

Smart Dermatology: Machine Learning Approaches for Skin Disease Identification

Mayur Srivastava¹, Prateek Saxena², Suman Kumar Mishra³ and Shan-e-Fatima⁴

¹Department of Computer Science and Engineering Khwaja Moinuddin Chishti Language University Tovar B 3 901 Sahara Siti Homs, IIM Road, Lucknow, Uttar Pradesh 226013

mayur851986@gmail.com

²Department of Computer Science and Engineering Khwaja Moinuddin Chishti Language University Tovar B 3 901 Sahara Siti Homs, IIM Road, Lucknow, Uttar Pradesh 226013

³Department of Computer Science and Engineering Khwaja Moinuddin Chishti Language University Tovar B 3 901 Sahara Siti Homs, IIM Road, Lucknow, Uttar Pradesh 226013

⁴Department of Computer Science and Engineering Khwaja Moinuddin Chishti Language University Tovar B 3 901 Sahara Siti Homs, IIM Road, Lucknow, Uttar Pradesh 226013

ARTICLE INFO

ABSTRACT

Received: 30 Dec 2024

Revised: 05 Feb 2025

Accepted: 25 Feb 2025

Skin conditions are among the most common health issues, affecting people for various reasons—such as bacterial or viral infections, allergies, or fungal growths. While laser and photonics-based technologies have made diagnosis faster and more precise, these methods are still costly and not accessible to everyone. That's where image processing and deep learning step in. By using just, a digital image of the affected skin area, along with a computer and a camera, we can build a low-cost, efficient dermatology screening system. Our method focuses on using pre-trained Convolutional Neural Networks (CNNs) to extract features after resizing the image. This eliminates the need for expensive equipment and speeds up diagnosis. Traditional diagnostic approaches like visual inspection or biopsy are not only time-consuming but can sometimes be inaccurate. We explore machine learning (ML) techniques like SVMs and deep learning models such as ResNet, EfficientNet, and MobileNet using benchmark datasets like ISIC. By applying preprocessing techniques like augmentation and feature extraction, the system's accuracy improves significantly. Experimental results show that CNN-based models consistently outperform older ML approaches in detecting melanoma. This demonstrates how AI-driven solutions can support dermatologists in making faster and more accurate diagnoses, ultimately leading to improved patient outcomes.

Keywords: Skin Disease; Convolutional Neural Networks (CNNs); Support Vector Machine; Dermatoscopic Image Analysis; Skin Cancer Classification; Image Preprocessing.

1. INTRODUCTION

The skin is the largest organ in the body, it shields us from the environment, helping regulate temperature, and even giving clues about our overall health. However, skin issues, ranging from mild rashes to chronic conditions like eczema, psoriasis, and even skin cancer, can significantly affect a person's well-being. Traditionally, dermatology has relied on visual assessments, physical examinations, and biopsies, which often involve lengthy processes and subjective interpretations. Fortunately, technological advancements are now transforming this field. The introduction of Smart Dermatology [18] is changing the way skin conditions are detected, monitored, and treated, making dermatological care more precise, accessible, and efficient. Smart Dermatology is all about using Computer vision, machine learning, and artificial intelligence, and wearable tech to enhance skin care. By combining these technologies, doctors can diagnose skin conditions faster and more accurately, create treatment plans tailored to each person, and even monitor skin health in real time. Unlike conventional methods that depend heavily on clinical expertise, Smart Dermatology leverages data driven insights to enhance diagnostic accuracy. Systems that are powered by AI can go through huge collections of skin images, spotting patterns and catching diseases that doctors might overlook with the naked eye. This technological shift is helping dermatologists make better-informed decisions while also offering patients quicker and more reliable results. The field of Smart Dermatology relies heavily on

cutting-edge technologies that make skin analysis faster and more precise. Some of the most impactful among these technologies are machine learning and artificial intelligence along with computer vision and image recognition, are making a big impact in dermatology. Teledermatology platforms are also becoming more popular, allowing people to get skin conditions diagnosed remotely with the help of advanced technology. One of the major benefits of Smart Dermatology is early disease detection, particularly for skin cancer. Conditions like melanoma, which can be life-threatening if not caught early, often start with subtle visual changes that are easy to miss [1]. AI-powered diagnostic tools, however, can detect the smallest irregularities in skin lesions, making early diagnosis far more achievable. This improves the chances of successful treatment and significantly increases survival rates. Additionally, Smart Dermatology promotes personalized skincare. Smart Dermatology enhances chronic skin condition management through wearable devices and AI-driven diagnostics, enabling real-time monitoring and reducing clinic visits. Teledermatology platforms improve accessibility, especially in underserved regions, making expert consultations more inclusive. However, challenges like data privacy and AI accuracy must be addressed to ensure reliability. While AI enhances efficiency, it should complement rather than replace human expertise. With advancements in deep learning and wearable technology, Smart Dermatology is set to revolutionize skin care, making it faster, more accurate, and patient centric. This study advances the field of Smart Dermatology by leveraging AI-driven diagnostic models for improved skin disease detection.

Our research Contributions are as follows: -

- We utilized a pretrained skin disease model from Kaggle, which was already trained on a comprehensive dataset, capturing essential characteristics through its optimized weights and biases.
- To make the diagnosis more accurate, on top of the pretrained model, we applied two machine learning techniques: Convolutional Neural Networks (CNN) and Support Vector Machines (SVM).
- By looking at how well both models work, we want to figure out which one does a better job at classifying skin diseases.
- Additionally, our study evaluates the integration of AI-based diagnostics with wearable technology and teledermatology platforms, addressing challenges such as model reliability, data privacy, and accessibility.
- With this research, we're working toward building a Smart Dermatology system that's not just more accurate but also secure and scalable.




Figure1. Labelled Diagram of Human Skin

The epidermis, dermis, and subcutis are the three main layers of human skin are depicted in the image shown in Fig.1. Skin color comes from cells called melanocytes, which, along with basal and squamous cells, are located in the outer layer of the skin. Beneath this outer layer, in the dermis, you'll find blood vessels, nerves, hair follicles, and sweat glands. Insulation and cushioning are provided by the subcutis, which is made up of connective tissue and fat. The main cellular elements of the epidermis are highlighted in a portion that has been zoomed in.

2. LITERATURE REVIEW

Many researchers have looked into skin disease detection using a variety of methods. They've applied different computer vision techniques like image processing, segmentation, and feature extraction to classify images. These

approaches also show up in numerous other studies. In [1], a two-phase approach that combines machine learning and computer vision achieved 95% accuracy for six skin diseases but was limited by a small dataset and potential bias. In [3], GoogleNet Inception v3 with transfer learning showed high accuracy, but relied only on image data and had small test sets, raising bias concerns. In [2], a Triple Net DCNN achieved 96.6% accuracy for melanoma detection but faced computational complexity and a narrow dataset scope. In [4], VGG16 variants classified skin lesions with 78.66% sensitivity, but struggled with overfitting and non-medical pretraining data. In [5], a CNN performed well for melanoma detection, but the small dataset raised overfitting concerns. In [6], a CNN with a Softmax classifier reached 70% accuracy but needed more data and lacked cross-validation. In [7], deep learning models like DenseNet 201 and ResNet 152 performed impressively but faced challenges with unbalanced data and rare cancers. In [9], a Multi-Class Multi Level classification algorithm achieved 96.47% accuracy but was hindered by a small, inconsistent dataset. In [8], a pretrained AlexNet with SVM reached 100% detection accuracy for three diseases but faced challenges with the small dataset and potential bias from internet images. In [11], different CNNs effectively classified face skin diseases, but dataset limitations and overfitting risks were highlighted. In [10], a hybrid approach combining image processing and CNN showed 85% accuracy for melanoma detection but focused only on melanoma and had a small dataset. In [12], ResNet-101 and Inception-v3 detected skin cancer with 84.09% and 87.42% accuracy, but lacked cross-validation and had possible class distribution bias. In [13], a deep learning CNN model detected skin cancer with high accuracy but faced dataset limitations and potential bias. In [14], Melanoma detection was enhanced using deep learning and traditional machine learning techniques. but required large datasets and had high computational costs. In [15], a CNN with transfer learning achieved dermatologist-level accuracy but struggled with rare skin conditions due to dataset limitations. In [16], a CNN achieved high accuracy for skin disease diagnosis but needed large datasets and computational resources. In [22], a review found that CNNs generally outperformed other models for skin cancer detection, but highlighted challenges with dataset diversity and heavy training requirements. In [19], a hybrid CNN and Fuzzy K-Means model classified skin cancer accurately but struggled with dataset size and misclassification in difficult cases. In [18], a review of ML algorithms found that CNNs outperformed dermatologists in skin cancer diagnosis, though concerns remained about real-world testing and dataset limitations. In [17], a systematic review concluded deep learning outperformed traditional ML for skin lesion classification but identified challenges like small datasets and racial bias. In [21], a CNN with data augmentation achieved 95.18% accuracy on skin lesions but faced overfitting and was limited to 7 classes. In [20], a dual-stage approach using ESRGAN and deep learning achieved 85.8% accuracy for skin cancer classification but focused on binary classification and depended on preprocessing. In [23], a custom CNN achieved 88.83% accuracy for melanoma subtypes but was tested only on a small dataset without cross-dataset validation. In [24], a CNN with ResNet-50 achieved 87.42% accuracy for skin disease detection but showed room for improvement in class separability and did not address.

3. MOTIVATION

Skin problems are a lot more common than most people think, and getting the right diagnosis quickly can make a huge difference in how they're treated. However, visiting a dermatologist can be a challenge, particularly for people who reside in isolated or rural locations with no access to medical treatment. As a result, many either ignore early symptoms or rely on random information they find online, which can lead to delayed treatment or unnecessary worry. The idea behind Smart Dermatology is simple—it's about using technology to bridge this gap. With the help of AI and advanced skin detection tools, people can get a better understanding of their skin issues and know when it's time to see a doctor. This isn't meant to replace dermatologists but to make skin health more accessible to everyone, no matter where they live. We believe that combining machine learning with image analysis has the power to make skin condition detection faster and more accurate. This can benefit both individuals and healthcare providers by speeding up diagnoses and improving their precision. The goal is straightforward: to make early detection easier, bring skin care closer to people, and ultimately help more individuals take better care of their skin.

4. PROPOSED SYSTEM MODEL

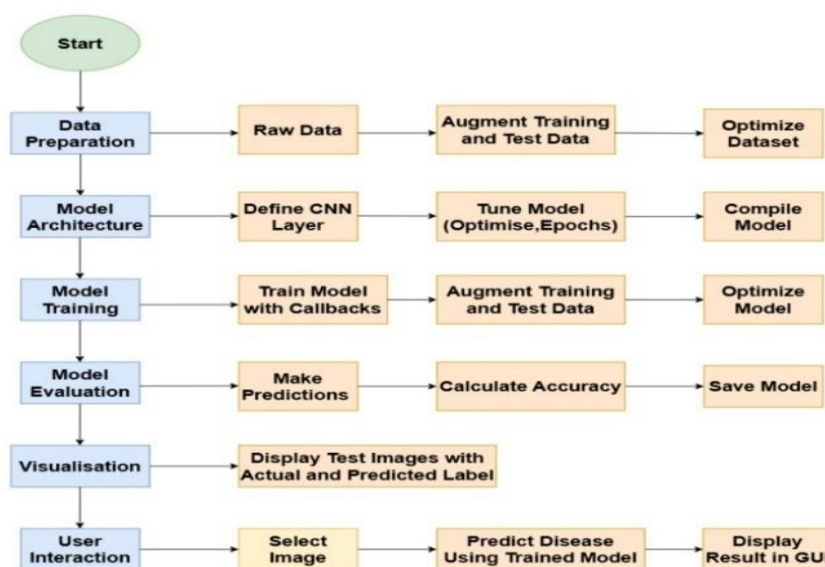


Figure.2. Flow chart for Skin Disease Prediction Model

The Model Architecture step follows, during which CNN layers are designed, hyper parameters such as epochs and optimization techniques are adjusted, and the model is put together. During the Model Training phase, which also includes data augmentation and Optimization, the model is then trained via callbacks and then saved, predictions are made, and accuracy is computed during the Model Evaluation step. During the Visualization phase, test images are displayed along with their actual and expected labels. Finally, in the User Interaction phase, users can select an image, forecast the disease using the trained model, and view the results in a graphical user interface (GUI). The flowchart shown in fig.2 systematically illustrates the steps involved in developing a CNN-based disease prediction system.

4.1 About the dataset

For classifying the different type of skin diseases [12], we use the ISIC Skin Cancer dataset with training (2239 images) and testing (118 images) sets across 9 classes (e.g., melanoma, nevus, and basal cell carcinoma).

	Class	No. of Image
0	squamous cell carcinoma	181
1	dermatofibroma	95
2	actinic keratosis	114
3	pigmented benign keratosis	462
4	nevus	357
5	vascular lesion	139
6	seborrheic keratosis	77
7	basal cell carcinoma	376
8	melanoma	438

Figure 3. Disease Classes and its no. of images

Fig. 3 gives an overview of how many images are available for each category of skin disease in the dataset, highlighting a significant class imbalance [24]. While pigmented benign keratosis (462 images) and melanoma (438 images) have the most images, seborrheic keratosis (77 images) and dermatofibroma (95 images) have the fewest, which could lead to performance issues due to insufficient training data. This imbalance may cause the model to favour majority classes, leading to poor classification accuracy for underrepresented conditions. To address this, data augmentation techniques like rotation, flipping, and synthetic image generation can be applied to increase diversity in minority classes. Additionally, class balancing through oversampling underrepresented categories or under sampling over

represented ones can help distribute training data more evenly. Implementing a weighted loss function during model training can further mitigate bias, ensuring better recognition of rare classes. Addressing these issues will improve model performance and enhance its ability to classify all skin conditions more accurately.

4.2 Main Approach

We utilized a pretrained skin disease model provided by Kaggle to train our prototype. This pretrained prototype had already been trained on a dataset, with its weights and biases capturing the dataset's essential characteristics. In the pretrained model, we first apply SVM, followed by CNN, and then compare the efficiency of both models to determine which performs better—the SVM-trained model or the CNN trained model.

4.2.1 SVM Approach

During preprocessing, images were resized to 180x180 pixels using the Image Data Generator from TensorFlow. No explicit augmentation was applied during training, but noise with a factor of 1.0 was added to the test features to assess the model's robustness [19]. The model architecture included dividing the training data, using 80% for training the model and the remaining 20% for validation. This was done to fine-tune the SVM. Once the SVM was trained, the final model was saved as `svm_model.pkl` for future reuse. TensorFlow and Keras were used for loading the pre-trained model and extracting features, while NumPy facilitated numerical operations on feature arrays. Scikit-learn was employed for implementing the SVM (using `sklearn.svm.SVC`), generating classification reports, and computing ROC curves. Matplotlib was used for visualizing the ROC curves and other plots, and Pickle enabled saving and loading the trained SVM model.




Figure 4. Confusion matrix – SVM model

Figure 4 displays the SVM model's confusion matrix. highlights some scattered predictions and shows an accuracy of 63%, suggesting the model performs moderately well but has noticeable struggles with certain classes. A confusion matrix is essentially a table that helps assess how well a classification model is doing by comparing the actual labels with the predicted ones. The rows of the table represent the true classes, and the columns represent what the model predicted. Each cell in the matrix shows how many instances ended up in a particular category, giving us an idea of how often a true label was matched to the correct predicted label. The diagonal cells, running from the top-left corner to the bottom-right, represent the instances that were correctly classified (true positives), giving us a clear picture of how accurate the model is for each class. Off-diagonal elements indicate misclassifications, where the model either falsely predicted a different class (false positives) or failed to recognize the correct class (false negatives). The color intensity in the matrix visually highlights classification frequencies, with darker shades representing higher counts, making it easier to identify patterns and areas where the model struggles. The classification report gives a clear overview of how well a machine learning model performs on test data. It looks at important metrics like precision, recall, F1-score, and support for each individual class to give a complete picture of the model's effectiveness.

Classification Report on Test Data:				
	precision	recall	f1-score	support
actinic keratosis	1.00	0.42	0.59	12
basal cell carcinoma	0.33	0.86	0.48	14
dermatofibroma	0.80	0.36	0.50	11
melanoma	0.56	0.64	0.60	14
nevus	0.85	0.85	0.85	13
pigmented benign keratosis	0.58	0.69	0.63	16
seborrheic keratosis	1.00	0.20	0.33	10
squamous cell carcinoma	1.00	0.58	0.74	12
vascular lesion	0.92	0.85	0.88	13
accuracy			0.63	115
macro avg	0.78	0.60	0.62	115
weighted avg	0.76	0.63	0.63	115

Figure 5. Classification report depicting performance of SVM approach

In Figure 5, Precision indicates the proportion of projected positives that were true, so higher values mean fewer false positives. Remember, however that, indicates how many of the actual positives were correctly identified, with higher recall meaning fewer false negatives. The F1-score is a balance between precision and recall, giving a more overall picture of how well the model is performing. Support refers to how many true instances there are for each class. The model does particularly well with certain classes, like vascular lesions, where it has a perfect precision of 1.00 and a solid recall of 0.85. Nevus also performs well, with the highest recall at 0.92, meaning the model did a great job of correctly identifying most of the actual nevus cases. However, it struggles with seborrheic keratosis, showing a precision, recall, and F1-score of 0.00, indicating complete misclassification. Basal cell carcinoma has high recall (0.80) but low precision (0.26), meaning it is frequently predicted but often misclassified. The overall accuracy of the model is 53%, suggesting it correctly classifies a little more than half of the instances. The macro average provides an equal mean of precision, recall, and F1-score for all classes, without considering how many instances there are in each class. On the other hand, the weighted average takes into account the imbalances between the classes by giving more importance to those that are larger. Despite some strong predictions, the model requires improvement, particularly for underperforming classes, possibly through better feature extraction, class balancing, or hyper parameter tuning.




Figure 6. ROC Curve for SVM Classifier Showing Model Performance

Fig. 6 shows the Receiver Operating Characteristic (ROC) curve for a multi-class classification problem using the one-vs-rest approach. The ROC curve contrasts the false positive rate at various thresholds with the genuine positive rate (sensitivity), which helps us assess how well the model is doing for each class. The Area Under the Curve (AUC) values give us an idea of how effectively the model can separate the classes, with higher AUC values (closer to 1) meaning better performance. Looking at the plot, seborrheic keratosis (AUC = 0.99) and nevus (AUC = 0.98) have

the highest AUC scores, indicating that these classes are accurately classified with few false positives. On the other hand, basal cell carcinoma (AUC = 0.83) and squamous cell carcinoma (AUC = 0.88) have lower AUC values, suggesting that the model has more trouble with these categories. Overall, the AUC values are strong, showing that the model does a good job distinguishing between most skin diseases, although there's room for improvement in the classes with lower scores. The model demonstrates variable performance across different classes, with some categories achieving high classification accuracy while others struggle. Melanoma detection is moderate, but the overall accuracy remains relatively low at 53%, indicating room for improvement. Certain classes, such as seborrheic keratosis, perform particularly poorly, likely due to limited training data (10%) and class imbalance. The model's ability to distinguish between similar skin conditions is inconsistent, affecting its reliability. Addressing these issues through data augmentation, class balancing, and model optimization could help enhance performance across all categories.

4.2.2 CNN Approach

The CNN Approach Preprocessing step involves using `image_dataset_from_directory` to load images from the directory, organize them into batches, resize them to (180, 180) pixels, and assign categorical labels (one-hot encoding). Data Augmentation with Augmentor includes Class Balancing, where 500 samples per class are generated to mitigate sparsity in underrepresented classes, ensuring a balanced dataset for accurate multi-class classification. Additionally, Improved Generalization is achieved through random rotations that simulate natural variations in skin lesion images (e.g., different angles of capture), this helps improve the model's strength and its ability to perform well on new, unseen test data.

Model: "sequential"

Layer (type)	Output Shape	Param #
rescaling (Rescaling)	(None, 180, 180, 3)	0
conv2d (Conv2D)	(None, 178, 178, 32)	896
max_pooling2d (MaxPooling2D)	(None, 89, 89, 32)	0
conv2d_1 (Conv2D)	(None, 87, 87, 64)	18496
max_pooling2d_1 (MaxPooling2D)	(None, 43, 43, 64)	0
conv2d_2 (Conv2D)	(None, 41, 41, 128)	73856
max_pooling2d_2 (MaxPooling2D)	(None, 20, 20, 128)	0
dropout (Dropout)	(None, 20, 20, 128)	0
flatten (Flatten)	(None, 51200)	0
dense (Dense)	(None, 128)	6553728
dropout_1 (Dropout)	(None, 128)	0
dense_1 (Dense)	(None, 9)	1161
Total params: 6,648,137		
Trainable params: 6,648,137		
Non-trainable params: 0		

Figure 7. CNN Sequential Model Summary

```
Epoch 17/25
169/169 [=====] - ETA: 0s - loss: 0.4253 - accuracy: 0.8422
Epoch 17: val_accuracy did not improve from 0.82480
169/169 [=====] - 104s 618ms/step - loss: 0.4253 - accuracy: 0.8422 - val_loss: 0.5676 - val_accuracy: 0.7921
Epoch 18/25
169/169 [=====] - ETA: 0s - loss: 0.4318 - accuracy: 0.8414
Epoch 18: val_accuracy did not improve from 0.82480
169/169 [=====] - 111s 659ms/step - loss: 0.4318 - accuracy: 0.8414 - val_loss: 0.5809 - val_accuracy: 0.7818
Epoch 19/25
169/169 [=====] - ETA: 0s - loss: 0.3620 - accuracy: 0.8587
Epoch 19: val_accuracy improved from 0.82480 to 0.82702, saving model to model.h5
169/169 [=====] - 105s 624ms/step - loss: 0.3620 - accuracy: 0.8587 - val_loss: 0.5403 - val_accuracy: 0.8270
Epoch 20/25
169/169 [=====] - ETA: 0s - loss: 0.3402 - accuracy: 0.8681
Epoch 20: val_accuracy improved from 0.82702 to 0.84410, saving model to model.h5
169/169 [=====] - 110s 652ms/step - loss: 0.3402 - accuracy: 0.8681 - val_loss: 0.4784 - val_accuracy: 0.8441
Epoch 21/25
169/169 [=====] - ETA: 0s - loss: 0.3195 - accuracy: 0.8811
Epoch 21: val_accuracy did not improve from 0.84410
169/169 [=====] - 108s 639ms/step - loss: 0.3195 - accuracy: 0.8811 - val_loss: 0.5756 - val_accuracy: 0.8003
Epoch 22/25
169/169 [=====] - ETA: 0s - loss: 0.3258 - accuracy: 0.8741
Epoch 22: val_accuracy improved from 0.84410 to 0.84781, saving model to model.h5
169/169 [=====] - 152s 983ms/step - loss: 0.3258 - accuracy: 0.8741 - val_loss: 0.4786 - val_accuracy: 0.8478
Epoch 23/25
169/169 [=====] - ETA: 0s - loss: 0.2737 - accuracy: 0.8974
Epoch 23: val_accuracy did not improve from 0.84781
169/169 [=====] - 137s 812ms/step - loss: 0.2737 - accuracy: 0.8974 - val_loss: 0.4913 - val_accuracy: 0.8456
Epoch 24/25
169/169 [=====] - ETA: 0s - loss: 0.2999 - accuracy: 0.8837
Epoch 24: val_accuracy did not improve from 0.84781
169/169 [=====] - 156s 927ms/step - loss: 0.2999 - accuracy: 0.8837 - val_loss: 0.5071 - val_accuracy: 0.8411
Epoch 25/25
169/169 [=====] - ETA: 0s - loss: 0.2992 - accuracy: 0.8904
Epoch 25: val_accuracy did not improve from 0.84781
169/169 [=====] - 125s 741ms/step - loss: 0.2992 - accuracy: 0.8904 - val_loss: 0.4811 - val_accuracy: 0.8463
```

Figure 8. Training Progress over Epochs in CNN Model

The Model Architecture consists of a CNN structure, built as a sequential model incorporating layers such as Conv2D, MaxPooling2D, Dense, and Dropout for regularization [15], [21]. The Adam optimizer helps speed up the training process, making it more efficient. At the same time, tools like Model Checkpoint and Early Stopping are used to improve the overall training and prevent the model from overfitting, ensuring it generalizes well. The model is trained for upto 25 epochs as shown in Fig.8 using a training dataset (train_ds) and validated on a separate dataset (val_ds). The fit method incorporates Model Checkpoint and Early Stopping callbacks to ensure an efficient and controlled training process. The Model Checkpoint saves the model weights to "model.h5" whenever the validation accuracy gets better. With save_best_only=True, it ensures that only the best version of the model is saved. Early Stopping keeps an eye on the validation accuracy and stops the training if there's no improvement after 5 epochs (patience=5). This helps avoid overfitting and also saves on computational resources by not running unnecessary training cycles. The provided image consists of two plots: model accuracy (left) and model loss (right), both plotted over training epochs. In the left plot, the training accuracy is represented by the blue curve, and the validation accuracy by the orange curve. The training loss is shown by the blue curve on the right plot, and the validation loss is shown by the orange curve.





Figure 9. ROC Curve - SVM Classifier Showing Model Performance

The growing accuracy and declining loss show that the model is learning efficiently. There is minimal overfitting, as validation accuracy and loss do not diverge significantly, suggesting a well-tuned model. Additionally, the Model Checkpoint and Early Stopping callbacks are likely contributing to optimal performance by preventing excessive training and ensuring the best model is retained. The AUC values range between 0.91 and 0.99, indicating strong classification performance for all classes. Higher AUC values, which are closer to 1.0, indicate that the model is better at telling apart positive and negative cases. In this case, seborrheic keratosis stands out as the best classified category with an AUC of 0.99, followed by melanoma at 0.96 and vascular lesions at 0.97, both of which also show a strong ability to differentiate between the two classes. Even the lowest AUC (0.91 for nevus) still demonstrates good classification capability. The ROC curves, shown in Fig. 9, are positioned near the top-left corner, which indicates that the model has high sensitivity (True Positive Rate) and a low False Positive Rate across all the classes. Additionally, as all curves lie well above the dashed diagonal line representing a random classifier (AUC = 0.5), the model's effectiveness in distinguishing between different skin disease categories is confirmed.

**Figure 10.** CNN Model's confusion matrix

As illustrated in Fig. 10, the model exhibits impressive classification accuracy, particularly in identifying certain skin conditions like melanoma, pigmented benign keratosis, squamous cell carcinoma, and vascular lesions. The highest values appearing along the diagonal of the confusion matrix highlight how effectively the model was able to predict these classes correctly. For instance, melanoma was correctly predicted 13 times, while pigmented benign keratosis, squamous cell carcinoma, and vascular lesions each had 12 accurate predictions. Other categories such as nevus (11 correct), dermatofibroma (10 correct), and actinic keratosis (10 correct) also showed promising results. However, there were a few notable misclassifications—for example, dermatofibroma was misclassified as actinic keratosis in three instances, and nevus was incorrectly labelled as melanoma three times. These errors suggest that there might be similarities in visual features between certain classes, which could lead to occasional confusion for the model. Still, the minimal number of misclassifications in categories like seborrheic keratosis, squamous cell carcinoma, and vascular lesions adds confidence to the overall accuracy and reliability of the model. The detailed classification report further supports these findings by using precision, recall, and F1-score to measure performance. In summary, despite some overlapping features among a few classes, the model's overall effectiveness and consistent accuracy across multiple skin disease categories underscore its strong potential for aiding dermatological diagnosis.

Classification Report:

	precision	recall	f1-score	support
actinic keratosis	0.82	0.75	0.78	12
basal cell carcinoma	0.85	0.79	0.81	14
dermatofibroma	0.71	0.91	0.80	11
melanoma	0.80	0.86	0.83	14
nevus	0.71	0.92	0.80	13
pigmented benign keratosis	0.85	0.69	0.76	16
seborrheic keratosis	0.80	0.80	0.80	10
squamous cell carcinoma	1.00	0.75	0.86	12
vascular lesion	0.92	0.92	0.92	13
accuracy			0.82	115
macro avg	0.83	0.82	0.82	115
weighted avg	0.83	0.82	0.82	115

Figure 11. Classification report depicting performance of CNN approach

5. RESULTS AND DISCUSSION

Compared to the traditional biopsy method, the diagnosing system we've proposed is much more efficient. It reduces both the cost and time required to detect melanoma. This study employs image processing and machine learning methods to identify skin cancer. We used the SVM technique to evaluate the skin lesion identification accuracy. The SVM algorithm helps improve the effectiveness and precision of melanoma detection. Area, perimeter, abnormality, irregularity, circularity, and diameter are among the shape features that are first extracted from the photos in the suggested system.




Figure 12. Comparative Analysis for accuracy prediction of two Models

The bar chart in figure 12 compares the accuracy of two machine learning models: CNN (Convolutional Neural Network) and SVM (Support Vector Machine). The models are shown on the x-axis, while accuracy values between 0 and 1 are shown on the y-axis. The CNN model, which is depicted by an orange bar, produced a far greater accuracy of 0.8000 than the SVM model, which is represented by a blue bar and achieved an accuracy of roughly 0.6314. This suggests that the CNN outperformed the SVM in terms of accuracy, likely due to its stronger ability to detect complex patterns, especially in tasks that involve image data. The bar chart compares the per-class F1-scores of two machine learning models—the Convolutional Neural Network (CNN) and the Support Vector Machine (SVM)—across various skin disease categories. The models are shown on the x axis, while accuracy values between 0 and 1 are shown on the y-axis. The CNN model, which is depicted by an orange bar, produced a far greater accuracy of 0.8000 than the SVM model, which is represented by a blue bar and achieved an accuracy of roughly 0.6314. This implies that the accuracy in CNN is far better than the SVM, most likely because of its enhanced capacity to recognize intricate patterns, particularly in jobs using picture data. Especially in classifying skin disease categories with higher F1-scores, reinforcing its superiority in handling complex patterns compared to SVM. F1 score comparison for the two models is indicated in Figure 13. The Y-axis represents the F1 score, which ranges from 0 to 1, while the X-axis displays the different types of skin diseases. Orange bars indicate the SVM model, and blue bars reflect the CNN model. CNN routinely achieves higher F1 scores than SVM in the majority of classes. CNN has made notable progress in classifications like "basal cell carcinoma," "seborrheic keratosis," and "squamous cell carcinoma." On the other hand, SVM does similarly well in a select instance, like "nevus." Overall, the trend indicates that CNN performs better than SVM in handling complicated patterns, as evidenced by its higher F1-scores for classifying skin disease categories. The study compares the performance of CNN (Convolutional Neural Network) and SVM (Support Vector Machine) across different skin disease categories. Along with the overall accuracy, it also looks at the macro and weighted average scores.




Figure 13. F1 score of skin diseases for two Models

Each class's precision, recall, F1 score, and support are included in the reports giving a detailed breakdown of how well each model performs. In every evaluation metric, the CNN model performs noticeably better than the SVM model. SVM only attains 63% overall accuracy, but CNN attains 82%. Additionally, CNN's weighted average (0.82) and macro average (0.62) are much higher than SVM's (0.62 and 0.63, respectively). This shows that CNN consistently performs well across all classes, but SVM performs poorly, even in classes with very low recall values, such as dermatofibroma and seborrheic keratosis. When looking into individual class performances, CNN continuously shows better recall and F1-scores in nearly every category, proving its superior case classification capabilities. For instance, CNN's F1-score for basal cell carcinoma is 0.81, but SVM's is only 0.48. CNN scores 0.83 compared to SVM's 0.60 for melanoma. All things considered, the comparison shows how well CNN can classify skin conditions, which makes it a better option than SVM for this task.

5.1 Disease Identified by Model

The image above shows a group of skin lesion pictures that were correctly identified by a Convolutional Neural Network (CNN) model [17]. Each picture represents a different skin condition, and both the actual diagnosis (ground truth) and the model's prediction are shown in green text, indicating that they match. The dataset includes a range of skin diseases like actinic keratosis, basal cell carcinoma, dermatofibroma, melanoma, nevus, pigmented benign keratosis, seborrheic keratosis, squamous cell carcinoma, and vascular lesions. The results show that the model performs really well in recognizing and correctly identifying these different conditions. The fact that the predicted labels match the actual ones across all nine types suggests that the CNN has effectively learned to pick up on the key patterns and features that make each disease distinct from the others. This proves that the model is quite accurate and dependable when it comes to analysing skin images. Its strong performance makes it a valuable tool for helping with automated skin disease diagnosis, which could make the process faster and more reliable.




Figure 14. Analysis of Predicted and actual outcomes

5.2 GUI for Disease Detection

The image in Fig. 15 illustrates a Skin Disease Detection System that leverages deep learning to recognize various skin conditions by analysing pictures that are uploaded. The system has a simple interface where users can upload a photo of their skin issue and get a possible diagnosis along with treatment suggestions. In this case, the system has analysed the uploaded image and predicted a specific skin condition, displaying its name along with advice on what to do next. The recommendations include basic care tips, suggestions on whether to keep an eye on the condition, and guidance on when it might be necessary to see a doctor. This kind of automated system can be really helpful for early detection and quick screening, giving people an idea of what their skin issue might be before they decide to visit a dermatologist. It doesn't replace professional medical advice, but it can definitely help in making more informed decisions about skin health.




Figure 15. Smart Dermatology Model

6. FUTURE SCOPE

The future of smart dermatology looks promising, with technology advancing at an incredible pace. In the coming years, AI-driven skin analysis will become even more accurate, helping people detect skin issues earlier than ever. Wearable devices with real-time skin monitoring might become common, giving users instant feedback on their skin health. Teledermatology is also expected to grow, making skin care more accessible for people in remote areas. Instead of waiting weeks for an appointment, users could get expert advice within minutes. Personalized skincare powered by AI and genetic data might become the norm, ensuring treatments are tailored to individual needs. Moreover, smart dermatology could integrate with other health tracking systems, linking skin health to overall well-being. We might see apps that not only analyse skin but also suggest lifestyle changes based on weather, pollution levels, or diet. While challenges like data privacy and accuracy still exist, ongoing research and innovation will continue to refine smart dermatology, making it a key player in the future of healthcare.

7. CONCLUSION

Smart dermatology is definitely transforming how we take care of our skin. Thanks to advancements in AI, machine learning, and wearable devices, diagnosing and managing skin issues has become quicker and more precise. It's changing the game by making skin care more efficient and reliable. People can now get instant assessments through apps, making skin care more accessible and convenient. This technology is also helping with early detection of serious issues like skin cancer, which can make a huge difference in treatment outcomes. While smart dermatology has made impressive progress, it's still evolving. There are challenges, like ensuring the accuracy of AI diagnoses and protecting user data, but constant improvements in technology are addressing these concerns. In the future, smart dermatology will likely become even more advanced, offering more personalized care and better connections between patients and doctors. Overall, this technology has the potential to make skin care smarter, more efficient, and accessible to more people, ultimately improving both skin health and quality of life.

REFERENCES

1. Kumar, V. B., Kumar, S. S., & Saboo, V. (2016, September). Dermatological disease detection using image processing and machine learning. In *2016 third international conference on artificial intelligence and pattern recognition (AIPR)* (pp. 1–6). IEEE.
2. Ge, Z., Demyanov, S., Chakravorty, R., Bowling, A., & Garnavi, R. (2017). Skin disease recognition using deep saliency features and multimodal learning of dermoscopy and clinical images. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11–13, 2017, Proceedings, Part III 20* (pp. 250–258). Springer International Publishing.
3. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115–118.
4. Lopez, A. R., Giro-i-Nieto, X., Burdick, J., & Marques, O. (2017, February). Skin lesion classification from dermoscopic images using deep learning techniques. In *2017 13th IASTED international conference on biomedical engineering (BioMed)* (pp. 49–54). IEEE.
5. Li, Y., & Shen, L. (2018). Skin lesion analysis towards melanoma detection using deep learning network. *Sensors*, 18(2), 556.
6. Rathod, J., Waghmode, V., Sodha, A., & Bhavathankar, P. (2018, March). Diagnosis of skin diseases using Convolutional Neural Networks. In *2018 second international conference on electronics, communication and aerospace technology (ICECA)* (pp. 1048–1051). IEEE.
7. Rezvantlab, A., Safigholi, H., & Karimijeshni, S. (2018). Dermatologist level dermoscopy skin cancer classification using different deep learning convolutional neural networks algorithms. *arXiv preprint arXiv:1810.10348*.
8. ALenezi, N. S. A. (2019). A method of skin disease detection using image processing and machine learning. *Procedia Computer Science*, 163, 85–92.
9. Hameed, N., Shabut, A. M., Ghosh, M. K., & Hossain, M. A. (2020). Multi-class multi-level classification algorithm for skin lesions classification using machine learning techniques. *Expert Systems with Applications*, 141, 112961.
10. Vijayalakshmi, M. M. (2019). Melanoma skin cancer detection using image processing and machine learning. *International Journal of Trend in Scientific Research and Development (IJTSRD)*, 3(4), 780–784.
11. Wu, Z. H. E., Zhao, S., Peng, Y., He, X., Zhao, X., Huang, K., ... & Li, Y. (2019). Studies on different CNN algorithms for face skin disease classification based on clinical images. *IEEE Access*, 7, 66505–66511.
12. Demir, A., Yilmaz, F., & Kose, O. (2019, October). Early detection of skin cancer using deep learning architectures: resnet-101 and inception-v3. In *2019 medical technologies congress (TIPTEKNO)* (pp. 1–4). IEEE.
13. Jinnai, S., Yamazaki, N., Hirano, Y., Sugawara, Y., Ohe, Y., & Hamamoto, R. (2020). The development of a skin cancer classification system for pigmented skin lesions using deep learning. *Biomolecules*, 10(8), 1123.
14. Daghrir, J., Tlig, L., Bouchouicha, M., & Sayadi, M. (2020, September). Melanoma skin cancer detection using deep learning and classical machine learning techniques: A hybrid approach. In *2020 5th international conference on advanced technologies for signal and image processing (ATSIP)* (pp. 1–5). IEEE.
15. Liu, Y., Jain, A., Eng, C., Way, D. H., Lee, K., Bui, P., ... & Coz, D. (2020). A deep learning system for differential diagnosis of skin diseases. *Nature Medicine*, 26(6), 900–908.

16. Shanthi, T., Sabeenian, R. S., & Anand, R. (2020). Automatic diagnosis of skin diseases using convolution neural network. *Microprocessors and Microsystems*, 76, 103074.
17. Kassem, M. A., Hosny, K. M., Damaševičius, R., & Eltoukhy, M. M. (2021). Machine learning and deep learning methods for skin lesion classification and diagnosis: a systematic review. *Diagnostics*, 11(8), 1390.
18. Das, K., Cockerell, C. J., Patil, A., Pietkiewicz, P., Giulini, M., Grabbe, S., & Goldust, M. (2021). Machine learning and its application in skin cancer. *International Journal of Environmental Research and Public Health*, 18(24), 13409.
19. Nawaz, M., Mehmood, Z., Nazir, T., Naqvi, R. A., Rehman, A., Iqbal, M., & Saba, T. (2022). Skin cancer detection from dermoscopic images using deep learning and fuzzy k-means clustering. *Microscopy Research and Technique*, 85(1), 339–351.
20. Gouda, W., Sama, N. U., Al-Waakid, G., Humayun, M., & Jhanjhi, N. Z. (2022, June). Detection of skin cancer based on skin lesion images using deep learning. In *Healthcare* (Vol. 10, No. 7, p. 1183). MDPI.
21. Shetty, B., Fernandes, R., Rodrigues, A. P., Chengoden, R., Bhattacharya, S., & Lakshmana, K. (2022). Skin lesion classification of dermoscopic images using machine learning and convolutional neural network. *Scientific Reports*, 12(1), 18134.
22. Dildar, M., Akram, S., Irfan, M., Khan, H. U., Ramzan, M., Mahmood, A. R., ... & Mahnashi, M. H. (2021). Skin cancer detection: a review using deep learning techniques. *International Journal of Environmental Research and Public Health*, 18(10), 5479.
23. Allugunti, V. R. (2022). A machine learning model for skin disease classification using convolution neural network. *International Journal of Computing, Programming and Database Management*, 3(1), 141–147.
24. Inthiyaz, S., Altahan, B. R., Ahammad, S. H., Rajesh, V., Kalangi, R. R., Smirani, L. K., ... & Rashed, A. N. Z. (2023). Skin disease detection using deep learning. *Advances in Engineering Software*, 175, 103361.