© March 2025| IJIRT | Volume 11 Issue 10 | ISSN: 2349-6002

Skin Cancer Detection and Classification using CNN

Aditi Shankar[1], Alaina Umme Hafiza[2]
[1]*SOET, CMR University*
[2]*SOET, CMR University*

***Abstract*—Skin cancer, particularly melanoma, remains a leading cause of death worldwide, with early detection crucial for improving survival rates. The development of machine learning-based diagnostic tools, particularly Convolutional Neural Networks (CNNs), offers significant potential to assist in non-invasive skin cancer detection. This study explores the integration of CNNs in building an intelligent skin cancer detection system capable of real-time image analysis and user feedback. A mobile application is proposed to offer users the ability to upload or capture images of suspicious skin lesions for immediate analysis. The system's core components include image preprocessing, feature extraction, and classification using machine learning models. Additionally, evaluation metrics such as accuracy, precision, recall and F1 score are used to assess model performance.**

***Index Terms*—CNN, Classification, Detection, TensorFlow, Keras, Skin Cancer**

I. INTRODUCTION

Our skin serves as a vital shield against environmental hazards, including chemicals and physical harm. Melanin, a pigment within the skin, plays a crucial role in absorbing ultraviolet (UV) radiation, thus mitigating the harmful effects of sun exposure. The skin comprises three layers: the epidermis (outermost), dermis (middle), and hypodermis (innermost). The epidermis acts as the first line of defense against the environment. The dermis, located beneath, contains connective tissues and sweat glands. The hypodermis, the deepest layer, is primarily composed of fat.

Skin cancer arises from the uncontrolled growth of skin cells, potentially spreading to other parts of the body. It is classified into three main types: basal cell carcinoma, squamous cell carcinoma, and melanoma. Melanoma, while less common, is the most aggressive form of skin cancer. Globally, significant numbers of new melanoma cases and related deaths are reported annually. Exposure to UV radiation from the sun is a primary risk factor for skin cancer. Melanoma has the characteristic of spreading to surrounding tissue. However, early detection and treatment significantly improve survival rates.

Identifying melanoma in its early stages is both challenging and essential. While many skin cancers fall into the non-melanoma category (such as squamous cell carcinoma and basal cell carcinoma), which are generally less aggressive and easier to treat, early detection remains critical for all types of skin cancer. Traditionally, a biopsy, involving the removal of a tissue sample for laboratory analysis, has been the standard diagnostic method. This procedure, however, can be uncomfortable, time-consuming, and costly. Furthermore, diagnostic accuracy is subjective and relies heavily on the expertise of the clinician. Even the most experienced dermatologists have a limited accuracy rate. The shortage of qualified dermatologists, especially in public healthcare systems, further complicates this issue.

This paper proposes a deep learning-based system designed to assist dermatologists. The system is trained to identify skin cancer by analyzing medical images, and it can be developed without requiring any programming expertise. This model achieves a high average diagnostic accuracy, demonstrating the potential of machine-assisted approaches to overcome challenges related to diagnostic delays, accuracy, and the shortage of dermatologists.

Numerous studies have investigated skin cancer detection and image classification using various techniques. These techniques include decision trees, Bayesian classifiers, support vector machines, and artificial intelligence-based methods, all aiming to improve the accuracy and efficiency of skin cancer diagnosis. The main objectives of this paper are:

1. To enable easy diagnosis of skin cancer even in inaccessible areas

2. To enhance the accuracy in the model's detection and classification mechanism.

II. RELATED WORKS

The field of skin cancer detection using image analysis has witnessed remarkable progress over the years, with researchers exploring a multitude of techniques [1]. The International Skin Imaging Collaboration (ISIC) event of 2018 stands as a pivotal benchmark, fostering advancements through its competitive challenge in skin cancer detection. Notably, the feasibility of mobile applications for skin cancer detection has also been demonstrated, highlighting the potential for accessible and widespread diagnostic tools. In all these endeavors, the primary focus has been on enhancing diagnostic accuracy through the application of diverse classification algorithms and methodologies.

The landscape of image classification underwent a significant transformation with the introduction of Convolutional Neural Networks (CNNs). Pioneered by Fukushima (1988) and further developed by LeCun (1990), CNNs revolutionized image analysis by emulating the human visual cognition system. These networks have since become the state-of-the-art approach for image classification tasks, showcasing unparalleled performance. While a vast body of literature exists on image classification, this review specifically concentrates on deep learning methods applied to skin cancer images.

The first significant breakthrough in skin cancer classification using a pre-trained GoogLeNet Inception V3 CNN model was achieved by Esteva et al. [20]. Their study utilized an extensive dataset of 129,450 clinical skin cancer images, including 3,374 dermatoscopic images, achieving a reported classification accuracy of 72.1±0.9\%. This work demonstrated the potential of deep learning to process large-scale clinical datasets for skin cancer detection.

Building on this foundation, Yu et al. [21] in 2016 developed a deep CNN architecture with over 50 layers, specifically trained on the ISBI 2016 challenge dataset for malignant melanoma classification. Their model achieved a best classification accuracy of 85.5\%, highlighting the benefits of deeper network architectures and specialized training datasets.

In 2018, Haenssle et al. [22] further advanced the field by utilizing a deep convolutional neural network to classify dermatoscopy melanocytic images into binary diagnostic categories. Their model achieved an impressive 86.6\% sensitivity and specificity, underscoring the potential of deep learning for accurate binary classification in dermatological imaging.

Han et al. in Ref. [24] employed a deep convolutional neural network to classify clinical images of 12 distinct skin diseases. Their study reported a best classification accuracy ranging around 96.0\% ± 1\%, showcasing the potential of deep learning for multi-disease classification in dermatology.

While a comprehensive review of all classifiers is beyond the scope of this paper, a systematic review of deep learning classifiers can be found in Ref. [25]. These studies collectively illustrate the significant advancements in skin cancer detection achieved through deep learning, highlighting the potential of these techniques for clinical applications.

III. APPROACH

Artificial Intelligence (AI) is transforming the healthcare industry, enhancing the capabilities of medical professionals and improving patient outcomes. AI algorithms, particularly those based on machine learning and deep learning, are used to analyze medical images and accurately diagnose diseases, such as cancer, cardiovascular conditions, and neurological disorders. In drug discovery and development, AI accelerates the process by predicting how compounds will interact with targets in the body, identifying potential drug candidates more efficiently. Personalized treatment plans are created by analyzing a patient's genetic information, medical history, and lifestyle factors, ensuring the most effective treatments tailored to individual needs. AI-powered robotic systems assist surgeons in performing complex procedures with greater precision, reducing the risk of complications.

A. CNN in Skin Cancer Detection

Convolutional Neural Networks (CNNs) are a specialized form of neural network particularly well-suited for image processing. Their advantage over traditional fully connected networks stems from their efficient handling of image data, leveraging sparse connections and weight sharing among pixels.

The architecture of a CNN's hidden layers typically includes convolutional layers, non-linear pooling layers, and fully connected layers as seen in Fig. 1. Multiple convolutional layers, followed by fully connected layers, form the core structure. The three key components of a CNN are: convolutional layers, pooling layers, and fully connected layers. Convolutional layers learn sets of weights, which are then simplified by pooling layers. This pooling process reduces the dimensionality of the output and effectively shrinking the input size. The output from the pooling layers is then fed into the fully connected layers for final classification or regression. A crucial part of the CNN is the convolutional layer, which employs diverse weight sets tailored to specific tasks, such as image segmentation, and utilizes multiple two-dimensional matrices.

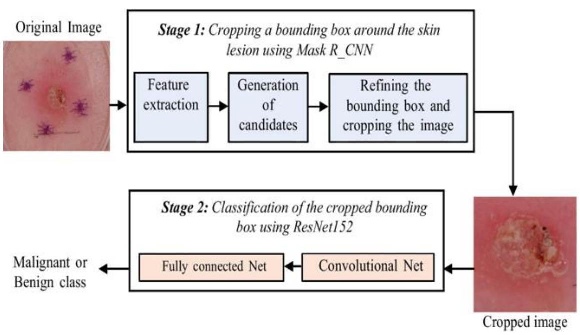

Fig. 1: Working of CNN Model

A multitude of studies have explored the application of Convolutional Neural Networks (CNNs) for skin cancer detection and classification. Helker et al. (2020) investigated the impact of integrating patient data, such as age, sex, and lesion location, with image features in CNN-based skin lesion classification. Their findings indicated that incorporating patient information can enhance the performance of CNN classifiers.

Hasan et al. (2019) employed CNNs to classify skin cancer images as malignant or benign, utilizing feature extraction techniques to analyze affected skin cells. Their approach yielded promising results, suggesting its potential as a benchmark for skin cancer detection. Brinker et al. (2018) conducted a comprehensive review of CNN-based skin cancer classification research, highlighting the effectiveness of CNNs while also acknowledging the challenges in comparing results due to the use of non-public datasets in some studies.

Gong and Xiao (2021) combined CNNs with Natural Language Processing (NLP) to create a user-friendly application for skin lesion classification and information provision through a chatbot interface. This application demonstrated high accuracy in image identification. Seung Seog Han et al. (2019) developed an algorithm using a type of CNN called R-CNN to detect potential skin cancer locations and predict malignancy, achieving high accuracy in diagnosing melanoma, particularly on the face.

These studies collectively illustrate the significant potential of CNNs for accurate and efficient skin cancer detection and classification, showcasing various approaches to enhance performance and contribute to the development of more reliable and accessible diagnostic tools. Fig. 2 is an average estimate of how CNNs generally work and its flow.

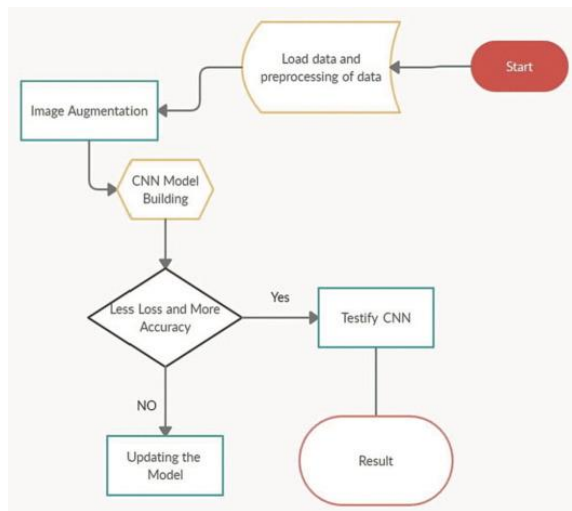

Fig. 2: System Flow Diagram

IV. METHODOLOGY

This section provides a comprehensive overview of the methodological framework employed in the design and implementation of our skin cancer detection system. It encompasses the key stages involved in realizing our objective of creating an accurate and accessible diagnostic tool, from the initial curation of a representative dataset to the final deployment on a mobile platform.

A. Data Collection

In this study, we leveraged a meticulously curated custom dataset comprising a diverse collection of skin lesion images. Each image was carefully categorized into one of two classes: "benign" or "malignant," representing the nature of the depicted lesion. This binary classification scheme, while simplifying the diagnostic task, ensured a clear and interpretable output for the end-user. The labeling process was conducted with utmost precision, often involving multiple expert dermatologists to ensure accuracy and minimize potential biases.

B. Data Preprocessing

MobileNet's streamlined architecture and reduced computational demands make it particularly well-suited for mobile deployment. Hence, we opted for the MobileNet model as the backbone of our skin cancer detection system.

To ensure our chosen network receives data in an optimal format for effective training and accurate classification of skin lesions as benign or malignant, we implemented a series of preprocessing steps:

1. Employing Data Generators: Data generators play a crucial role in streamlining the data feeding process for our neural network. These generators efficiently handle the reading of images from our designated source folders, converting them into float32 tensors compatible with TensorFlow and Keras, and pairing them with their corresponding labels ("benign" or "malignant").

2. Normalization: Normalization is a vital preprocessing step that scales the pixel values of our images to a standardized range. In our implementation, we normalized the images by scaling their pixel values from the original range of [0, 255] to a normalized range of [0, 1].

3. Resizing: To maintain consistency and compatibility with the input requirements of our chosen MobileNet architecture, we resized all images to a uniform dimension of 224x224 pixels. This resizing step ensures that all input images conform to the expected input shape of the network.

C. Build the Model

Constructing our skin cancer detection model involved leveraging the power of transfer learning. With the help of the Hub module, we seamlessly layered a linear classifier on top of the pre-trained MobileNet feature extractor. This approach allowed us to capitalize on the extensive knowledge already embedded within the MobileNet architecture, fine-tuned on a massive dataset of diverse images.

Initially, we configured the feature extractor to be non-trainable. This strategic decision expedited the training process, as the pre-trained weights remained fixed. However, to potentially enhance the model's accuracy and tailor it more precisely to our specific task of skin cancer detection, we enabled fine-tuning. Fine-tuning involves allowing the feature extractor's weights to be adjusted during training, as shown in Fig. 3, enabling the model to adapt more closely to our custom dataset.

```
Building model with https://tfhub.dev/google/tf2-preview/mobilenet_v2/feature_vector/2
Model: "sequential_4"
_________________________________________________________________
Layer (type)                 Output Shape              Param #
=================================================================
keras_layer_10 (KerasLayer)  (None, 1280)              2257984
_________________________________________________________________
flatten_11 (Flatten)         (None, 1280)              0
_________________________________________________________________
dense_22 (Dense)             (None, 512)               655872
_________________________________________________________________
dropout_11 (Dropout)         (None, 512)               0
_________________________________________________________________
dense_23 (Dense)             (None, 2)                 1026
=================================================================
Total params: 2,914,882
Trainable params: 656,898
Non-trainable params: 2,257,984
_________________________________________________________________
```

Fig. 3: MobileNet Summary

D. Train the Model

During the training phase, it is crucial to validate each step by employing a validation dataset. Initially, our model achieved an accuracy of 81\% after 15 epochs. This demonstrates a strong foundational performance, but with additional fine-tuning, the accuracy can be further improved. By adjusting hyperparameters, increasing the complexity of the model, or utilizing advanced techniques such as data augmentation and regularization, we can enhance the model's predictive capabilities.

E. Testing and Validation

We picked 5 images at random and tested them on the validation dataset and observed the following output as shown in Fig. 4. To monitor the performance of a machine learning model, it is crucial to plot the training and validation accuracy and loss metrics over the course of the training process. These plots in Fig. 5 provide visual insights into how well the model is learning and generalizing to unseen data.

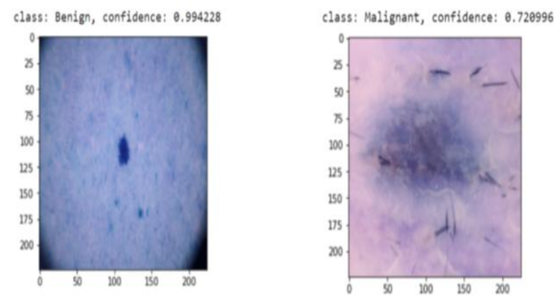

Fig. 4: Test Results

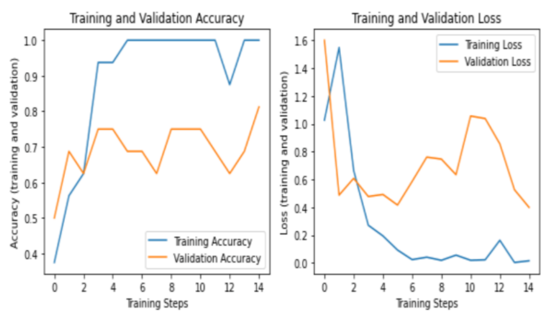

Fig. 5: Graph Results

The following confusion matrix was obtained as shown in Fig. 6 after validation and testing using the required dataset:

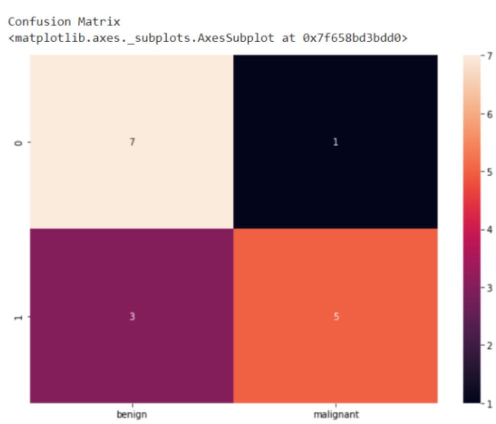

Fig. 6: Confusion Matrix

F. Integrate into Android Application

The TensorFlow Lite Model is then uploaded to Android Studio and is integrated into the code. To integrate the model into the Android project, start by placing the \verb|model.tflite| and \verb|labels.txt| files into the \verb|assets/| directory. Then, add the TensorFlow Lite dependency to your \verb|app/build.gradle| file with the implementation line: implementation org.tensorflow:tensorflow-lite:2.8.0'|. Next, load the model and labels in the code by initializing the \verb|Interpreter| and reading the labels into a list. This setup allows the Android application to leverage TensorFlow Lite for running machine learning models, enabling advanced functionality and improving user experience.

G. Output

Users can upload an image of a lesion or mole, which undergoes segmentation, resizing, and color normalization. The image can be taken in real-time or selected from the user's gallery. The output is then displayed as either "Malignant" or "Benign" (e.g., Fig. 8 and Fig. 9) along with a confidence rate in decimals (e.g., [0.98758]). This helps users evaluate the model's dependability and diagnostic accuracy.

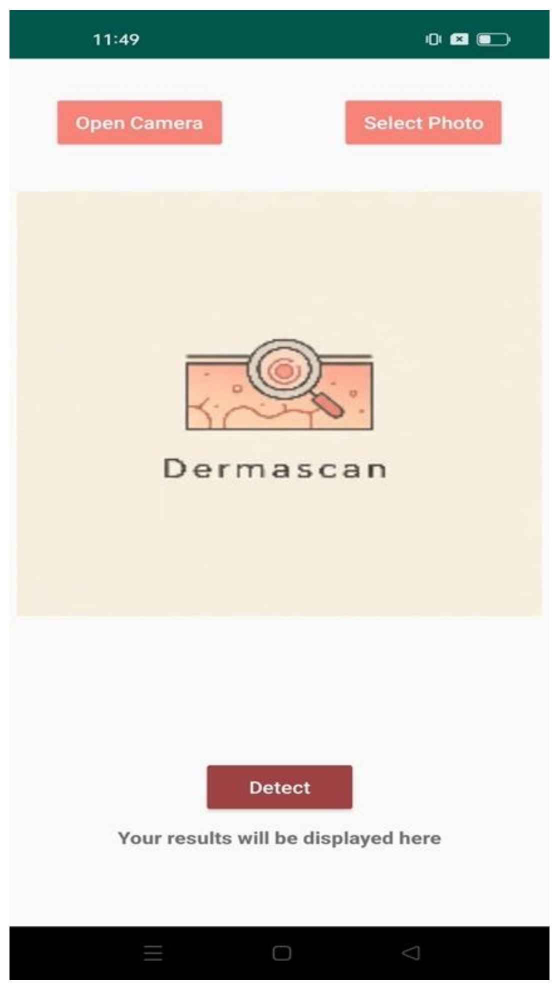

Fig. 7: Home Page

Fig. 8: Result Page- Malignant

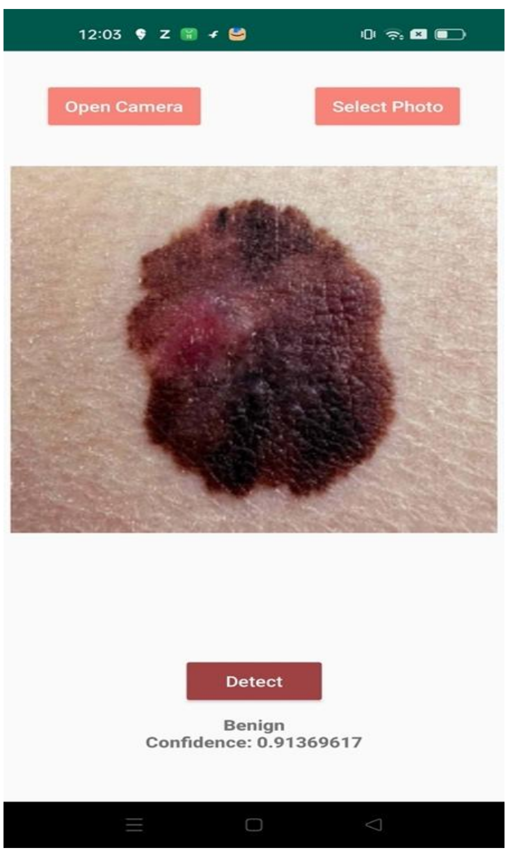

Fig. 9: Result Page- Benign

V. IMPLEMENTATION

Our first step involved the meticulous curation of a custom dataset comprising diverse skin lesion images, carefully labeled as either benign or malignant. We sought to gather a diverse array of skin lesion images, capturing the subtle variations in color, texture, and shape that distinguish benign from malignant tumors. Each image underwent careful review and labeling by experienced dermatologists, ensuring the accuracy and reliability of our dataset.

With our dataset in place, the next crucial decision was the selection of a deployment framework that could translate our trained CNN model into a practical mobile application. TensorFlow Lite emerged as the ideal candidate, offering a streamlined approach to running machine learning models on resource-constrained devices like smartphones. Its lightweight architecture and optimized performance enabled us to achieve real-time inference without sacrificing accuracy, a critical factor for a user-friendly diagnostic tool. The ability to perform computations locally on the device also addressed concerns regarding data privacy and network connectivity.

The development of the Android application itself was a labor of thoughtful design and meticulous coding. We chose Kotlin and Java, languages renowned for their flexibility and stability in Android development, to construct an interface that would be both intuitive and accessible. We understood that the application's success hinged on its ease of use, particularly for individuals who may not be technically proficient. Thus, we focused on creating a seamless user experience, guiding users through the process of capturing or uploading skin lesion images with clear instructions and visual cues.

The application's integration with the device's camera was carefully implemented to ensure optimal image capture. We incorporated features such as automatic focus and adjustable lighting to minimize variations in image quality. For users who preferred to utilize existing images, we provided seamless access to the device's gallery, allowing them to select images with ease. Once an image was acquired, the magic of our TensorFlow Lite-powered CNN began. In a matter of seconds, the model analyzed the image, extracting relevant features and comparing them to the patterns learned during training. The result was a clear and concise classification: benign or malignant.

VI. CONCLUSION

Skin cancer, a significant global health concern, necessitates early detection for effective treatment and improved patient outcomes. This project aimed to develop an accessible and user-friendly tool for preliminary skin lesion assessment, leveraging the power of deep learning and mobile technology. By harnessing a Convolutional Neural Network (CNN) trained on a custom dataset and deployed using TensorFlow Lite within an Android application built with Kotlin and Java, we successfully created a system capable of classifying skin lesions as either benign or malignant.

The CNN, trained on a carefully curated dataset of skin lesion images, demonstrated the ability to extract intricate features and patterns indicative of malignancy. TensorFlow Lite facilitated efficient deployment on mobile devices, enabling real-time classification without the need for server-side processing. The Android application, with its intuitive interface, allows users to capture or upload images of skin lesions and receive immediate feedback on the likelihood of malignancy.

This project contributes to the growing field of mobile health by providing a readily accessible tool for skin cancer risk assessment. While not intended to replace professional medical diagnosis, it empowers individuals to take proactive steps in monitoring their skin health and seeking timely medical attention when necessary. The integration of deep learning and mobile technology opens up exciting possibilities for democratizing healthcare access and improving early detection rates for skin cancer.

REFERENCES

[1] Divya D. J., Prakasha S,"\textit{Clustering Techniques for Medical Imaging",}International Journal of Innovative Technology and Exploring Engineering (IJITEE) (ISSN: 2278-3075) vol. 9, December 2019

[2] Fumio Nogata, Yasunai Yokota, Yoko Kawamura, Tetsuya Mouri, William R. Walsh, Takahiko Kawamura, Nigishi Hotta, "\textit{Towards the application of AI Technology to Assess Risk of Aneurysm Rupture Based On Medical Imaging", }International Journal of Computer and Information Technology (ISSN:2279-0764) vol. 8, July 2019.

[3] Janis Born, David Beymer, Deepta Rajan, Adam Coy, Vandana V. Mukherjee, Matteo Manica, Prashanth Prasanna, Deddeh Ballah, Michal Guindy, Dorith Shaham, Pallav L Shah, Emmanouil Karteris, Jan L. Robertus, Maria Gabrani and Michal Rosen-Zvi, "\textit{On the role of artificial intelligence in medical imaging of Covid-19", }Patterns 2, June 11 2021.

[4] V. A Ashwath, O.K. Sikha and Raul Benitez, "\textit{TS-CNN: A Three Tier Self-Interpretable CNN for Multi-Region Medical Image Classification"}, IEEE Engineering in Medicine and Biology Society Section, vol. 11, July 2023.

[5] Zhang Li, Jiehua Zhang, Tao Tan, Xichao Teng, Xiaoliang Sun, Hong Zhao, Lihong Liu, Yang Xiao, Byungjae Lee, Yilong Li, Qianni Zhang, Shujiao Sun, Yushan Zheng, Junyu Yan, Ni Li, Yiyu Hong, Junsu Ko, Hyun Jung, Yanling Liu, Yu-cheng Chen, Ching-wei Wang, Vladimir Yurovskiy, Pavel Maevskikh, Vahid Khanagha, Yi Jiang , Li Yu, Zhihong Liu, Daiqiang Li, Peter J. Schüffler, Qifeng Yu, Hui Chen, Yuling Tang, and Geert Litjens, "\textit{Deep Learning Methods for Lung Cancer Segmentation in Whole-Slide Histopathology Images",}IEEE Journal of Biomedical and Health informatics, vol.25, February 2021.

[6] Chaitanya Kolluru, Juhwan Lee, Yazan Gharaibeh,Hiram G. Bezerra, David L. Wilson,"Learning With Fewer Images via Image Clustering: Application to Intravascular OCT Image Segmentation", IEEE Access, VOLUME 9, February 11, 2021.

[7] Siyuan Tang , Rong MA, Qingqian LI, Yingchun Bai, and Shijun Chen, "Classification of Benign and Malignant Pulmonary Nodules Based on the Multi Resolution 3D DPSECN Model and Semi Supervised Clustering", IEEE Access, VOLUME 9, February 18, 2021.

[8] Xiangxia LI , Bin LI, Fang Liu, Hua Yin,and Feng Zhou, "Segmentation of Pulmonary Nodules Using a GMM Fuzzy C-Means Algorithm", IEEE Access, VOLUME 8, February 3, 2020.

[9] Ying Chen, Yerong Wang , Fei HU, and Ding Wang, "A Lung Dense Deep Convolution Neural Network for Robust Lung Parenchyma Segmentation", IEEE Access, VOLUME 8,May 11, 2020.

[10] Dan Wang , Junfeng Wang, Yihua D, and Peng Tang, "Adaptive Solitary Pulmonary Nodule Segmentation for Digital Radiography Images Based on Random Walks and Sequential Filter", IEEE Access, VOLUME 5, February 14, 2017.

[11] Baihua Zhang, Shouliang QI, Patrice Monkam, Chen LI, Fan Yang , YU-Dong Yao, and Wei Qian,"Ensemble Learners of Multiple Deep CNNs for Pulmonary Nodules Classification Using CT Images", IEEE Access, VOLUME 7, August 7, 2019.

[12] Jan L. Bruse, Maria A. Zuluaga, Abbas Khushnood, Kristin McLeod, Hopewell N. Ntsinjana, Tain-Yen Hsia, Maxime Sermesant, Xavier Pennec, Andrew M. Taylor, and Silvia Schievano, "Detecting Clinically Meaningful Shape Clusters in Medical Image Data: Metrics Analysis for Hierarchical Clustering Applied to Healthy and Pathological Aortic Arches", IEEE Transaction On Biomedical Engineering, VOL. 64, NO. 10, OCTOBER 2017.

[13] IN Wang, Shuaizong SI , Enuo CUI, Hai Zhao, Dongaiang Yang , Shengchang Dou, and Jian Zhu, "A Fast and Efficient CAD System for Improving the Performance of Malignancy Level Classification on Lung Nodules", IEEE Access, VOLUME 8, February 27, 2020.

[14] Tao Han, Virginia Xavier Nunes,Luis Fabricio De Freitas Souza Adriell Gomes Marques, Iagson Carlos Lima Silva, Marcos Araujo Ferreira Junior, Jinghua Sun, and Pedro P. Reboucas Filho," Internet of Medical Things—Based on Deep Learning Techniques for Segmentation of Lung and Stroke Regions in CT Scans", IEEE Access, VOLUME 8, April 14, 2020.

[15] Debanjan Konar, Bijaya K. Panigrahi, Siddhartha Bhattacharyya, Nilanjan Dey, and Richard Jiang, "Auto-Diagnosis of COVID-19 Using Lung CT Images With Semi-Supervised Shallow Learning Network", VOLUME 9, February 11, 2021.