Skin Cancer Detection Using Deep Learning

Varshini[1], Sudarshan K[2], Swathi[3], Harshini P K[4], Sanvi S Shetty[5]
[1,3,4,5]*Third year B.E Student, Department of Information Science and Engineering, Srinivas Institute of Technology, Mangaluru, India.*
[2]*Associate Professor, Department of Information Science and Engineering, Srinivas Institute of Technology, Mangaluru, India.*
doi.org/10.64643/IJIRTV12I12-206750-459

***Abstract*—With an always rising number of cases and a fast progression, skin cancer, especially melanoma, has transformed into a very critical public health issue. To detect it as soon as possible, this automated system based on a Deep Learning model: MobileNetV2 is created. The small size and high accuracy of this model MobileNetV2 was the key reason for considering it. During the training of the model, the dataset used was over 10,000 images. To increase the reliability, all the images went through several preprocessing steps, which includes resizing and a method called "data augmentation" for training the model to recognize melanoma in different situations. The system has achieved an accuracy of about 90% for detecting melanoma. In order to create a user interface that is easy to use without reading a manual first, Gradio was chosen and implemented an upload button that lets the users to upload the images directly from their hard drive. After uploading a photo, the user immediately receives a result and a confidence score. The aim is to support the doctors in filtering out potential melanomas to avoid them from getting lost in a backlog and to ensure that patients with severe skin cancer do not have to wait for too long.**

***Index Terms*—Deep Learning, Gradio, Melanoma Detection, MobileNet**

I. INTRODUCTION

Melania is indeed a concern, primarily due to its rapidity in spreading. Early detection significantly increases survival rate. The difficulty lies within identifying dangerous spots at on set since they appear virtually indistinguishable from benign moles. As such, when numerous lesions exist on the body which maybe cancerous or noncancerous (and virtually indiscernible) even dermatologists have to perform extensive and time-consuming visual inspections, which inevitably contain potential for error.

As well, for individuals residing in rural settings, access to specialists can be limited. Thus, when treatment is required, many patients are forced to travel long distances to receive care - and unfortunately delay treatment. A remote medical consultation could resolve this issue. It would provide a means for receiving a professional's opinion quickly regarding a possible serious condition via remote examinations and imaging.

I have created a melanoma-detecting system utilizing mobile learning technology to assist with detecting melanoma. Specifically, I used the MobileNet V2 model, as it is very lightweight and fast. Unlike a large, high-end computer, my model can run on any basic phone or laptop.This will be a huge help for many rural clinics where specialist visits may be rare. In order to test the effectiveness of this method, I used a dataset of over 10,000 images of skin. In order to make this as accurate as possible, I used a "data augmentation" method, basically showing the AI the same photograph but rotated or with different amounts of brightness to make sure it actually learns what a melanoma looks like and isn't just saying all of the pictures in its database are melanomas. It currently gets a 90% accuracy in distinguishing melanoma from normal molems, but in order to make this easy to access and use by many users, I built a Gradio tool so anyone can upload a photograph and receive a result with a confidence score. In conclusion, this would give doctors a much better shot at getting patients help, and would make the whole screening process much quicker.

As it stands right now, spotting melanoma is still a very manual and subjective process. Dermaticians usually look at a picture of the skin, have a tool called a dermo scope, and they have to go through the standard ABCD rules. But eventually, accuracy comes down to how much experienced you are and that particular doctor. And if you're in a crowded Indian hospital, single doctor might see hundreds of patients in one shift, fatigue is a very real factor. It's only human to have a small margin of error when you're that busy

The scenario is even tougher in rural areas where specialists are few and far apart. People end up stuck in long waiting lines, and the delays can turn a treatable case into a dangerous one. Because early-stage melanoma looks mostly similar to a harmless mole, it's incredibly easy to miss out during a quick and manual check. Essentially, our current system is just too slow and places a massive burden on the shoulders of individual experts. That is why we need a faster and more automated approach

The system provides an AI-powered solution that uses the mobilenetv2 architecture to detect melanoma automatically. Instead of just depending on a manual eye-test, this system examines skin images and classifies them with high accuracy. We chose MobileNetV2 because it's not only efficient but works well even on basic devices.

Which make it great for small clinics in remote or rural areas. The model that is trained on a large set of over 10,000 images and is designed to handle different types of skin tones and lighting conditions properly.

With an accuracy of about90%, it's a very reliable tool for diagnostic support. Also, we are using Gradio to create a simple interface. So even someone who's not a tech-expert can also upload a photo and get an. Instant result. It helps doctors quickly identify high-risk patients. It's also reducing delays and making the whole process of clinical decision-making much smoother. This makes the whole process of diagnostic more accurate.

II. PROCEDURE

A. Review Stage

This work focused on using MobileNetV2 for melancoma classification is submitted electronically in a standard two-column format. All technical drawings like our CNN architecture diagram or snapshots of the Gradio interface are placed within these columns along Performance tables. For complex data or high-resolution skin images that need more detail in the figures, we scaled the figures to spans both columns so that the reviewers can see all the visual features clearly.

B. Final Stage

Once the research paper is accepted, we will submit the final "camera-ready" version. Along with this we will also provide the signed copyright transfer form required for the final publication phase. We've made sure that the final document accurately includes everything from the mobilenet V2 methodology and results to all the technical documentation required for the journal's records.

C. Figures

The figures in this paper show exactly how the system is built and how it works for diagnosis. It includes a detailed flowchart to explain the whole process, starting from taking the image and processing it and leading up to the final deep learning classification results. We've also added snapshots of the Gradoweb interface to show how the model actually works in real-time. These visuals along with some sample results from our melanoma model. Keras can prove the 90% accuracy and clearly show how the system tells the difference between different types of skin lesions.

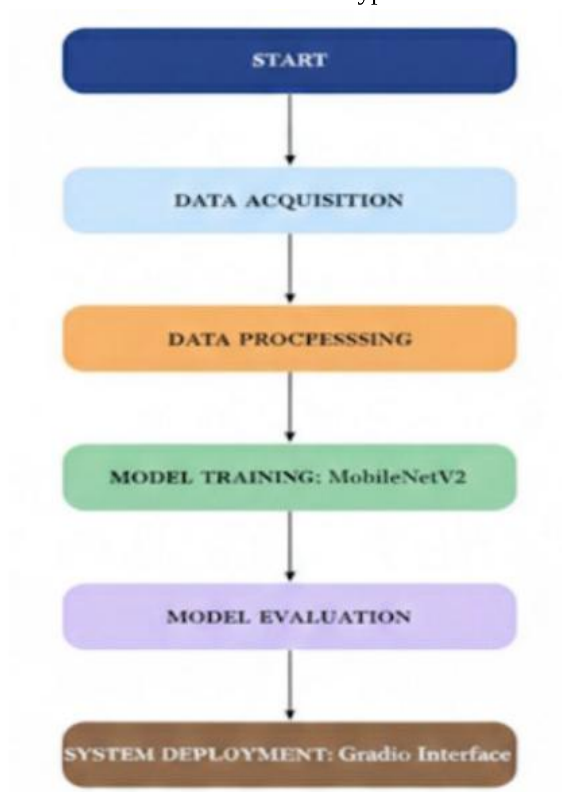

Figure.1: System Flowchart

III. LITERATURE SURVEY

A lot of progress has been made lately in using deep and lightweight CNNs to spot skin cancer. Here is a quick summary of the most relevant research papers we looked at:

Chatervedi et alo (2019) build a classification system based on the HAM10000 dataset and a MobileNet-based model. We can see that using lightweight models is useful here, because the results show that these types of models can do better at classifying skin lesions accurately without the heavy computing power that deeper models require

Innani et alo 20 (2023) present a 2-step process that first segments the Skin lesion and then classifies it. We see that segmenting the lesion first and using only the relevant parts of the skin image can go a long way in making the classification more accurate, because it removes all the distractions caused by the background.

A research paper in BMC Medical Imaging (2020) uses a very deep model, ResNet 15, on the ISIC 2017 dataset. We find that even though their results are very strong, their results also show that deep models need a lot of computing power. This gives us a clearer understanding of why it is good to create more Optimized models which can do well with less computing power if we want to use them in real-world settings.

Then an article was found a 2023 article in mdpi Diagnosties which gives us a comparison of different models, VGG, ResNet, and MobileNet. They find out that while Deep models are accurate MobileNetV2 offers the best balance between speed and performance. They recommended that future work should focus on better Data augmentation and to get systems ready for real-time use.

We have found an interesting recent paper published on arrived (2024) about Transfer learning using VgG16 and VGG19. It is found that with the correct kind of Data Augmentation and Training, the accuracy can goas high as 8.18% showcasing how powerful transfer learning could be for this task.

The review which the publishing house Springer Nature released in 2025 examined the newest directions like transformer models and "Explainable A. I." It has been observed that these systems must have greater transparency and reliability before medical professionals can place their trust in the A. I. for clinical use. Medical professionals require high confidence in the decisions made by the A. I. because clear functions allow for better understanding in the clinical environment.

Research from the Pima hybrid model study (2023-) explored a different path when traditional texture extraction (GLCM) was joined with deep learning by the investigators. Many people believe that this hybrid approach functioned with high success, particularly for visual records with low clarity. This strategy suggests that the blending of traditional features with A. I. makes the detection process stronger because the A. I. utilizes multiple sources of data for the analysis.

Various investigations which were printed in MDPI Sensors (2022-24) gave attention to ensuring melanoma detection operates on IoT and small portable tools. Experts claim that efficient structures such as MobileNet or ShuffleNet are required for the melanoma detection to happen in real-time. The decision to use MobileNetV2 for this specific research endeavor is supported by these results because MobileNetV2 matches the needs of the hardware perfectly.

Akinrinade and Du (2025) used deep learning and few-shot learning methodologies to improve the accuracy of skin cancer detection in limited resource settings. The study employed transfer learning and GANs to address imbalances in the dataset to obtain high accuracy skin cancer classification. A deep neural network was used to examine texture features of lesions for developing hybrid models for the early detection of skin cancer. The findings suggest the promise of such methods in incorporating into digital health platforms for timely screening and treatment of skin cancer.

In a paper published in advance of the 2025 LangFest conference, S. R. S, J. R, and R. K. developed, along with the team at Vikram Hospital, a system called ResNetVision, based on a Convolutional Neural Network (CNN) for combating skin diseases, which are rapidly increasing. The method of manually diagnosing diseases is often unreliable and, thus, has been replaced by image preprocessing and feature extraction to automate the process. The large number of patient images in this study were analyzed to identify potential lesions rapidly in order to save lives.

M. Kumar, R. Kumar, A. Maheshwari, R. Gupta, S. Sharma and N. Kumar (2025) presented a machine learning based approach in conjunction with deep neural networks for the classification of seven different skin diseases. An enormous dataset of 10200 images was used and optimal features were extracted and the same were classified with accuracy 98.56%. The whole process of computerized image analysis was made accessible on web and the output obtained in few seconds helps in early diagnosis which increases chances of survival.

Afifi et al. (2024) design and implement a mobile tool called “SkinKo” for skin diagnosis using only a smartphone camera, utilizing the EfficientNetB0V2 architecture and the powerful Flutter framework. This high-efficiency substitute for in-clinic skin assessments are now deployable in the pockets of general practitioners and patients alike through training on the HAM10000 dataset, encouraging ongoing monitoring by the patient and facilitating earlier referral to skin specialists.
In order to automate the analysis of skin lesions, Khamsa et al. (2023) in their research paper introduced a multi-task model that merges Yolo which is used for the localization of lesions and VGG16 for diagnostic categorization of skin lesions into different types. They employed transfer learning for training and validation purposes on a dataset which included 3,297 dermo scopic images. The model achieved an accuracy of 83% and also managed to ensure 83% success rate in skin cancer screening, reducing human involvement and making the process less subjective.

Suneetha et al. (2024) proposed a hybrid approach of deploying pre-trained ResNet101 V2 architecture and evolutionary search algorithms in tuning features for skin cancer classification with 95.5% of accuracy in detecting melanoma. The above model automates the diagnostic process leading to timely clinical interventions thus preventing the spread of malignant cells. The mixing and matching of pre-trained deep learning architectures with optimisation techniques specific to them significantly improves the accuracy and reliability of skin cancer classification system.
Augustine et al. (2025) compared and fine-tuned various architectures such as DenseNet, InceptionV3, and MobileNet to improve transfer learning techniques for skin cancer screening. An evolutionary perspective was used to select the best models for classification. The results show that a carefully designed deep learning model can drastically improve the early detection and clinical decision making

IV. METHODOLOGY

The procedure for building this melanoma detection system follows a clear step by step pipeline, starting from collecting data to finally deploying the model.

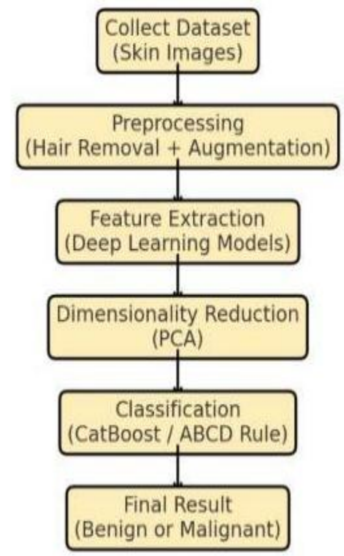

Figure.2: Methodology Diagram of the Proposed System

A. Data Acquisition and Dataset Description
A large Dataset of over 10k skin images was used for this project. These pictures are dividing into two main groups, Melanoma(malin) and non-melanoma (benigne). The dataset is quite unique covering many different lesion shapes, colors, and textures exactly what we need to train a trustworthy model.

B. Image Preprocessing
Experts claim that several significant procedures are completed before the neural network receives the images to guarantee that the information remains tidy and uniform.

- Resizing: The resizing of all raw images to a standard 224×224 pixels is performed by the system.
- Normalization (scaling down): Scaling pixel values to a range between 0 and 1 through division by 255 is how normalization occurs because the model learns with greater speed during the training phase.
- Data Augmentation: Many researchers believe that data augmentation techniques must be used to stop the model from experiencing "over fitting"

regarding the training data. Random rotating, flipping, zooming, and changing the brightness of the images are included in these techniques. Such variety in methods assists the AI when the AI identifies various skin conditions in the physical world.

C. Model Architecture: MobileNet V2

The foundation for the intellectual core of the system relies upon the MobileNet V2 architecture. Many experts believe that MobileNet V2 uses "inverted residual blocks," which allows the model to stay lightweight and rapid. While the model remains lightweight, MobileNet V2 functions without losing much precision, which is a reason the model was selected. Achievement for the system depends on this speed so that the system operates in real-time. A Global average Pooling layer and a Dense layer featuring a sigmoid as the activation function are added at the conclusion of the network. These components provide a probability score (p) for the final classification task. we've added a Global average Pooling layer and a Dense layer with a sigmoid as the activation function to give us a probability score (p) for the final classification.

D. Classification Logic and Risk Assessment

The system decides if a lesion is dangerous based on the probability (p) it calculates:

- The system decides if the lesion is dangerous based on the probability (p) it calculates.
- Non-Melanoma ($p<0.50$). Melanomen ($p>0.5$).
- If the system detects Melanoma, then it takes one step further by categorizing the risk into three levels based on confidence level:
- Stage 1 (Low Risk): $0.50 < p < 0.60$
- Stage 2 (Moderate Risk): $0.60 \leq p < 0.80$
- Stage 3 (High Risk): $p \geq 0.80$

E. System Implementation and Deployment

After the model was trained, we saving melanoma. Model. Keras and then used the Gradio library to build an interface for it.

This implementation really makes the system very user friendly. Anyone can just take a skin image and receive immediate report showing the classification, the risk stage, and some basic medical advice generated by the AI.

V. RESULT AND DISCUSSION

A. Performance Metrics

We examined the system based on how well it classified skin images and found that the model has an overall 90% accuracy. This means that it is effective in distinguishing between malignant melanomas and harmless lesions.

The model's exceptional performance on new data is largely because of the comprehensive data augmentation practices we applied during training, such as rotating and zooming photos. It's important to note that by building our model on 10,000 images, it was able to pick up the features very effectively.

B. Real-Time Inference and UI Feedback

Using the "Gradio" interface, we were able to test the system with great success in real-time. When we upload a photo, the system processes it and it gives a result in 3 to 5 seconds.

The output includes the following features:

1. Classification: The output clearly labels the image as "melanoma" or "non-Melanoma."
2. Confidence: This gives a percentage showing how confident the model is about its choice.
3. Risky Level: If melanoma is detected, it correctly categorizes it into Stage 1 (Low), Stage 2 (Moderate), or Stage 3 (High) according to its probability.

C. Error Analysis

Even though the system is highly reliable, there's still that 13% error margin to keep in mind. From what we can see, the AI mostly makes mistake when the lighting in a photo is uneven or if there's hair covering the skin lesion.

In order to solve this problem, we implemented specific preprocessing steps to normalize and clean up image data to increase accuracy and reliability. It is not a perfect solution for every single bad photo, but it makea huge difference.

If compared to basic models that don't utilize a preprocessing step like this one it had really improved the error rate.

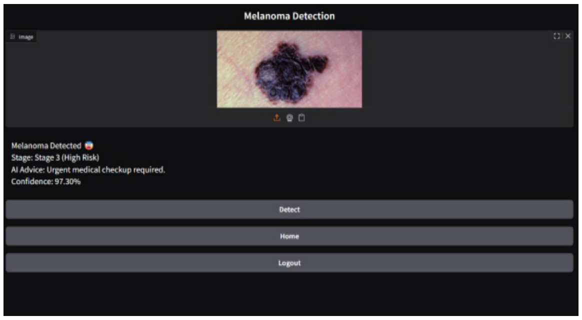

Figure.3: Prediction page as melanoma

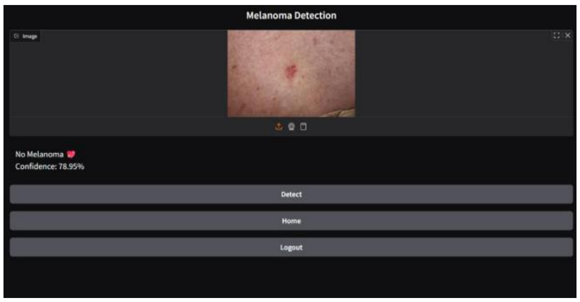

Figure.4: Prediction page as non-melanoma

D. Discussion

The results obtained really highlight those lightweight models specifically MobilenetV2, are a perfect fit into a medical support tool. They hit that area where you don't need a high-end supercomputer to get professional-grade results. This is really helpful for smaller clinics that might not have a big budget for expensive hardware. By providing a quick risk assessment and clear advice right away, the system does exactly what we set out to do: help catch skin cancer early and ensure doctors can focus their attention on the most urgent cases for a biopsy.

VI. CONCLUSION

In a few words, this project demonstrates that deep learning methods are more than capable of spotting melanoma with a high precision. It made use of the MobileNetV2 architecture and achieved a high accuracy rate of around 90% when distinguishing between melanomap and ordinary marks on the skin. However, the real success of this project is the Accessibility component. Thanks to the integrated Gradio web interface, you do Not have to be a techexpert to use it. Anyone can simply upload a photo and, within seconds, see a result along with a confidence score. A huge part of the performance came down. The preparation of data and the expansion of information stages were handled with great attention so that the skin image model could see the very small and difficult details in the pictures of human skin. Because the preparation of data was done carefully, the skin image model remains very strong when the skin image model encounters different types of skin damage.

Experts claim that the skin image model learns better because the expansion of information was used during the training time. Even though the physical inspection by a medical professional is considered by the medical world to be the highest level of truth, the skin image model functions as a very important assistant for the medical professional. This skin image model is very useful for the fast medical inspection process because the skin image model works without delay. It has been observed that the skin image model helps to identify people who have dangerous health problems. Since the skin image model identifies the skin damage early, the medical professional can focus on the individuals with low income or those who need help the most. Individuals with high-risk conditions receive the medical support those individuals require from the hospital because the skin image model provides the results quickly.

Although the skin image model is not a replacement for a medical professional, the skin image model helps the medical professional do the work better. This skin image model shows how technology can help people when the skin image model is built with care. By using these methods, the skin image model maintains a high level of accuracy across various lesions. Many people believe that the skin image model provides a necessary service in areas where a medical professional is not always available. Because the skin image model was tested extensively, the skin image model does not make many mistakes. When the skin image model is used correctly, the skin image model saves many lives by finding problems early.

REFERENCES

[1] I. R. Ashraf *et al.*, "Region-of-Interest Based Transfer Learning Assisted Framework for Skin Cancer Detection," *IEEE Access*, vol. 8, pp. 147858–147871, 2020.

[2] S. P. Maniraj and P. Maran, "A Hybrid Deep Learning Approach for Skin Cancer Diagnosis

Using Subband Fusion of 3D Wavelets," *Journal of Supercomputing*, vol. 78, no. 10, pp. 12394–12409, 2022.

[3] N. C. F. Codella *et al*., "Skin Lesion Analysis Toward Melanoma Detection: A Challenge at the 2017 International Symposium on Biomedical Imaging," *IEEE Transactions on Medical Imaging*, vol. 38, no. 8, pp. 1884–1896, 2019.

[4] I. Pacal *et al*., "A Novel CNN-ViT–Based Deep Learning Model for Early Skin Cancer Diagnosis," *Biomedical Signal Processing and Control*, 2025.

[5] S. E. Farea and M. Al-Antari, "A Hybrid Deep Learning Skin Cancer Prediction Framework," *Engineering Science and Technology, an International Journal*, vol. 57, p. 101818, 2024.

[6] M. A. Khan, A. Sharif, T. Saba, and M. Raza, "Skin Lesion Segmentation and Classification: A Unified Deep Learning Framework," *IEEE Access*, vol. 8, pp. 129395–129406, 2020.

[7] R. P. Amaral *et al*., "Automated Melanoma Detection Using Deep Learning: A Cross-Sectional Study of Dermatology Images," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 4, pp. 1772–1783, 2023.

[8] T. Mendonça *et al*., "Deep Learning for Melanoma Detection in Dermoscopy Images," *IEEE Latin America Transactions*.

[9] O. Akinrinade and C. Du, "Skin Cancer Detection Using Deep Machine Learning Techniques," *Intelligence-Based Medicine*, vol. 11, p. 100191, 2025.

[10] S. R. S, R. K. S, V. J. S, and R. R, "ResNetVision: Harnessing Deep Learning for Early Skin Cancer Detection," in *Proc. 2025 International Conference on Computing and Communication Technologies (ICCCT)*, Chennai, India, 2025, pp. 1–5.

[11] M. Kumar, J. Kumar, and P. Kukana, "A Comprehensive Approach for Skin Cancer Detection and Classification Using Machine Learning and Deep Neural Network Algorithm," in *Proc. 2025 International Conference on Cognitive Computing in Engineering, Communications, Sciences and Biomedical Health Informatics (IC3ECSBHI)*, Greater Noida, India, 2025, pp. 984–989.

[12] S. M. Afifi, R. Kaur, H. GholamHosseini, O. Azzam, and R. Taha, "Deep Learning-Powered Mobile Application for Early Skin Cancer Detection," in *Proc. 2024 11th International Conference on Soft Computing & Machine Intelligence (ISCMI)*, Melbourne, Australia, 2024, pp. 318–322.

[13] D. Khamsa, L. Pascal, B. Zakaria, M. Lokman, and M. Y. Zakaria, "Skin Cancer Diagnosis and Detection Using Deep Learning," in *Proc. 2023 International Conference on Electrical Engineering and Advanced Technology (ICEEAT)*, Batna, Algeria, 2023, pp. 1–6.

[14] K. Suneetha, M. Srivastava, S. Singh, K. K. Brar, L. H. Alzubaidi, and R. Singh, "A Comparative Analysis of Skin Cancer Detection Using Shallow and Deep Learning Features," in *Proc. 2024 3rd Edition of IEEE Delhi Section Flagship Conference (DELCON)*, New Delhi, India, 2024, pp. 1–4.

[15] M. L. P. Augustine, S. Spanogianopoulos, and M. S. Sharif, "A Machine Learning Framework for Skin Cancer Detection Using Transfer Learning," in *Proc. 2025 IEEE 6th International Conference on Image Processing, Applications and Systems (IPAS)*, Lyon, France, 2025, pp. 1–6.