

Skin Guard: Intelligent Detection and Management of Skin Cancer

Lahari Valluri¹, Kranthi Madala², Thanuja Nadikattu³, Eswar Avisetti⁴ and Mehaetab Shaik⁵

^{1,3-5}IVyr B.Tech, VR Siddhartha Engineering College, Department of EIE, Vijayawada, India

Email: laharivalluri18@gmail.com, nadikattuthanuja9@gmail.com, eswaravisetti@gmail.com, mehaetabshaik@gmail.com

²Asst. Professor, Department of EIE, Siddhartha Academy of Higher Education (Deemed to be university) Vijayawada, A.P, India

Email: kranthimadala16@gmail.com

Abstract— Skin cancer continues to be one of the most prevalent types of cancer worldwide, with a steadily increasing number of cases reported annually. It is generally classified into two major types: benign lesions, which are harmless, and malignant lesions, which can be life-threatening if not diagnosed and treated promptly. Early detection is critical, as it significantly improves the chances of successful treatment and can greatly reduce the burden on healthcare systems. However, traditional diagnostic methods such as visual inspection and biopsy often depend heavily on clinical expertise and may not be readily accessible in all settings. This highlights the need for efficient, automated systems that can aid in the early identification of skin cancer. Techniques such as dropout regularization and learning rate scheduling were also explored to further optimize the model.

The findings suggest that with appropriate tuning and preprocessing, ANN models can serve as valuable tools in assisting dermatologists and healthcare providers in the early and reliable diagnosis of skin cancer. In conclusion, this work demonstrates the practical application of artificial neural networks in skin lesion classification and reinforces the growing role of AI in medical diagnostics. In response to this need, this project explores the use of Artificial Neural Networks (ANNs) for the classification of skin lesions using image-based data. The study leverages a publicly available Kaggle dataset composed of high-resolution dermoscopic images labelled as either benign or malignant. The primary aim is to evaluate the effectiveness of ANN models in classifying skin lesions by learning complex patterns from the visual data. A series of preprocessing steps were applied to the dataset to enhance model performance. These include image normalization, resizing, and augmentation techniques such as rotation, flipping, and contrast adjustment, which help in increasing data diversity and preventing overfitting. Feature extraction was performed directly from the pixel data, allowing the neural network to learn hierarchical features during training.

The ANN architecture was designed and optimized by experimenting with different numbers of hidden layers, neurons, activation functions (such as ReLU and sigmoid), and optimizers (including Adam and SGD). The model was trained and validated using various evaluation metrics, including accuracy, precision, recall, F1-score, and a confusion matrix to assess its predictive capabilities. The results indicate that the ANN model is effective in distinguishing between benign and malignant lesions, achieving a high level of accuracy on the validation set. This confirms the potential of ANNs in medical image classification tasks, especially in dermatology.

Additionally, the project investigate how changes in model parameters and training strategies impact performance. The model developed in this study provides a foundation for future research, including the use of more complex architectures like Convolutional Neural Networks (CNNs), integration with real-time diagnostic systems, and deployment in clinical environments.

***Index Terms*— Benign Lesions, Malignant Lesions, Artificial Neural Network (ANN), Machine Learning (ML), Image processing.**

I. INTRODUCTION

Many new cases of skin cancer are detected each year, making it one of the fastest expanding cancer types. It shows up in the skin's outer parts because UV radiation from tanning beds has accumulated over time. Skin cancer can be separated into less serious benign lesions and more serious malignant lesions which may develop into other cancers if they are left alone. Spotting and defining skin lesions as early as possible helps treat patients without delay and can save many lives.

Dermatologists rely on a visual check first and then they perform biopsies to determine the form and location of the lesion. On the other hand, these approaches may take time, depend too much on doctors' training and be limited by the number of medical experts in rural or underdeveloped areas. Also, looking at a lesion by hand can sometimes result in an error, since different lesions may appear close to each other.

Thanks to recent progress in Artificial Intelligence (AI) and Machine Learning (ML), doctors can now use medical imaging more accurately. Especially, Artificial Neural Networks (ANNs) are considered the best for patterns such as classifying images. By adjusting the weights among their different layers of synthetic neurons, ANNs, like the human brain, can learn from data and spot complicated relationships.

Working with ANNs, this project creates a model capable of labelling images of both benign and malignant skin lesions taken from Kaggle's dataset. The model seeks to identify features in dermoscopic images that differ the two skin conditions. The idea is to show that using deep learning can support dermatology and the creation of automated tools for skin cancer screening.

This goal is achieved by doing data preprocessing (changing size, standardizing and generating more images), creating a model, optimizing it and measuring its performance through standard classification scores. Different ANN structures and training strategies are examined to determine which ones result in top classification accuracy. Most places report an increase in skin cancer and the World Health Organization (WHO) forecasts nearly 4 million skin cancer cases each year.

It develops in the epidermis as a result of long exposure to both natural sunlight and tanning beds. For basal cell carcinoma (BCC) and squamous cell carcinoma (SCC), there is a slow growth, though melanoma (less common) can quickly spread and metastasize, so detecting it early is vital.

AI and ML have made a big difference in dermatology, especially for CAD. From large amounts of images, AI models can study quality aspects objectively and, at times, surpass the performance of human experts.

Doctors divide skin lesions into various categories. Seborrheic keratoses and nevi are examples of benign lesions which are not likely to be dangerous. If melanoma, BCC or SCC are not treated, they can spread through the skin to other parts of the body.

Early spotting of skin cancers is very important, as almost all those with early stage melanoma can be treated successfully. Regular methods for diagnosing stroke involve the following:

A doctor will examine the patient's eyes, as well as look at previous medical records. The initial phase is for the physician to examine the patient's skin and ask about the risk factors in the patient's medical history. Having had skin cancer in the family Being in the sun or using a tanning bed A person's skin type and how many times they have gotten a sunburn.

Dermoscopic Examination: Using dermoscopy or dermatoscopy, a clinician is able to examine skin lesions with magnified and polarized light. 3.Assessing Risk and Monitoring Process: Pictures may be taken of low-risk lesions and watched over a time period.

The next step is a biopsy (the histopathological examination). If a lesion appears suspicious, it is usually biopsied by removal. A portion of the tissue is sent to a pathology lab for study under a microscope. Here is where you get the best confirmation of if the lesion is cancerous or benign.

II. LITERATURE SURVEY

- In this section, I searched for past articles and studies online to get a clear idea for the next part of the work. Both AI and ML have played a growing part in dermatology, mostly in detecting skin cancer. Correctly

identifying diseases as early as possible is very important for patients and AI-based systems assist dermatologists with their decisions. The analysis here dives into significant literature, describing different ways researchers work with datasets and discuss the difficulties in deploying ANNs for skin lesion identification.

- Thiong'o and Rutka (2023), investigate how the use of digital twin technology can help manage the various connections between childhood brain tumors, treatment and outcomes in the brain [1]. The researchers have pointed out that certain problems in pediatric neuro-oncology are, among others, the irregular outcomes after treatment and the lack of models that can match treatment to patients exactly.
- Esteva et al. (2017) [2], developed a CNN that is reputed to classify skin lesions as effectively as expert dermatologists. The authors used a huge database of over 129,000 dermatological images covering more than 2,000 distinct diseases while building their model. Research papers included by the authors demonstrate how CAD has improved in dermatology. Initially, models were built with specially designed features and common machine learning techniques which could not deal well with the different look of images.
- Holzinger et al. (2017) [3] study in depth what makes AI-based explanations in medicine both transparent and reliable. Authors argue that though AI models succeed on test data, making them useful for many fields, they lack a way to explain their reasoning which is crucial for clinical settings where decisions are justified to both professionals and
- LeCun, Bengio and Hinton (2015) [4] highlight the major developments and core ideas that led to deep learning becoming the main approach for image understanding and visual perception tasks. They present a historical background, starting with basic multilayer perceptrons and ending with convolutional neural networks (CNNs) which are ideal for working with images.
- Schöhrer and coworkers (2023) [5], provide a newly curated dataset that describes antibodies developed to detect important parts of the RAS pathway. The experts analyzed in this paper point out that scientists and clinicians need validated antibodies for accurate basic research and clinical studies. All these techniques—Western blotting, immunoprecipitation, immunohistochemistry and flow cytometry—use antibodies to study protein activity and expression.
- S. C. Y. Lu (2013) [6], Lu studies the performance of various efficient initialization methods for K-means in 2013 [8]. According to the paper, despite minimizing the variance within clusters, implementation errors in initialization can make solutions less accurate or more challenging to achieve. Using random initialization which is simple, often leads to a wide range of results in clustering. By choosing initial centers using an approach from probabilities based on distance, K-means++ enhances the performance measured both theoretically and experimentally.
- Kawahara, BenTaieb, and Hamarneh (2016) [7] use deep features from CNNs to differentiate between different types of skin lesions in dermoscopic images. To save time, they use pre-existing learning, applying AlexNet and VGG, as feature extractors. The researchers applied these networks, first trained on ImageNet, to medical images. Then, the networks' top features were used to classify lesions with SVMs. This method addresses the usual problem with medical imaging which is the difficulty in finding large, annotated datasets.
- Mirikharaji et al. (2022) [8], reviewed 177 studies on deep learning methods for segmenting skin lesions. The different factors analyzed were data, preprocess methods, model designs and methods used to measure performance. Researchers identify three major obstacles: the presence of artifacts, inconsistencies in lesion appearance and the different standards for obtaining images. You can find a rich summary of the surveyed works in the interactive table.
- Tschandl et al. (2018) [9], The dataset used for this study contains 10,015 dermoscopic images from seven distinct categories of pigmented skin lesions. Its purpose is to stand as a reference data set for modeling and testing in dermatology.
- Ahmed, A. Ali (2025) [10], This work concentrates on combining attention mechanisms—self, soft and hard attention—with the pre-trained Xception model for identifying skin cancer as normal or malignant. It is shown that by using attention mechanisms, the model does better at detecting benign and malignant skin lesions.
- Fatima Al Zegair, Nathasha Naranpanawa, Brigid Betz-Stablein, Monika Janda, H. Peter Soyer and Shekhar S. Chandra (2023) [11] describe how machine learning can be used to spot melanoma and identify 'ugly duckling' naevi. It points out that using computer-aided diagnosis systems improves the accuracy and choice of treatment in detecting skin cancer.

- Hossain Shahriar, MstShapna Akter, Sweta Sneha and Alfredo Cuzzocrea(2023)[12] outlined a deep learning approach for correctly identifying many kinds of skin lesions in this paper. Many deep learning networks such as ResNet-50, VGG-16 and InceptionV3, are applied to the HAM10000 dataset for classifying the lesions.

III. PROPOSED MODEL

The system developed here tries to build an effective and precise skin cancer detection model using Artificial Neural Networks (ANN). The system is built to place skin lesions into two major groups: benign (non-cancerous) and malignant (cancerous) using images. For this reason, the system depends on a publicly accessible Kaggle data set which features dermoscopic pictures labeled by diagnoses. At the start, the data is analyzed on Kaggle for quality, consistency and the number of instances in each class. Images are often resized, normalized and augmented (rotated, flipped or brightness adjusted) to support a better performance and more stable training process. This technique helps to correct the problem of skewed classes and enhance the capacity to use new data. Feature extraction from preprocessed images covers color, texture and shape. These features are used by the ANN model to tell benign lesions from malignant ones.

A seven-layer multi-layer feedforward Artificial Neural Network is used for this classification task. It covers: An input layer to process the drew features. Hidden layers have functions called ReLU(Rectified Linear Unit) that change in a non-linear way.

In binary classification, the function uses a sigmoid and for multi-class, the function uses a SoftMax activation in the output layer. Overfitting is lowered by using dropout, stabilizing learning with batch normalization and modifying the learning rate using scheduling techniques.

The model's training set is only a portion of the data, whereas a distinct validation set is used to change hyperparameters like learning rate, an epoch limit and batch size. With the help of accuracy, precision, recall, F1-score and a confusion matrix, performance measures are taken to see how effective the model is at telling the two types of data apart.

When the model is trained, it can review a dermoscopic image and determine if it is benign or malignant. Having these prediction results allows healthcare professionals to evaluate patients with a quick initial diagnosis, helping them notice problems sooner.

Future improvements may include adding clinical information (eg, patient age and lesion location) along with the image data using a multi-modal approach. Also, if we include tools like Grad-CAM, we can display which locations in the image affected the model, boosting transparency and helping doctors put their faith in the system.

IV. METHODOLOGY




Fig 1: Block Diagram of the proposed methodology

Dataset Description: Kaggle's dataset has thousands of photos grouped into two groups: benign and malignant. Eighty percent of the data was used for training, ten percent for validation and the last ten percent for testing. Each image had a different resolution, so I had to preprocess them before they could be used for ANN.

Preprocessing Steps: The size of the images was changed to 128x128 so they would be standard. The pixel values were rescaled to include only 0 and 1 to help improve how the model trained.

New images created using rotation, flipping and zoom were used to diversify the dataset and help prevent learning only from the existing training set. The data was put into a NumPy array format and saved for productive operation.

Model Architecture: A TensorFlow/Keras model was built with all neural connections in place. The architecture ensured an input layer was present for receiving the flattened data. Each hidden layer used ReLU and a dropout layer was incorporated to reduce overfitting. To classify samples as benign or malignant, the output layer used a SoftMax function.

V. RESULTS

Training Process and Results: The Adam optimizer and categorical cross-entropy loss were used to make the model. Training in batches of 32 and with a learning rate of 0.001 was used. The model completed 30 epochs of training. In fig 2, the accuracy settled at 50%, showing there is still room for improvement. A confusion matrix was applied to the model's predictions on test images. Figures 3 and 4 display the output from our analysis.

```

21/21 ————— 3s 122ms/step
Test Accuracy: 0.5045
Classification Report:

```

	precision	recall	f1-score	support
Benign	0.55	0.48	0.51	360
Malignant	0.46	0.53	0.49	300
accuracy			0.50	660
macro avg	0.51	0.51	0.50	660
weighted avg	0.51	0.50	0.51	660

Fig 2: Accuracy of the trained model



Fig 3: Predicted class of image is benign




Fig 4: Predicted class of image is malignant

VI. CONCLUSION

Since the ANN's accuracy is just 50%, additional changes need to be made to improve its ability to predict results. It shows that both the model design and the readiness of the dataset must be given careful attention. Although several major preprocessing steps were put in place, attempting new methods and adding powerful models might increase classification performance. Reliable and accurate medical image analysis depends highly on the quality and how the dataset is organized.

FUTURE WORK

Hyper parameter Tuning: Modifying learning rate, changing batch size and playing with the number of network levels are ways to improve both training and performance.

Use of Pre-trained Models: This approach helps to reduce the time needed for development and helps improve outcomes by using learned patterns from enormous datasets in existing architectures such as VGG16 and ResNet50.

Dataset Expansion: Using greater amounts and types of training images will help cut down any bias and boost the model's general performance. By putting together various models, we may achieve better returns since the possible errors of separate models are reduced.

Enhancements made to the basic neural network architecture by adding more layers and using CNN-based models increase the model's sensitivity to small changes seen in medical images. Identifying the best combination of hyperparameters can be done by performing grid search, random search or Bayesian optimization techniques.

Transfer Learning Enhancement: Using Efficient Net or Dense Net, in addition to VGG16 or ResNet50, can bring even better results if you train them with your data.

Data Quality Improvement: At the same time, it's important to pay attention to image sharpness and proper marking of every image.

REFERENCES

- [1] Grace M. Thiong'o and James T. Rutka, "Digital Twin Technology: The Future of Predicting Neurological Complications of Pediatric Cancers and Their Treatment".
- [2] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun, "Dermatologist-level classification of skin cancer with deep neural networks".
- [3] Andreas Holzinger, Chris Biemann, Constantinos S. Pattichis, and Douglas B. Kell, What do we need to build explainable AI systems for the medical domain?.
- [4] Y. LeCun, Y. Bengio, and G. Hinto, "Deep Learning for Computer Vision: A Brief Review.
- [5] "Regine M. Schoenherr, Dongqing Huang, Uliana J. Voytovich, Richard G. Ivey."A dataset describing a suite of novel antibody reagents for the RAS signaling network".
- [6] S. C. Y. Lu,"A comparative study of efficient initialization methods for the k-means clustering algorithm". Jeremy Kawahara, Aïcha BenTaieb, and Ghassan Hamarneh, "Deep Features to Classify Skin Lesions".
- [7] Mirikharaji, Z., Navarrete-Dechent, C., Lee, T. K., & Hamarneh, G. (2022). A Comprehensive Survey of Deep Learning-Based Skin Lesion Segmentation: Datasets, Methods, and Challenges.
- [8] Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 Dataset: A Large Collection of Multi-Source5 Dermatoscopic Images of Common Pigmented Skin Lesions.
- [9] D. Ahmed, A. Ali (2025). Enhancing Skin Cancer Classification with Attention Mechanisms Integrated into the Xception Transfer Learning Model.
- [10] Fatima Al Zegair, Nathasha Naranpanawa, Brigid Betz-Stablein, Monika Janda, H. Peter Soyer, Shekhar S. Chandra (2023). A Review of Machine Learning Applications in the Detection of Melanoma and Identification of 'Ugly Duckling' Naevi.
- [11] MstShapna Akter, Hossain Shahriar, Sweta Sneha, Alfredo Cuzzocrea (2023)," A Deep Learning Approach for Accurate Classification of Skin Lesions Using the HAM10000 Dataset".