

ARTICLE

Skin Sentinel: An artificial intelligence-based skin cancer detection system

Bhupesh Kumar Namdeti^{1†}, Anoosha Shaga², Vayunandan Kumar Konakalla³, Chittibabulu Sape⁴, Veeramallu Bobba⁵, and Satyanarayana Gunupusala^{6†*}¹Department of Electrical and Electronics Engineering, Vidya Jyothi Institute of Technology, Hyderabad, Telangana, India²Department of Information Technology, VNR Vignana Jyothi Institute of Engineering and Technology, Hyderabad, Telangana, India³Fidelity Investments, Dallas, Texas, United States of America⁴Department of Computer Science and Engineering, Aditya University, Surampalem, Andhra Pradesh, India⁵Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Andhra Pradesh, India⁶Department of Computer Science and Engineering, Vasireddy Venkatadri Institute of Technology, Nambur, Andhra Pradesh, India

[†]These authors contributed equally to this work.

***Corresponding author:**
Satyanarayana Gunupusala
(snarayana.5813@gmail.com)

Citation: Namdeti, B. K., Shaga, A., Konakalla, V. K., Sape, C., Bobba, V. & Gunupusala, S. (2026). Skin Sentinel: An artificial intelligence-based skin cancer detection system. *Int J Systematic Innovation*, 10(4): 026140046.
[https://doi.org/10.6977/IJoSI.202608_10\(4\).0008](https://doi.org/10.6977/IJoSI.202608_10(4).0008)

Received: April 4, 2026

Revised: May 15, 2026

Accepted: June 2, 2026

Published online: July 3, 2026

Copyright: © 2026 Author(s). This is an Open-Access article distributed under the terms of the Creative Commons Attribution License, permitting distribution, and reproduction in any medium, provided the original work is properly cited.

Publisher's Note: AccScience Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Abstract

Skin cancer is the most common malignancy worldwide, and effective early detection remains essential despite advances in prevention and photoprotection strategies. This study introduces an intelligent system for skin cancer diagnosis via deep learning methodologies. The suggested hybrid model combines the best aspects of MobileNet and long short-term memory (LSTM) networks to improve skin lesion image classification. The technology was designed not only to help doctors make accurate diagnoses, but also to support real-world healthcare tasks, including uploading images, automatically analyzing them, suggesting treatments, and scheduling appointments for patients and doctors. The experimental findings indicate that the proposed hybrid model achieves 93% accuracy, surpassing several current models, including support vector machine, convolutional neural network, Visual Geometry Group, ResNet, and MobileNet. Combining MobileNet with LSTM improves the ability to extract and classify features. Early detection of skin cancer can make therapy more effective and lower death rates. The suggested method shows great promise for use in real-world medicine. For future work, the system can be improved by using increasingly diverse datasets, more advanced deep learning architectures, and real-time clinical deployment to make diagnoses more accurate and reliable. The suggested system is intended to facilitate early skin lesion analysis and help medical practitioners with initial diagnosis by acting as an artificial intelligence-assisted screening and decision-support tool.

Keywords: Dermatological manifestation; Supervised machine learning; Convolutional neural network; MobileNet; ResNet; Hybrid algorithm

1. Introduction

Image processing is a method by which an image is enhanced or manipulated to bring out required features using editing operations. With such a tremendous rise in popularity in recent years, it has emerged as one of the major areas of research in science and engineering. This, without a doubt, is one of the best medical image types for scholars to study, since the dawn of image processing (Al-Masni *et al.*, 2018). Because there are not many dermatologists, especially in rural regions, the global prevalence of dermatological illnesses calls for simple diagnostic options. It is possible to achieve promising results in image identification and accurate skin condition classification by utilizing artificial intelligence (AI), such as convolutional neural network (CNN) and MobileNet architectures (Alfred *et al.*, 2015). Healthcare workflow is optimized, and telemedicine is improved by incorporating AI in line with digital transformation trends. By tackling accessibility issues and democratizing healthcare, this study seeks to develop a multilingual, AI-based tool for first-time dermatological diagnosis. This study improves patient outcomes and enhances the utility of AI in dermatology by enabling early therapy. The remainder of this paper is organized as follows. Section 1.1 presents the problem statement, while Section 1.2 outlines the research objectives. Section 2 reviews the relevant literature, and Section 3 describes the dataset, methodology, and deep learning models employed in the study. Section 4 presents and discusses the results, followed by the conclusions and future research directions in Section 5. The performance of the proposed deep learning models was evaluated based on the results obtained.

1.1. Problem statement

Despite advances in medical technology, there remains a considerable shortage of dermatological services, especially for initial diagnosis. Treatments and diagnoses are delayed due to a shortage of dermatologists and geographic barriers to specialized care. This service gap disproportionately impacts underprivileged and rural communities, which frequently leads to worsening health outcomes. There are two main issues: first, an unequal distribution of dermatological knowledge; second, the increasing prevalence of skin diseases worldwide. Hence, the created tool needs to serve as a preliminary diagnostic procedure that is easily accessible, fast, and accurate. The main aim is to create an AI-driven diagnostic instrument that can rapidly provide initial evaluations of dermatological disorders. The tool seeks to reduce the burden on the healthcare system by facilitating rapid triage and prioritization, potentially reducing patient wait times

and enabling more equitable workloads across available dermatological treatments.

1.2. Aims and objectives

The aim of the study is to create an improved deep learning model for skin cancer identification to diagnose skin diseases at their early stages, which would increase dermatological therapy's accessibility and effectiveness. This goal is supported by the following particular objectives:

- (i) To effectively classify skin disorders from photographs using CNNs and MobileNet architectures. The study aimed to evaluate whether these AI models could achieve diagnostic performance comparable to or exceeding that of experienced clinical specialists for identifying and classifying various skin lesions.
- (ii) To make an interface that will suit both healthcare providers and patients with ease. A patient should be able to upload their skin photo along with any relevant clinical information, while medical personnel should find it easier to review diagnostic report submissions.
- (iii) To provide a language support feature to gain knowledge in diverse languages. The tool will thus guide in English and Telugu as well to cater to wider population groups, while eliminating linguistic barriers to accessing dermatology care.
- (iv) To evaluate how the AI tool affects patient throughput, precision, and speed throughout the first diagnosis procedure. Furthermore, it is essential to evaluate the tool's impact on patient satisfaction, the timeliness of patients' receipt of a preliminary diagnosis, and the efficiency of the diagnostic process.
- (v) To determine if AI technology's ability to triage effectively would reduce the workload of dermatologists. If the technology is used to identify situations requiring an urgent response, healthcare systems could better distribute and allocate resources.

By accomplishing these goals, the research will enhance patient care by providing prompt and accurate preliminary diagnoses, thereby reducing the current burden on dermatologists. Figure 1 shows the major types of skin cancer considered in this study.

2. Related work

Numerous studies have been cited regarding the diagnosis of skin cancer. Commonly, features based on asymmetry, border, color, and diameter, as well as gray level, are used to extract statistical and textural features with greater focus, and are then used with classification algorithms such as support vector machine (SVM), MobilenetV2, ResNet, and Visual Geometry Group (VGG). Skin lesion segmentation in dermoscopy images via deep, full-resolution models.




Figure 1. Types of skin cancer

The study by Al-Masni *et al.* (2018) focuses on using convolutional networks (ConvNets) to segment skin lesions in dermoscopy images. However, it is challenging to segment the skin lesions due to variations across frames and hazy boundaries. This segmentation plays a crucial role in differentiating melanoma. In addition, they focused primarily on 14 features: age, infection, gender, days, body mass index, smoking status, region, bleeding, diagnosis, growth, family history, allergy, pain level, and itchiness level. Fuzzy-rough cognitive networks were evaluated using PH2 and ISBI datasets, which are available on Kaggle. Fully convolutional networks, U-Net, and SegNet were compared with this fuzzy-rough cognitive networks approach. The overall segmentation accuracy for this approach was 94.03% on the ISBI 2017 test data and 95.08% on the PH2 dataset, respectively. When distinguishing different kinds of skin lesions, including benign and malignant cases such as carcinoma and seborrheic keratosis, fuzzy-rough cognitive networks performed better than other methods. Therefore, the authors concluded that segmentation performance can be enhanced by identifying more precise and significant features using the full spatial resolution of the input images.

Alfed *et al.* (2015) reported that higher death rates are a result of the growing number of melanoma occurrences. In particular, when melanoma is detected at Stage I, early identification is essential to improving survival rates. Segmenting melanoma tumors accurately is still difficult, though. The study provided an artificial bee colony (ABC) algorithm, a unique method for melanoma identification in digital photographs. This approach is fast, flexible,

easy to use, and requires fewer parameters than previous algorithms. PH2, the ISBI 2016 and 2017 challenges, and Dermis were among the datasets used to test the proposed approach. Images with different brightness, resolution, and other properties can be found in these datasets. To reduce noise in the photos, morphological filtering was applied. The best threshold value for melanoma detection was found using the ABC algorithm. In high-specificity settings, the technique outperformed ground-truth photographs verified by dermatologists. Across the four databases, the suggested approach showed excellent average accuracy (95.24–97.61%) and Jaccard's coefficient (83.56–85.25%) for melanoma detection. A comparison of the suggested algorithm with other melanoma detection techniques found in the literature demonstrated its superiority. Additionally, for diagnosing skin cancer, the lesion location must be accurately segmented. The suggested method showed that lesion detection could be performed entirely automatically without the need for a training period.

Aljanabi *et al.* (2015) proposed a novel method for melanoma identification that uses the ABC algorithm. This technique performed well across a variety of datasets, suggesting it has the potential to improve the efficiency and accuracy of melanoma identification.

Ech-Cherif *et al.* (2019) suggested a mobile skin cancer diagnosis system using deep neural networks (DNNs) for dermoscopy applications. However, diagnosing skin cancer using dermoscopic pictures can be difficult for both patients and medical personnel due to the variety of ways that skin lesions can manifest. Moreover, patients

can benefit from patient-centric health management by using their smartphones for early detection through mobile teledermatoscopy. This work uses transfer learning to train a Binary MobileNetV2 model to classify benign and malignant skin lesions. The trained model achieved an overall accuracy of 91.33% on skin lesion classification with a batch size of 32. To create an iOS mobile application, the author used the trained model and the Core Machine Learning library. According to recent studies from renowned institutes, deep CNNs have proven substantial success in skin cancer identification. Promising outcomes from this study's application of DNNs and transfer learning suggest the possibility of precise, easily accessible skin cancer diagnosis via mobile platforms. The study emphasizes how mobile teledermatoscopy can help detect skin cancer early and provide patient-centered healthcare. The proposed mobile skin cancer detection system achieves high accuracy in classifying skin lesions through deep learning, specifically transfer learning. The development of mobile applications for skin cancer diagnosis creates new opportunities for prompt, easily accessible healthcare interventions, especially in places where traditional medical facilities are scarce.

Waheed *et al.* (2017) proposed an effective machine learning method for identifying melanoma from dermoscopic images. Melanoma skin lesions are identified by their distinguishing characteristics. The suggested method begins by extracting various color and texture features from dermoscopic images, based on distinct structures and varying intensities of melanocytic lesions. To identify melanoma in dermoscopic images, the classifier was trained on the derived features in the second stage. This research also discusses the significance of color and textural features in melanoma detection. The public PH2 dataset was used to assess the accuracy, sensitivity, and specificity of the suggested approach.

Hoshyar *et al.* (2020) proposed a novel method for early melanoma detection using a multiclass SVM. Seborrheic verruca, basal cell cancer, nevus myticus, squamous cell cancer, and solar keratosis (actinic keratosis) were the five categories of skin lesions. The proposed method used an automated procedure to select and match the obtained images to more likely types to classify the melanoma type. Among the most effective tools for handling classification difficulties was the multiclass SVM. The method relies on using certain training examples to help it learn at each step. Here, the gradient, contrast, and edge properties of the color and texture were retrieved. A collection of images featuring all five forms of melanoma was included in the proposed system for testing and categorization needs. According to the simulation results, among the five

varieties, the suggested SVM strategy achieved relatively high accuracy.

Pham *et al.* (2018) made two major advances using CNNs and data augmentation, including a classification model that improved performance on skin cancer classification. Their second contribution was a demonstration of how to use image augmentation to address limited data. Additionally, the impact of varying numbers of augmented samples on the performance of different classifiers was investigated using image augmentation. The largest public dataset, comprising 6,162 training and 600 testing photos, was used in this study. Furthermore, the impact of every image enhancement on three different classes was investigated. Additionally, it was noted that, in contrast to conventional procedures, the behavior and abilities of each class ID were influenced differently by each augmentation. Therefore, data augmentation can be employed to further improve the efficacy of skin cancer categorization.

Rezaoana *et al.* (2020) proposed an automated method for skin cancer identification and categorization. In the study, the behavior and capabilities of deep CNNs were observed and evaluated, along with the categorization of nine distinct types of skin cancer. There were nine different forms of skin cancer in the dataset used in this investigation: basal cell carcinoma, squamous cell carcinoma, nevus, actinic keratosis, pigmented benign keratosis, dermatofibroma, seborrheic keratosis, melanoma, and vascular lesions. The study's primary goal was to build a model that uses CNNs to diagnose and categorize skin cancer. Skin cancer diagnosis involves deep learning and image processing. The number of images was increased by applying different image segmentation techniques. A transfer learning strategy was applied to further increase the precision and effectiveness of categorization tasks. The overall accuracy of this suggested CNN-based model was 79.45%.

Zhang *et al.* (2020) used image processing techniques to detect skin cancer at an early stage. This makes use of optimization and optimum convolutional neural networks. The algorithm of choice for optimization was used to obtain the optimal weights and biases within the network. The purpose of this is to minimize errors in both the intended and network outputs. The study used the Dermquest and DermIS digital datasets. The results of the investigation were contrasted with those from various techniques, including the spot mole tool, the semi-supervised approach, regular CNN, AlexNet, VGG-16, LIN, Inception-v3, and ResNet. The study yielded the best outcome for skin cancer, according to the final data.

Carcagni *et al.* (2019) proposed a semi-supervised, self-advised learning model for automatic melanoma detection

using dermoscopic images. To optimize the separation of labeled data, a deep belief architecture was built using both labeled and unlabeled data, with fine-tuning performed via an exponential loss function. In parallel, the impact of incorrectly categorized data was mitigated by a self-advised SVM method, improving the classification outcomes. The redundancy and generalization properties of the model were enhanced by training polynomial- and radial basis function-based simulated annealing SVMs and deep networks on training samples randomly selected via a bootstrap technique. Then, using least-squares estimation weighting, the findings were combined. A total of 100 dermoscopic images were used to evaluate the proposed model. When the suggested model with deep neural processing was compared with well-known classification methods, it outperformed most popular approaches, including k-nearest neighbors, artificial neural networks, SVMs, and semi-supervised algorithms such as expectation-maximization and transductive SVMs.

Vijayalakshmi (2019) provided a fully automated approach to dermatological disease recognition from lesion photographs, in contrast to traditional, medical-personnel-based identification. The three stages of the model's design were data collection and augmentation, model design, and prediction. It also combined several AI algorithms—such as CNNs and SVMs—with image processing techniques to create a more effective architecture, increasing accuracy to 85%.

Subramanian *et al.* (2021) developed a CNN model with precision and accuracy exceeding 80% and a false negative rate below 10%. Dermatoscopic images of skin lesions comprise the HAM10000 dataset, which was used in the investigation. There were 10,015 colored photos in the dataset overall, with most having a resolution of 600×450 . A number of research articles and techniques were examined and tested. Using the HAM10000 dataset, an accuracy of at least 80% was attained. The study's ultimate outcome demonstrated that the most effective method for detecting skin cancer was a routine CNN.

Shah & Ross (2009) used the MED-NODE dataset of digital images to develop a system for identifying melanoma. The dataset included a variety of artifacts, which were first removed by preprocessing. Next, the active contour segmentation approach was applied to extract the region of interest. From the segmented portion, a variety of color features were retrieved, and three classifiers—decision tree, k-nearest neighbor, and naïve Bayes—were used to assess the system's performance. The system outperformed other classifiers, achieving an accuracy of 82.35% with the decision tree.

Hameed *et al.* (2018) presented a review of computer-

aided diagnosis systems and examined current practices across various stages of these systems. An analysis and reporting of the statistics and outcomes of the most significant and recent implementations were performed. They compared the performance of recent work across parameters such as accuracy, datasets, computational time, color space, machine learning techniques, etc., and summarized the results in a table to provide an overview of emerging researchers in the field of computer-aided skin diagnosis systems. The various components of computer-aided skin cancer diagnosis systems present research problems that were also emphasized.

Masood *et al.* (2013) provided statistics and findings from the most significant implementations to date. They examined the results of various classifiers developed specifically for diagnosing skin lesions and compared their performance. Indications of several factors that affect the procedure's performance were presented where available. They examined findings from various models and propose a framework for comparing evaluations of skin cancer diagnostic models. Some shortcomings of current studies were highlighted, and recommendations for additional research were provided.

Hasan *et al.* (2019) reported the use of automatic skin cancer detection using CNNs to categorize skin cancer images as benign or malignant. In the study, feature extraction techniques were used to extract the characteristics of skin cells impacted by cancer. The retrieved features were sorted in the next step using CNNs. Using the publicly available dataset, this method yielded an accuracy of 89.5% and a training accuracy of 93.7%. The method can be regarded as a standard for identifying skin cancer, as evidenced by the experimental and assessment sections.

Deep learning has updated more efficient architectures, such as EfficientNet, vision transformers (ViT), explainable AI (XAI), and federated learning, for medical image analysis. EfficientNet is a family of models that use increased depth within their components, along with an asymmetric, layer-optimized parameter-scaling scheme, to improve classification performance. ViTs have performed well as feature extractors because the components themselves capture global image features. XAI techniques enhance transparency by visualizing how models make decisions. Federated learning enables single-model training across thousands of medical institutions while keeping data private. These new methodologies have a highly promising future in the field of systems for diagnosing skin cancer.

3. Methodology

In this study, we first describe the workflow of the

proposed system to identify skin cancer. Then we discuss the system's working principle. Figure 2 shows that the workflow begins with the collection of skin cancer images. To improve consistency and quality, preprocessing was applied to the gathered photos. In this step, pixel values were normalized, contrast was improved, and images were resized. Data augmentation, including rotation and scaling, was used to improve the dataset's robustness. The CNN and MobileNet architectures were trained and validated on the preprocessed dataset. During the training phase, the model was trained using images and their corresponding diagnostic labels to learn patterns associated with different skin disorders. Validation tests were conducted throughout the training process to continuously evaluate the model's performance and assess its ability to generalize to unseen data.

3.1. Dataset

The dataset used in this research was obtained from the publicly available skin lesion image datasets on Kaggle. The data included nine distinct classes of skin lesions: actinic keratosis, basal cell carcinoma, dermatofibroma, melanoma, nevus, pigmented benign keratosis, seborrheic keratosis, squamous cell carcinoma, and vascular lesion (Figure 3). To ensure that all images contained the same content and were compatible with deep learning models, they were resized to 224×224 pixels before training. The data were split into a 70% training set and a 30% testing/

hold-out set. Other patient-level information, such as age, gender, body mass index, smoking status, family history, allergies, pain level, and itchiness level, was also included to enhance model learning. Class imbalance was addressed using data augmentation techniques, thereby improving model generalization.

3.2. Data preprocessing

The dermoscopic images obtained were preprocessed before modeling to enhance image quality and model performance. First, all images were converted into a fixed size of 224×224 pixels. Pixel normalization scaled the image values into the same range. Lesion visibility was further improved utilizing contrast enhancement and noise reduction techniques. Additionally, data augmentation methods such as rotation, zooming, horizontal flipping, scaling, and brightness adjustment were applied to increase dataset diversity and minimize/avoid overfitting. Furthermore, the patient-related data were anonymized to maintain patient confidentiality and comply with ethical considerations regarding medical data. The details of the dataset used for training and testing are summarized in Table 1.

From Figure 4, we observe that the system includes roles such as admin/owner, patients, and doctors, each with unique functionalities and interfaces. Admin/owner role is primarily responsible for overseeing the system's operation. The AI model may be managed and trained




Figure 2. Flowchart for the proposed methodology




Figure 3. Images of skin lesions. (A) Actinic keratosis. (B) Basal cell carcinoma. (C) Dermatofibroma. (D) Melanoma. (E) Nevus. (F) Pigmented benign keratosis. (G) Seborrheic keratosis. (H) Squamous cell carcinoma. (I) Vascular lesion.

Table 1. Dataset used for training and testing

Train/Test	Classification	Images	Percentage
Training	Actinic keratosis, basal cell carcinoma, dermatofibroma, melanoma, nevus, pigmented benign keratosis, seborrheic keratosis, squamous cell carcinoma, vascular lesion	7,000	70%
Testing	Actinic keratosis, basal cell carcinoma, dermatofibroma, melanoma, nevus, pigmented benign keratosis, seborrheic keratosis, squamous cell carcinoma, vascular lesion	3,000	30%
Total		1,0000	

with fresh data by the admin, guaranteeing that it keeps getting better. They are also responsible for monitoring user behavior, protecting data, and ensuring the system runs smoothly. Feedback is another key responsibility. After registering and logging in, the patient can use the system to post skin problems with pictures, including details such as location, time, prior medical treatment history, and knowledge. These will be analyzed by the analysis system, yielding a superficial result, a diagnosis, and even treatment recommendations.

The patient’s role includes using the system. After registering and logging into the system, he/she can post skin problems with their pictures, along with details such as location, time, prior medical treatment history, and knowledge. These will be analyzed by the analysis system, which will provide a superficial diagnosis and even treatment recommendations, and send them to the specialists who will help with that treatment. Patients may also schedule an appointment with a dermatologist in the function. Patients may update all their appointments




Figure 4. Block diagram of the complete project

and profiles on the website and use it to follow up and communicate with their providers. Doctors, in this case, will use the site to ensure there is no patient congestion and that the needy are cared for. Figure 4's overview of the project will ensure that each and every task in the system is designed to work hand in hand with one another, allowing the dermatological diagnoses to be leveraged for value and effectiveness through user-friendly interfacing.

3.3. Activation function

After applying an activation function to our model, we used the rectified linear unit (ReLU). The motive for using the RELU activation function is that it introduces nonlinearity, which helps the network learn complex patterns in data. Mathematically, RELU is straightforward to compute, and its gradient is well-defined. RELU helps mitigate the vanishing gradient problem.

The ReLU activation function is defined as $f(x) = \max(0, x)$, where (x) represents the input. It outputs the input value when (x) is positive and returns 0 when (x) is negative or equal to zero.

3.4. Max pooling

Next, we applied max pooling to our image to preserve its significant features while minimizing its size. This also helps reduce the model's computational cost. We can speed up training by reducing the number of parameters and operations needed in subsequent layers by downsampling the feature maps. Additionally, this will improve the

features.

3.5. Flatten and fully connected layer

After performing all these convolution operations, we flatten the output into a one-dimensional array and pass it to a fully connected layer, where we use the ReLU activation. For the last layer, we used softmax because we want to classify nine different types of skin lesions. The raw score was converted into a probability distribution over several classes using the softmax method. This enables the model to forecast each class's probability, resulting in more comprehensible and direct predictions.

3.6. Loss function

Following this, we aimed to reduce our loss function during training to improve our model's accuracy and performance. To do this, we used categorical cross-entropy, a loss function appropriate for multiclass classification problems where the output labels are one-hot encoded vectors.

3.7. Model training and validation

The proposed deep learning models were implemented in TensorFlow and a Keras-based framework (in Python, version 3.10). It was trained using a batch size of 32 and for 50 epochs. The network parameters were optimized using the Adam optimizer with a learning rate of 0.001. The loss function is categorical cross-entropy for multiclass skin lesion classification. To evaluate model performance and

minimize training bias, validation accuracy and validation loss were continuously monitored during training. While this work used a 70/30 train-test split, future studies would add methods such as k-fold cross-validation and external validation to enhance robustness and clinical reliability.

3.8. Proposed models

3.8.1. Hybrid (MobileNet and long short-term memory)

The hybrid algorithm, specifically designed for this dermatological project, cleverly blends MobileNet with long short-term memory (LSTM). By combining the best features of the two designs into one, the instrument gains strength and power, which is essential for tracking changes in skin over time. The hybrid model in question is based on MobileNet, which has been shown to be efficient on high-resolution imagery. In fact, it is well-suited to extracting all levels of detail essential to understanding which patterns are critical for recognizing different conditions from images of the skin, a subfield of computer vision known as prevalent object detection. This algorithm suggests how skin conditions change over time in the medical field. This information about the disease course can be effectively obtained using LSTM-based concurrence patterns, providing insight into the disease trajectory and enabling more accurate prognosis and therapy planning. LSTM is a specialized form of recurrent neural network designed to process sequential data and learn long-term dependencies within input sequences. When combined, LSTM and MobileNet offer such a potent tool. LSTM is added in the temporal aspect of how these conditions change, while MobileNet is dedicated more to the precise identification of skin abnormalities in static photos. Such

synergy makes the hybrid algorithm a more effective tool in dermatological diagnostics, offering improved diagnosis and a better understanding of the disease. Figure 5 shows how the CNN works.

Long short-term memory networks are commonly used for sequential and time-series data; in this work, an LSTM layer was incorporated after MobileNet-based feature extraction to capture deeper correlations among the extracted feature vectors. MobileNet is used to effectively extract spatial features from dermoscopic images, while the LSTM layer helps learn feature relationships, thereby enhancing classification performance. The fused combination outperformed the individual CNNs MobileNet and ResNet. (Chittibabulu *et al.*, 2026).

3.8.2. Convolutional neural network

We chose the CNN algorithm for the present dermatological diagnostics due to its incredible ability to parse large numbers of images at high speed. Thus, deep learning algorithms such as CNNs are ideally suited to parse dermatological images because they are designed to learn from visual data. The algorithm employed in this project comprises multiple layers that detect specific features in the skin picture data it examines. As more layers follow, more complex patterns, such as skin lesion textures or shapes, become evident, whereas the preceding layers identify fewer patterns or features. To identify a sufficient number of issues, hierarchical feature extraction is required. On the other hand, CNN receives a large amount of skin lesion data along with its annotation for network training.

Consequently, the network can establish connections between image elements and various skin diseases. To




Figure 5. The working principle of a convolutional neural network

do this, one requires learning using backpropagation and optimization techniques that adjust the network's weights to reduce discrepancies between real and anticipated diagnoses. After training, the CNN uses the learned features to evaluate new skin photos and identify the most likely dermatological disease.

3.8.3. MobileNetV2

MobileNet is a family of lightweight deep learning models designed for efficient on-device inference. These models were developed by researchers at Google, with the primary aim of enabling high-performance deep learning tasks on resource- constrained devices such as smartphones, Internet of Things devices, and embedded systems. Due to its simplified architecture, which strikes a compromise between performance and computing resource requirements, it stands out among CNNs. Compared with conventional convolutions, depth wise separable convolutions have significantly fewer parameters and require less computational power; this is the primary innovation in MobileNet. Filters that aggregate inputs from every channel of the preceding layer were applied in a conventional convolution layer. This procedure was divided into two layers, using depth wise separable convolutions: a depth wise convolution followed by a pointwise convolution. There is just one filter used on each input channel by the depth wise convolution, and the result of the depth wise convolution is combined across channels by the pointwise convolution using a 1×1 convolution. With this method, the computational cost and model

size were significantly reduced without a significant loss in performance. MobileNet is accurate and not overly compact, making it widely applicable for the purposes provided, such as object and face recognition and image identification.

From Figure 6, we note that multiclass classification consists of a classifier head (on the right) and a contracting path (on the left). The contracting route follows the usual architecture of a convolutional network by continually performing two 3×3 convolutions (unpadded convolutions), each of which is followed by a ReLU and a 2×2 max pooling operation with stride 2 for downsampling. The network is made deep by repeatedly repeating these three processing steps, or "blocks," which culminate in a set of completely connected layers (classifier stage). To build local weighted sums (also known as "feature maps") at each layer, the convolutional layers compute filters that are repeatedly applied across the entire dataset to maximize training efficiency. Then, the nonlinear layers enhance the nonlinearity of the feature maps. Finally, by using a pooling layer to subsample non-overlapping portions of the feature maps, the network can combine local features to identify complex properties. At each step of the down sampling procedure, the number of feature channels was doubled. Using max pooling, the largest element from the updated feature map was selected. An alternative method for calculating the average pooling is to take the largest element. Following the final block of the MobileNetV2 stage, the feature map was vectorized and passed to the classifier stage, which consisted of a single fully connected




Figure 6. Basic MobileNetV2 architecture

layer. We combine these attributes to form a model using fully connected layers.

3.8.4. ResNet

Among CNNs, ResNet, also known as residual networks, is another family of efficient models well-known for its exceptional performance across a variety of computer vision tasks. By adding shortcut connections, or skip connections, ResNet solves the vanishing gradient issue that arises in deep neural networks and allows the model to bypass layers during training. This facilitates training far deeper networks by allowing the network, rather than explicitly learning the essential underlying mapping, to learn residual functions. These skip connections are found in residual blocks, which make up ResNet's design. Usually, each block has identity shortcuts and many convolutional layers. These shortcuts make it possible for gradients to flow more naturally during training by adding a block's input straight to its output. In addition to helping with very deep network training, this architectural design also contributes to increased accuracy as network depth increases. Because of the use of global average pooling and the elimination of fully connected layers at the end, which reduces the number of parameters, ResNet models are computationally efficient despite their depth. Efficiency is also increased by methods such as bottleneck layers, which employ 1×1 convolution to reduce dimensionality prior to applying larger convolutional filters (Subrahmanyam *et al.*, 2025).

As shown in Figure 7, the ResNet architecture facilitates training of deeper networks and helps to resolve the

vanishing gradient issue.

ResNet50 usually reduces the spatial dimensions of the feature maps by employing global average pooling instead of completely linked layers at the end. As a result, overfitting is avoided, and the number of parameters is decreased.

4. Results and discussion

This section outlines the dataset used for our solution, the evaluation metrics, and the results obtained in evaluating its efficiency. Using the same data, we compared three different models in this study. With skin pictures, this skin cancer recognition model performs incredibly well. The network was trained over 50 epochs, each with 32 batches, starting from the beginning. Before an experiment is run, training data (usually 70%) and test data (usually 30%) are always separated. While the current study used a 70/30 train-test split as an initial assessment method, stronger validation strategies are critical in medical AI applications. In future work, k-fold cross-validation and repeated runs of the experiments should be incorporated to validate the robustness and generalizability of our methodology across clinical datasets on other patients.

4.1. Experimental metrics

An experimental environment was built on a machine with a Windows 11 64-bit operating system, an Intel(R) Core™ i5-1135G7 CPU @ 2.40 GHz, and 8 GB of RAM. Our proposed solution was implemented in Python 3.8.8 using external libraries and modules such as Keras, TensorFlow,




Figure 7. Basic ResNet architecture

NumPy, Pandas, OS, CV2, mysql.connector, and sklearn. For creating and implementing machine learning models, TensorFlow offers an extensive and adaptable ecosystem of tools and libraries, whereas Keras has an intuitive interface for creating neural networks. XAMPP Server was used for server deployment, and MySQL was used as the database.

4.2. Evaluation metrics

Accuracy, precision, specificity, sensitivity, and recall are the key metrics for evaluating the classification algorithms used in our skin cancer detection problem. Accuracy is the proportion of actual outcomes among the total evaluated events; precision is the proportion of accurately predicted positives; and recall is the proportion of true positives relative to all log anomalies. Specificity measures a model's ability to correctly identify negative instances (true negatives) among all actual negative instances.

The performance based on these prediction values was evaluated not only for accuracy but also for other measures, including precision, recall, sensitivity, and specificity (Equations 1–4), as well as a confusion matrix. These metrics help better understand how well a model performs in medical image classification. In the future, more precise evaluation metrics, including the receiver operating characteristic curve, the F1-score, class-wise precision–recall plots for each class, and melanoma sensitivity analysis, should be implemented to improve clinical validity and robustness assessment.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN} \quad (1)$$

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

$$Specificity = \frac{TN}{TN + FP} \quad (4)$$

A true positive is one that is predicted to be anomalous, a false positive is one that is predicted to be normal, a false negative is one that is falsely predicted to be normal, and a true negative is one that is predicted to be normal. Figures 8–10 show the confusion metrics for the CNN, MobileNet, and ResNet models.

Figure 11 shows a graphical representation of the process

for training and evaluating the hybrid model network's precision/loss. After a specific number of epochs, the test's correctness and loss remain almost constant, whereas the validation's reliability and loss fluctuate a lot initially for a few epochs before stabilizing at a fixed value.

The graphical representation of the training process and the evaluation of correctness/loss for the CNN model is shown in Figure 12. The evaluation accuracy changes significantly as the number of epochs increases, while the validation loss remains constant. In contrast, the test accuracy and loss are nearly comparable as the number of epochs rises.

Figure 13 shows the training and validation accuracy/loss for the MobileNet model. The test reliability and loss, however, nearly equalize as the number of epochs grows.

Figure 14 shows the training procedure and the evaluation accuracy/loss for the ResNet model. While the validation loss remains constant, the evaluation accuracy changes dramatically as the number of epochs increases. As the number of epochs increases, however, test reliability and loss approach equality.

Slight variations between training and validation performance during testing suggested the potential for overfitting. Preprocessing involved applying data augmentation strategies to reduce overfitting. Dropout layers, batch normalization, early stopping mechanisms, and more robust regularization techniques to enhance model generalization on unseen clinical data are potential future enhancements.

Figure 15 compares the performance of various models; we can deduce that the hybrid model, when paired with MobileNet and LSTM, offers the best performance (Table 2).

4.3. Ethical considerations

This AI-based skin cancer detection system is intended to be used only as an initial screening and decision-support tool. There are many ethical concerns that need to be resolved before clinical implementation in real-world settings. These include protecting patient privacy, obtaining informed consent when using medical images, ensuring fairness in datasets across different skin tones and ethnicities, and transparency in AI decision-making. If datasets are only from a specific population, models may be biased in their predictions on other underserved populations. Thus, future work should focus on augmenting the diversity of datasets used to train our models; integrating and improving explainable AI techniques; expanding fairness-




Figure 8. Confusion matrix for the convolutional neural network model




Figure 9. Confusion matrix for the MobileNet model




Figure 10. Confusion matrix for the ResNet Model




Figure 11. Hybrid model training and validation accuracy/loss. (A) Accuracy. (B) Loss.




Figure 12. Convolutional neural network model training and validation accuracy/loss. (A) Accuracy. (B) Loss.




Figure 13. MobileNet model training and validation accuracy/loss. (A) Accuracy. (B) Loss.




Figure 14. ResNet model training and validation accuracy/loss. (A) Accuracy. (B) Loss.




Figure 15. Bar graph showing the performance of the models
Abbreviation: CNN: Convolutional neural network.

Table 2. Comparison of deep learning models

Models	Accuracy	Precision	Recall	Specificity	Sensitivity
Hybrid model (MobileNet + long short-term memory)	0.9388	0.9031	0.9181	0.998	0.9973
Convolutional neural network	0.9294	0.9091	0.9266	0.928	0.9746
MobileNet	0.9156	0.8971	0.9191	0.9915	0.9576
ResNet	0.9071	0.8923	0.8091	0.9364	0.9661

aware learning checkpoints to improve adoption without bias; and collaborating with medical expert(s) to ensure responsible deployment of such systems.

5. Conclusion

We initially concentrated on choosing appropriate algorithms and setting up the process for our model in the skin cancer classification using deep learning methods (CNN) project. Our objective was to create a reliable system that can recognize skin photos with high accuracy, categorize them using training data, and produce accurate results. Using deep learning techniques on these skin lesions makes it simple to diagnose skin cancer. Deep learning models have been trained using the skin image collection. The results of the experiment demonstrate that, given the dataset we used, MobileNet, in conjunction with the LSTM model, achieved the highest accuracy (93%) among the models. The medical field can find substantial uses for the system. The suggested method can be used to determine whether skin cancer is present. It seems very promising to classify skin lesions using deep learning and other sophisticated machine learning techniques. Combining several types of information could help doctors make better diagnoses. While the proposed model demonstrates promising classification performance, the system should be considered an initial screening and decision-support tool rather than a replacement for professional dermatological diagnosis. Practical application in healthcare requires clinical validation with real-world patient data and consultation with medical experts.

This extension's future focus on AI-driven dermatological diagnoses using a cross-breed MobileNet-LSTM model is extensive and promising. Improving the model's precision and generalizability is a crucial area for development. This might be achieved by adding more diverse images to the training collection, representing a wider range of skin types, conditions, and ethnicities.

Putting advanced information-augmentation techniques into practice can also improve the model's ability to handle the variability of real-world skin images. Adding additional information types, like statistics or silent historical data, to support the image-based conclusion is another possible direction. This appears to result in more accurate and customized evaluations. Future work should focus on improving classification accuracy, transparency, and privacy preservation by incorporating advanced deep learning architectures such as EfficientNet, ViT, XAI, and federated learning. In addition, broader, richer clinical datasets that improve the generalizability of our models across skin tones, age groups, and demographic groups should be incorporated.

Acknowledgments

None.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

Author contributions

Conceptualization: Anoosha Shaga, Satyanarayana Gunupusala

Formal analysis: Chittibabulu Sape, Veeramallu Bobba

Investigation: Vayunandan Kumar Konakalla, Chittibabulu Sape

Methodology: Bhupesh Kumar Namdeti, Satyanarayana Gunupusala

Validation: Satyanarayana Gunupusala, Chittibabulu Sape

Visualization: Veeramallu Bobba, Satyanarayana Gunupusala

Writing-original draft: Anoosha Shaga, Satyanarayana Gunupusala

Writing-review & editing: Bhupesh Kumar Namdeti, Veeramallu Bobba

Availability of data

The data may be made available from the corresponding author upon reasonable request for academic and research use.

References

- Alfred, N., Khelifi, F., Bouridane, A., & Seker, H. (2015). Pigment network-based skin cancer detection. In: *Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 7214–7217). IEEE.
<https://doi.org/10.1109/EMBC.2015.7320056>
- Aljanabi, M., Özok, Y. E., Rahebi, J., & Abdullah, A. S. (2018). Skin lesion segmentation method for dermoscopy images using artificial bee colony algorithm. *Symmetry*, 10(8), 347.
<https://doi.org/10.3390/sym10080347>
- Al-Masni, M. A., Al-Antari, M. A., Choi, M., Han, S., & Kim, T. (2018). Skin lesion segmentation in dermoscopy images via deep full resolution convolutional networks. *Computer Methods and Programs in Biomedicine*, 162, 221–231.
<https://doi.org/10.1016/j.cmpb.2018.05.027>
- Carcagnì, P., Leo, M., Cuna, A., Mazzeo, P. L., Spagnolo, P., Celeste, G., & Distanti, C. (2019). Classification of skin lesions by combining multilevel learnings in a DenseNet architecture. In: *Lecture Notes in Computer Science* (pp. 335–344). Springer.
https://doi.org/10.1007/978-3-030-30642-7_30
- Chittibabulu, S., Gantayat, H., Gayathri, V. V., Barik, S. K., Chunduru, M. R., & Satyanarayana, G. (2026). Deep transfer-learning framework for automated cardiovascular disease prediction from ECG images. In: *2026 7th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI)* (pp. 1019–1026). IEEE.
<https://doi.org/10.1109/ICMCSI67283.2026.11412486>
- Ech-Cherif, A., Misbhaudhin, M., & Ech-Cherif, M. (2019). Deep neural network based mobile dermoscopy application for triaging skin cancer detection. In: *2019 2nd International Conference on Computer Applications & Information Security (ICCAIS)* (pp. 1–6). IEEE.
<https://doi.org/10.1109/CAIS.2019.8769517>
- Hameed, N., Shabut, A., & Hossain, A. (2018). A computer-aided diagnosis system for classifying prominent skin lesions using machine learning. In: *Proceedings of the International Conference on Computer and Electrical Engineering* (pp. 186–191). IEEE.
<https://doi.org/10.1109/CEEC.2018.8674183>
- Hasan, M., Barman, S., Islam, S., & Reza, A. W. (2019). Skin cancer detection using convolutional neural network. In: *Proceedings of the International Conference on Computing and Information Technology* (pp. 254–258).
<https://doi.org/10.1145/3330482.3330525>
- Hoshayar, A. N., Hoshayar, A. N., Al-Jumaily, A., Hoshayar, A. N., & Hoshayar, A. N. (2014). The beneficial techniques in preprocessing step of skin cancer detection system comparing. *Procedia Computer Science*, 42, 25–31.
<https://doi.org/10.1016/j.procs.2014.11.029>
- Masood, A., & Al-Jumaily, A. A. (2013). Computer aided diagnostic support system for skin cancer: A review of techniques and algorithms. *International Journal of Biomedical Imaging*, 2013, 1–22.
<https://doi.org/10.1155/2013/323268>
- Pham, T., Luong, C., Visani, M., & Hoang, V. (2018). Deep CNN and data augmentation for skin lesion classification. In: *Lecture Notes in Computer Science* (pp. 573–582). Springer.
https://doi.org/10.1007/978-3-319-75420-8_54
- Rezaeana, N., Hossain, M. S., & Andersson, K. (2020). Detection and classification of skin cancer by using a parallel CNN model. In: *2020 IEEE International Women in Engineering (WIE) Conference on Electrical and Computer Engineering (WIECON-ECE)* (pp. 380–386). IEEE.
<https://doi.org/10.1109/WIECON-ECE52138.2020.9397987>
- Shah, S., & Ross, A. (2009). Iris segmentation using geodesic active contours. *IEEE Transactions on Information Forensics and Security*, 4(4), 824–836.
<https://doi.org/10.1109/TIFS.2009.2033225>
- Subrahmanyam, V., Sudhakar, P., Rao, B. M., Ramakrishna, V. J., Vemula, M. K., & Satyanarayana, G. (2025). An innovative approach for detecting the freshness of fruits and vegetables using transfer learning techniques. In: *2025 4th International Conference on Innovative Mechanisms for Industry Applications (ICIMIA)* (pp. 999–1005). IEEE.
<https://doi.org/10.1109/ICIMIA67127.2025.11200855>
- Subramanian, R. R., Dintakurthi, A., Kumar, S., Reddy, K., Amara, S., & Chowdary, A. (2021). Skin cancer classification using convolutional neural networks. In: *2021 International Conference on Confluence* (pp. 13–19). IEEE.
<https://doi.org/10.1109/Confluence51648.2021.9377155>
- Vijayalakshmi, M. M. (2019). Melanoma skin cancer detection using image processing and machine learning. *International Journal of Trend in Scientific Research and Development*, 3(4), 780–784.
<https://doi.org/10.31142/IJTSRD23936>

Waheed, Z., Waheed, A., Zafar, M., & Riaz, F. (2017). An efficient machine learning approach for the detection of melanoma using dermoscopic images. In: *Proceedings of the International Conference on Computing, Communication and Coding Technologies (C-CODE)* (pp. 316–319). IEEE.

<https://doi.org/10.1109/C-CODE.2017.7918949>

Zhang, N., Cai, Y., Wang, Y., Tian, Y., Wang, X., & Badami, B. (2020). Skin cancer diagnosis based on optimized convolutional neural network. *Artificial Intelligence in Medicine*, 102, 101756.

<https://doi.org/10.1016/j.artmed.2019.101756>