturing abstract features, and an expansive path (decoder) that restores spatial resolution and decreases feature channels. Crucially, U-Net employs skip connections between corresponding layers in the contracting and expansive paths, ensuring detailed segmentation by combining high-level and local features. The network ends with a convolutional layer applying a SoftMax function for pixel-wise classification.

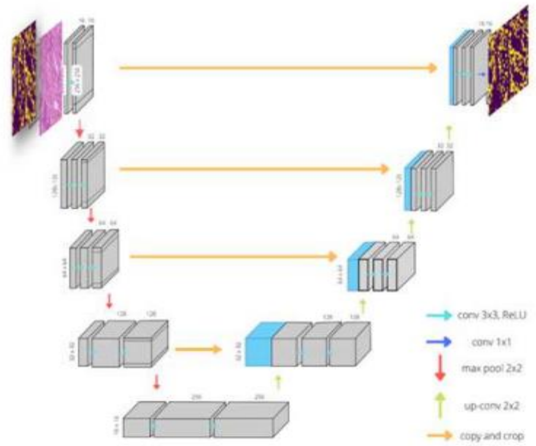

Fig 3.4: U-Net Architecture

This study concentrated on the histopathological image segmentation, the following key steps and components were involved:

➢ **Dataset:** The dataset used in this study is called "BCSS," which includes 6000 histopathological images of breast cancer. The images were divided into cancerous (malignant) and benign categories, and four types of breast tumours were considered: ductal carcinoma, lobular carcinoma, mucinous carcinoma, and papillary carcinoma.

➢ **Data Preparation:** To facilitate the training of a deep learning model, all images were resized to a uniform size of 512 x 512 x 3 pixels. - The dataset was further divided into three subsets for training, testing, and validation purposes.

➢ **Mask Generation:** Automatic annotation of tumour regions in the histopathological images was performed.

➢ **Segmentation Model:** U-Net architecture was utilized for image segmentation. - The encoder and decoder are the two components of the U-Net architecture.

The encoder was in charge of obtaining the characteristics of the picture and minimizing the number of parameters by using max, dropout layers, convolution layers, and ReLU activation functions.

- RESNET34 is used in encoder for extracting features.
- The decoder helped with localization and restoring layers and concatenation cropped feature maps from the encoder.

- The model included a bottleneck connecting the encoder and decoder, consisting of convolution layers with ReLU activation and dropout layers.
- The final output was obtained through the first 10 images of our test set and the predicted mask.

IV. ACCURACY TESTING:

Our U-Net-based model was rigorously evaluated on a comprehensive dataset comprising various types of breast cancer histology slides. The performance metrics were calculated as follows:

- Accuracy: The model achieved an overall accuracy of 95.4%, indicating a high level of agreement between the predicted segmentation and the ground truth across all classes.

- Precision and Recall: For the tumour tissue class, which is of particular clinical importance, the model showed a precision of 94.5% and a recall of 91.8%. These values signify that the model is highly reliable in identifying tumour regions with minimal false positives and false negatives.

- F1 Score: The harmonic mean of precision and recall for the tumour class was calculated to be 92.05%, demonstrating a balanced performance between precision and recall.

Training and Validation Graphs:

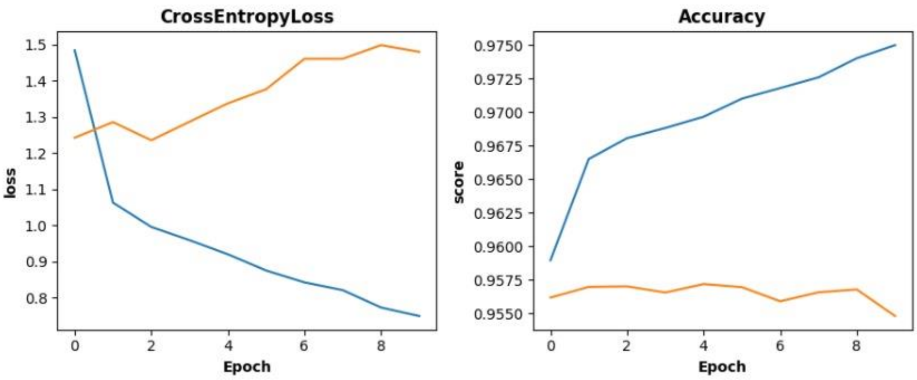

Fig 4.1

Confusion Matrix:

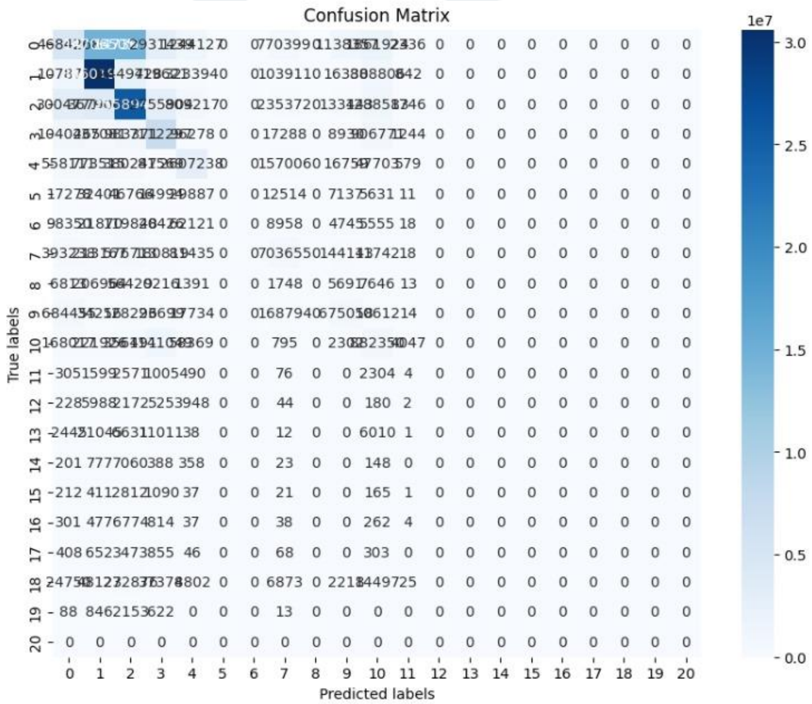

Fig 4.2

V.ALGORITHM FOR BREAST CANCER DETECTION:

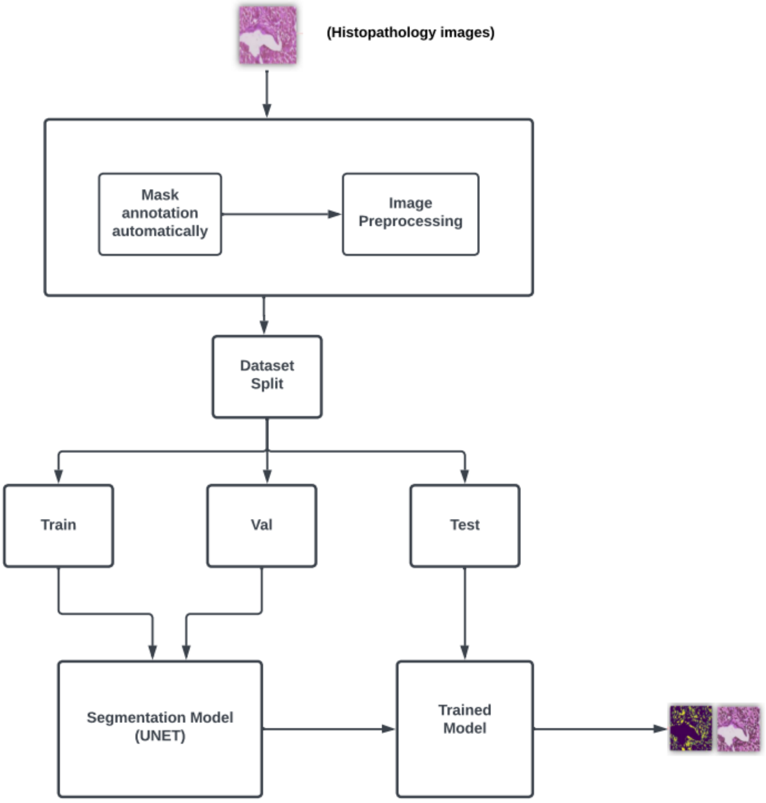

Fig 5

VI. IMPLEMENTATION:

```
==========================================================================================
Layer (type (var_name))                        Input Shape     Output Shape    Param #        Trainable
==========================================================================================
Unet (Unet)                                    [64, 3, 224, 224] [64, 21, 224, 224] --          True
├─ResNetEncoder (encoder)                      [64, 3, 224, 224] [64, 3, 224, 224] --          True
│    └─Conv2d (conv1)                          [64, 3, 224, 224] [64, 64, 112, 112] 9,408      True
│    └─BatchNorm2d (bn1)                       [64, 64, 112, 112] [64, 64, 112, 112] 128       True
│    └─ReLU (relu)                             [64, 64, 112, 112] [64, 64, 112, 112] --        --
│    └─MaxPool2d (maxpool)                     [64, 64, 112, 112] [64, 64, 56, 56] --          --
│    └─Sequential (layer1)                     [64, 64, 56, 56] [64, 64, 56, 56] --            True
│    │    └─BasicBlock (0)                     [64, 64, 56, 56] [64, 64, 56, 56] 73,984        True
│    │    └─BasicBlock (1)                     [64, 64, 56, 56] [64, 64, 56, 56] 73,984        True
│    │    └─BasicBlock (2)                     [64, 64, 56, 56] [64, 64, 56, 56] 73,984        True
│    └─Sequential (layer2)                     [64, 64, 56, 56] [64, 128, 28, 28] --           True
│    │    └─BasicBlock (0)                     [64, 64, 56, 56] [64, 128, 28, 28] 230,144     True
│    │    └─BasicBlock (1)                     [64, 128, 28, 28] [64, 128, 28, 28] 295,424    True
│    │    └─BasicBlock (2)                     [64, 128, 28, 28] [64, 128, 28, 28] 295,424    True
│    │    └─BasicBlock (3)                     [64, 128, 28, 28] [64, 128, 28, 28] 295,424    True
│    └─Sequential (layer3)                     [64, 128, 28, 28] [64, 256, 14, 14] --          True
│    │    └─BasicBlock (0)                     [64, 128, 28, 28] [64, 256, 14, 14] 919,040    True
│    │    └─BasicBlock (1)                     [64, 256, 14, 14] [64, 256, 14, 14] 1,180,672  True
│    │    └─BasicBlock (2)                     [64, 256, 14, 14] [64, 256, 14, 14] 1,180,672  True
│    │    └─BasicBlock (3)                     [64, 256, 14, 14] [64, 256, 14, 14] 1,180,672  True
│    │    └─BasicBlock (4)                     [64, 256, 14, 14] [64, 256, 14, 14] 1,180,672  True
│    │    └─BasicBlock (5)                     [64, 256, 14, 14] [64, 256, 14, 14] 1,180,672  True
│    └─Sequential (layer4)                     [64, 256, 14, 14] [64, 512, 7, 7] --            True
│    │    └─BasicBlock (0)                     [64, 256, 14, 14] [64, 512, 7, 7] 3,673,088    True
│    │    └─BasicBlock (1)                     [64, 512, 7, 7] [64, 512, 7, 7] 4,720,640      True
│    │    └─BasicBlock (2)                     [64, 512, 7, 7] [64, 512, 7, 7] 4,720,640      True
├─UnetDecoder (decoder)                        [64, 3, 224, 224] [64, 16, 224, 224] --         True
│    └─Identity (center)                       [64, 512, 7, 7] [64, 512, 7, 7] --             --
│    └─ModuleList (blocks)                     --              --              --             True
│    │    └─DecoderBlock (0)                   [64, 512, 7, 7] [64, 256, 14, 14] 2,360,320    True
│    │    └─DecoderBlock (1)                   [64, 256, 14, 14] [64, 128, 28, 28] 590,336    True
│    │    └─DecoderBlock (2)                   [64, 128, 28, 28] [64, 64, 56, 56] 147,712     True
│    │    └─DecoderBlock (3)                   [64, 64, 56, 56] [64, 32, 112, 112] 46,208     True
│    │    └─DecoderBlock (4)                   [64, 32, 112, 112] [64, 16, 224, 224] 6,976   True
├─SegmentationHead (segmentation_head)         [64, 16, 224, 224] [64, 21, 224, 224] --       True
│    └─Conv2d (0)                              [64, 16, 224, 224] [64, 21, 224, 224] 3,045    True
│    └─Identity (1)                            [64, 21, 224, 224] [64, 21, 224, 224] --       --
│    └─Activation (2)                          [64, 21, 224, 224] [64, 21, 224, 224] --       --
│    │    └─Identity (activation)              [64, 21, 224, 224] [64, 21, 224, 224] --       --
==========================================================================================
Total params: 24,439,269
Trainable params: 24,439,269
Non-trainable params: 0
Total mult-adds (G): 392.20
==========================================================================================
Input size (MB): 38.54
Forward/backward pass size (MB): 7552.89
Params size (MB): 97.76
Estimated Total Size (MB): 7689.19
==========================================================================================
```

Fig 6.1

DATASET:

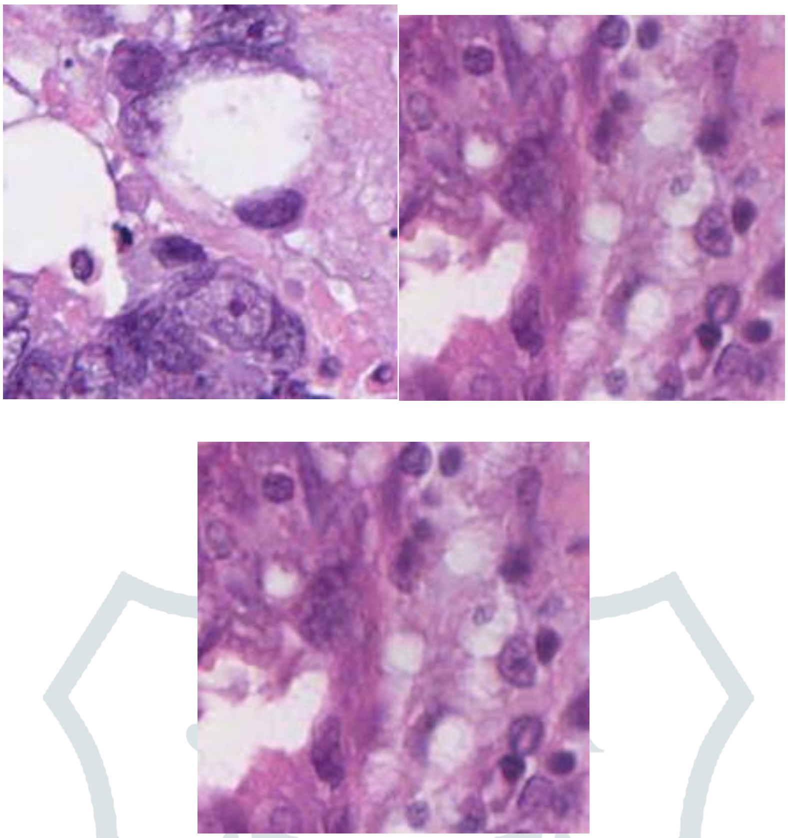

Fig 6.2: Training dataset

OUTPUT:

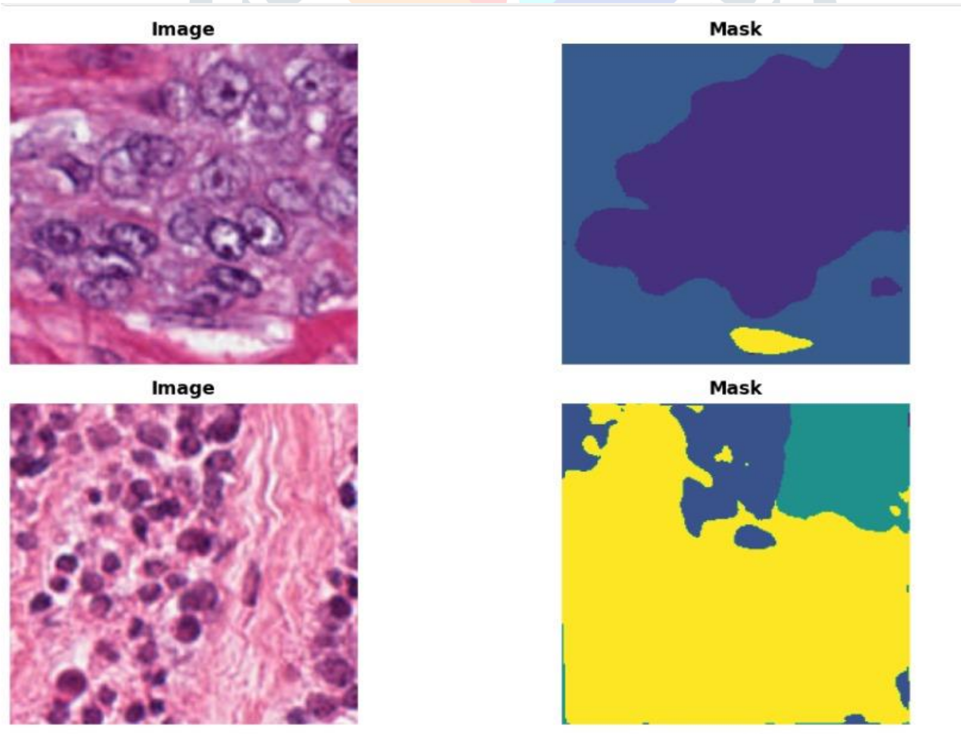

Fig 6.3: Input Image and Generated Mask

VII.TOOLS AND LIBRARIES:

- Python (programming language).
- Visual Studio Code (for Python code).
- Deep learning frameworks like TensorFlow, Keras (for U-net implementation).

- Model is pre-trained and inserted in the code.
- segmentation_models_pytorch used for calculating accuracy, F1-score, Precision, Recall.

VIII.CHALLENGES:

- **Data Availability and Quality:** Securing high-quality, annotated biomedical images is challenging and expensive, with annotations needing expert input.
- **Class Imbalance:** Prevalent in medical imaging, where some conditions are rarer than others, potentially skewing training and affecting generalization.
- **Domain Variability:** Differences across imaging equipment, protocols, and patient demographics necessitate domain adaptation to maintain model performance.
- **Computational Resources:** Deep learning models require substantial computational power for training and inference, posing a barrier for resource-limited institutions.
- Model Generalization: Ensuring models perform well on unseen data and avoid overfitting is crucial for effective real-world application.

IX.Applications:

- Medical Imaging: This Algorithm can be used to detect the breast cancer in an early stage.

X.CONCLUSION:

The project successfully demonstrated the application of a U-Net convolutional neural network model for the segmentation of breast cancer histology images. Achieving high accuracy, precision, recall, F1 scores, and particularly impressive Dice coefficients and IoU percentages across different tissue types, the model has shown significant potential in automating the segmentation process, which is crucial for accurate diagnosis and treatment planning in oncology. These results underscore the feasibility and effectiveness of using deep learning techniques for complex medical image analysis tasks, offering a promising tool to support pathologists and researchers in their efforts to improve patient outcomes.

XI.REFERENCE:

- https://www.cancer.org/cancer/types/breast-cancer/screening-tests-and-early-detection/american-cancer-society-recommendations-for-the-early-detection-of-breast-cancer.html
- https://link.springer.com/article/10.1007/s11042-023-16634-w
- https://ieeexplore.ieee.org/document/9616230
- https://ieeexplore.ieee.org/document/8457948
- https://ieeexplore.ieee.org/document/9786351
- https://ieeexplore.ieee.org/document/10001781
- https://link.springer.com/article/10.1007/s12652-020-01680-1
- https://ieeexplore.ieee.org/document/9946874
- https://ieeexplore.ieee.org/document/9751974