Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 07 May 2026
Revised 25 June 2026
Accepted 07 July 2026

Natural Resources for Human Health 2026, Vol. 6 (11s): 412–424

KEYWORDS:
Skin cancer, Early detection, Artificial intelligence, Skin cancer classification, Automated screening methods, Convolutional neural networks, Renet-50.

AI-Dermis: Intelligent Skin Cancer Detection and Patient Monitoring using IoT Technology and Learning Model

Manisha Sudam Sagar[1*], Sudhir Kumar Meesala[2]

[1]Department of Computer Science & Engineering, School of Computer Science & Engineering (SOCSE), Sandip University, Nashik, Maharashtra, India
[2]Department of Computer Science & Engineering, School of Computer Science & Engineering (SOCSE), Sandip University, Nashik, Maharashtra, India
manishasagar3322@gmail.com, mskcec@gmail.com

ABSTRACT: Skin cancer presents in various forms, including squamous cell carcinoma (SCC), basal cell carcinoma (BCC), and melanoma. Established risk factors include ultraviolet (UV) radiation exposure from solar or artificial sources, lighter skin pigmentation, a history of sunburns, and a family history of the disease. Early detection and prompt intervention are crucial for achieving a favorable prognosis. Traditionally, treatment modalities include surgery, radiation therapy, and chemotherapy. Recent advancements in immunotherapy have revolutionized skin cancer diagnosis, but manual identification remains time-consuming. Artificial intelligence (AI) has shown potential in skin cancer classification, leading to automated screening methods. To support dermatologists, we improved the model for classifying images. This model is able to recognize seven different kinds of skin lesions. On the ISIC dataset, an analysis has been done. This study offers a novel approach to early skin cancer diagnosis based on image processing. Our approach leverages the high accuracy of a specific convolutional neural network architecture, utilizing transfer learning with pre-trained data to further enhance detection performance. Our findings demonstrate that the employed ResNet-50 transfer learning model achieves a remarkable accuracy of 97%, while ResNet50 without augmentation gives an accuracy of 81.57% and an F1-score of 75.75%.

1 INTRODUCTION

Skin cancer is the most common type of cancer in the world. It happens when cells in the skin grow too quickly. These malignant cells display variations in characteristics and aggressiveness, with the potential to infiltrate deeper tissues or disseminate to distant anatomical regions [1]. The skin's high sensitivity to environmental factors, especially ultraviolet (UV) radiation from sunlight, plays a major role in this widespread occurrence [2].

1.1 The importance of classifying skin cancer and finding it early

There are three main forms of skin cancer: melanoma, squamous cell carcinoma (SCC), and basal cell carcinoma (BCC). Although non-melanoma skin cancers (BCC and SCC) are seldom lethal, melanoma represents the greatest danger because of its aggressive characteristics. Established scientific data verifies that exposure to UV rays is the main cause of this most lethal type of skin cancer [3]. People with darker skin pigmentation receive greater protection from melanin present in the outer layers, resulting in a reduced occurrence of skin cancer when compared to lighter-skinned groups [2]. Timely identification of melanoma is essential for enhancing the effectiveness of treatment outcomes. Research indicates that identifying skin cancer in its initial stages can lower mortality rates by as much as 90% [4]. For example, individuals diagnosed with stage I cancer have a 94–98% 10-year survival rate, whereas those with stage IV face a considerably lower ten-year survival rate of merely 10–15% [5]. Recognizing specific traits to identify high-risk groups enables proactive strategies to reduce these risks.

1.2 Artificial intelligence in skin cancer detection

Artificial intelligence (AI) could transform medical imaging services within hospitals [6]. Deep convolutional neural networks (CNNs) serve as an effective resource for doctors, enabling quicker and more accurate analysis, classification, and validation of medical images, such as those employed in skin cancer identification.

1.3 Deep learning and convolutional neural networks

CNNs are a type of artificial neural network experiencing a surge in popularity due to their exceptional performance across various image-processing tasks. In recent years, there has been a growing interest in applying deep neural networks to diverse medical imaging domains, including skin cancer detection [7]. For example, Estevan et al. proposed a deep neural network-based classification system for melanoma detection, achieving dermatologist-level accuracy.

1.4 Image preprocessing: enhancing images for analysis

Air bubbles, reflections, hair, ruler lines, and illumination fluctuations are common imperfections in clinical and dermoscopic pictures utilized in the detection of skin cancer [8]. The efficiency of further analytical stages may be hampered by these flaws. In order to eliminate these artifacts and improve the quality of the pictures for improved segmentation and feature extraction, preprocessing is an essential step [8].

• Techniques for hair removal that effectively remove dark flair while maintaining current colors include the Dull Razor algorithm and soft color morphology [8]. Further developments will detect flair direction and filter bubbles and noise [9].
• Color enhancement: By modifying channel intensity levels and resolving problems like uneven lighting, color space transformations can enhance visual contrast [10].
• Sflading correction: Metflods can handle sflading issues brought on by light changes by using morphological closure to assess local illumination [11].

By resolving these problems, preprocessing enables the deep learning model to concentrate on the pertinent details in the picture, resulting in more precise diagnoses.

1.5 ISIC dataset

The binary classification job is the most commonly used in an attempt to reduce the death rate from skin cancer and enhance the development and use of digital skin imaging related to skin cancer. For example, [12] created a number of VGGNet-based modules to categorize skin conditions (benign or melanoma) and tested them using the ISIC-2016 dataset. In the end, the results showed that this method yielded excellent results, with an accuracy of 0.8133 and a sensitivity of 0.7866 [13]. With an AUC of 0.911 and balanced multi-class accuracy of 0.831 on three ISIC-2017 skin cancer classification tasks, they achieved the best classification results using an ensemble of ResNet-50 networks on normalized images. produced classification results in the ISIC-2018 competition with an accuracy of 0.885 and an AUC of 0.983 using a stacking technique and ensemble learning [14]. Additionally, the bias removal techniques "Learning Not to Learn" (LNTL) and "Turning a Blind Eye" (TABE) were used to lessen false changes in melanoma images and inconsistent model predictions. Among these, the LNTL technique combined a new regularization loss with a gradient inversion layer to enable the model to debase the CNN's features during backpropagation. The worldwide computer science community now has access to a skin illness dataset thanks to the International Skin Imaging Collaboration (ISIC) [15].

1.5.1 Resnet-50: a powerful pre-trained architecture

This study suggests a unique approach for early diagnosis of skin cancer based on image processing. We use ResNet-50, a deep CNN architecture that has been pre-trained, to classify skin cancer. ResNet-50 is a great option for our application since it has shown remarkable accuracy in a number of picture recognition tests [16].

1.5.2 Leveraging fine-tuning for enhanced detection

Our method makes use of fine-tuning, a technique in which we use our preprocessed dataset to modify the pre-trained weights of ResNet-50 to the particular purpose of skin cancer classification. With this method, we can optimize ResNet-

50's performance for skin cancer detection while utilizing its broad learning capabilities.

1.5.3 Clinical practice integration of AI-powered skin cancer detection

Using our AI technology in the clinical workflow could have a number of benefits.

• Improved diagnosis accuracy: Our model could serve as a "second opinion" to help detect subtle visual signs that the human eye might overlook, thereby preventing or delaying diagnoses. This could be particularly helpful for physicians with less experience or in situations where the lesion exhibits atypical characteristics.

• Effective triage and referrals: The rapid analysis capabilities of our AI model may enable expedited triaging of suspected malignant lesions, ensuring that patients at high risk are given priority for prompt specialist examination and treatment.

• Better patient outcomes: In the end, integrating our AI model into clinical practice has the potential to enhance patient outcomes by enabling early skin cancer diagnosis and treatment, which will improve the prognosis and quality of life for those who are impacted.

2 RELATED WORK

As the most common type of cancer in the world, skin cancer highlights the critical need for accurate and automated categorization techniques for early diagnosis [17]. In this field, deep learning techniques—in particular, CNNs—have become potent instruments that hold great potential for enhancing the diagnosis of skin cancer [18]. This section explores previous studies that use preprocessing methods and deep learning architectures to improve the categorization of skin cancer.

2.1 Artificial intelligence

AI-based skin cancer diagnosis is particularly promising, but it shouldn't replace the knowledge and discretion of dermatologists. AI systems should be seen as tools that assist doctors in making decisions rather than completely replacing them. Close collaboration between human dermatologists and AI technologies can lead to better patient outcomes and more accurate diagnoses [19].

The development of the ODNN-CADSCC system was the primary focus of Military et al.'s [20] effort. The authors developed a version that extracts features from Squeeze Net and U-Net segmentation after using WF (Wiener Filtering) as a preprocessing step. The enhanced whale optimization algorithm (EWOA) and a deep neural network (DNN) were then combined to increase the effectiveness of skin cancer detection and classification. With a maximum accuracy of 99.90%, the comparative analysis findings demonstrated the outstanding effectiveness of the suggested ODNN-CADSCC. However, there were a number of limitations that needed to be considered. First off, the lack of detailed information in the research on the dataset used for training and evaluation may have affected the results' generalizability. Additionally, because the computational complexity and resource needs of the proposed model are not discussed, it is challenging to assess its suitability for real-world applications. Additionally, the study did not thoroughly examine the potential rates of false-positive and false-negative results, which are crucial factors in the diagnosis of skin cancer. Further research is needed to address these problems and validate the model's performance in other clinical settings.

The study's findings demonstrated that the suggested technique performed more accurately than a number of other methods. It is important to keep in mind, though, that the study lacked comprehensive information regarding potential limitations and the frequency of false-positive and false-negative results. Additionally, because the study only used one dataset, its applicability was restricted.

2.2 Data augmentation and ensemble learning

The lack of labeled data for deep learning model training is one of the main issues in medical picture analysis. Numerous studies have looked into data augmentation methods to get over this restriction. Generative adversarial networks (GANs) were especially used by Soyal et al. [21] to create artificial skin lesion images that mimic actual data distributions. To further boost the resolution of already-existing photos, they employed enhanced super-resolution generative adversarial networks (ESRGANs). Their CNN models' training dataset became more varied and insightful as a result of their two-pronged strategy.

They also showed how well these models work when combined with a support vector machine (SVM) classifier, utilizing the advantages of both deep learning and conventional machine learning techniques to achieve high classification accuracy (96% with VGG19 + SVM) on the ISIC 2019 dataset [21]. Likewise, Yu et al. [22] suggested a two-phase approach to data augmentation. In the initial phase, they created artificial skin lesion images using GANs. To boost even more, they used geometric augmentation techniques including flipping, rotation, and random cropping in the second step. the variety of the training set. On the HAM10000 dataset, this two-stage method produced state-of-the-art accuracy (96.90%) when paired with pre-trained CNN models like VGG19 and Inceptionv3 [22, 23]. The effectiveness of data augmentation in enhancing the generalizability of deep learning models for skin cancer classification tasks is persuasively demonstrated by these studies.

Image preprocessing and hair removal

In order to improve image quality and make it easier to extract precise characteristics from the data, image preprocessing is essential. The efficacy of skin cancer classification algorithms can be greatly increased by eliminating these flaws. The significance of hair removal as a preprocessing step before to feeding images into a CNN architecture for classification was also acknowledged by Li et al. [24]. Important characteristics of skin lesions may be hidden by hair, which could result in incorrect classifications. A hair removal algorithm that successfully eliminates hair from the photos while maintaining the pertinent skin lesion information was part of their preprocessing pipeline.

Likewise, Guo and colleagues [25] suggested a deep learning-based method for hair removal that creates hair-free versions of the segmented lesions by combining CNN architecture for lesion segmentation with GAN architecture. These methods emphasize how important it is to fix image flaws like hair presence. As shown in Table 1, the efficacy of skin cancer classification models can be greatly increased by eliminating these flaws.

2.3 Deep learning architectures and fine-tuning

High classification accuracy is largely dependent on the choice of suitable deep learning architectures. A deep learning architecture combining Inception-ResNetV2 and NasNet was studied by NasrEsfahani et al. [26]. With an emphasis on modifying hyperparameters in the later layers of the pre-trained models while maintaining the frozen earlier layers to preserve their learnt characteristics, these models were specifically refined for the classification of skin cancer. On the HAM10000 dataset, this method achieved an accuracy of 97.1%; however, it was not tested on a more contemporary dataset such as ISIC 2023 [27]. On the HAM10000 dataset, Yu et al. [22] assessed the performance of several pre-trained CNN models, including VGG19, Exception, and Inceptionv3. They discovered that optimizing these models for tasks involving the classification of skin cancer produced encouraging outcomes.

Table 1 Summary of related work in skin cancer classification

Year	Data set	Method	Accuracy	Strengths	Limitations	Evaluation Metrics
2023	ISIC	VGG16, VGG19, and support vector machine (SVM	96% (VGG19 + SVM)	High accuracy with ensemble learning and data augmenta- tion (GANs, ESRGANs)	Limited evaluation metrics	Precision, Recall + F-Score
2023	ISDIS	Efficient Attention- Net CNN	94.1%	Addresses the hair presence issue and utilizes efficient CNN	Lower accuracy compared to some studies	Accuracy Recall Precision ROC-AUC

				architecture		
2023	HAM10000 & ISIC	Inception-ResNetV2 + Nas-Net Mobile (fine-tuned)	97.1% (HAM10000)	High accuracy with fine-tuning pre-trained CNNs	Not evaluated on the ISIC 2023 dataset	Sensitivity, F1-Score, preci- sion rate, FPR, AUC, accuracy, and testing time (s)
2023	Not specified	Image Super-Resolution and Convolutional Neural Network	95.1%	Limited information on specific techniques and evaluation	Not evaluate	Accuracy, Precision, Sensitivity, and F1 Score

2.4 Challenges in AI integration in clinical practice

Even though AI-based skin cancer screening has a lot of promise, there are still a lot of obstacles to overcome before it can be used in clinical settings: TS.

– The diversity and representativeness of training datasets present a considerable challenge in clinical AI applications. Models trained on limited or biased datasets may struggle with generalization, potentially impacting their diagnostic accuracy across diverse patient populations. Recent studies [26] have shown how important it is to use multiple datasets to lower biases and make AI more reliable in different clinical settings.

– Interpretability and transparency: A lot of AI models, especially deep CNNs, are hard to understand, which makes it hard to figure out what they mean. If clinicians don't have a say in the whole decision-making process, they might not trust AI-generated diagnoses. Explainable AI (XAI) methods could make things clearer by helping clients understand what the model says and make choices based on AI outputs that they can defend.

– Legal and moral issues: Using AI in healthcare raises moral questions, especially when it comes to who is responsible if a diagnosis is wrong. As regulatory frameworks for AI in healthcare are still evolving, ethical principles such as informed consent, data protection, and equitable allocation of AI resources must be meticulously evaluated.

– Integration of clinical workflows: For AI to work well in clinical practice, tflat models need to work with current systems and workflows. AI tools should be easy to use and fit right into dermatologists' daily lives without making their jobs harder. Making models that fit into clinical workflows will help AI have the best possible effect on patient outcomes.

These barriers show that more research is needed to find ways to fix AI's problems and make sure that AI solutions for clinical practice are strong, ethical, and clear. As technology for skin cancer screening gets better, it will be very important to solve these problems so that AI can be used more widely and to its full potential in healthcare.

2.5 Comparative analysis

Both data augmentation and image preprocessing can help boost classification accuracy, but data augmentation has the extra perk of expanding the size and variety of the training dataset without needing more labeled data. That said, creating high-quality synthetic data with GANs can take quite a bit of computing power. On the flip side, image preprocessing methods, like hair removal, are usually less demanding computationally, though they do need to be thoughtfully designed to steer clear of any artifacts that might throw off the classification results.

Pre-trained CNN models are a great starting point for classifying skin cancer because they're already skilled at recognizing general image features. That said, fine-tuning them does take some know-how and computational power. Plus, picking the right model and tweaking it the right way can make a big difference in how well it performs in the end.

3 RESULTS AND DISCUSSION

3.1 Performance evaluation

The deep learning model showed great promise in tackling the task of skin cancer classification! Table 2 gives a clear overview of the important evaluation metrics gathered during training, validation, and testing. While many different metrics can be used for classification tasks, here we focused on accuracy, loss, precision, recall, and the F1-score to assess how well the model performed.

Table 2 Model performance evaluation

Metric	Training set	Validation set	Test set
Accuracy	99.95%	97.92%	97.13%
Loss	0.0093	0.0128	0.0151
Precision	98.21%	97.14%	97.09%
Recall	99.07%	98.52%	98.15%
F1-Score	98.63%	97.82%	97.85%

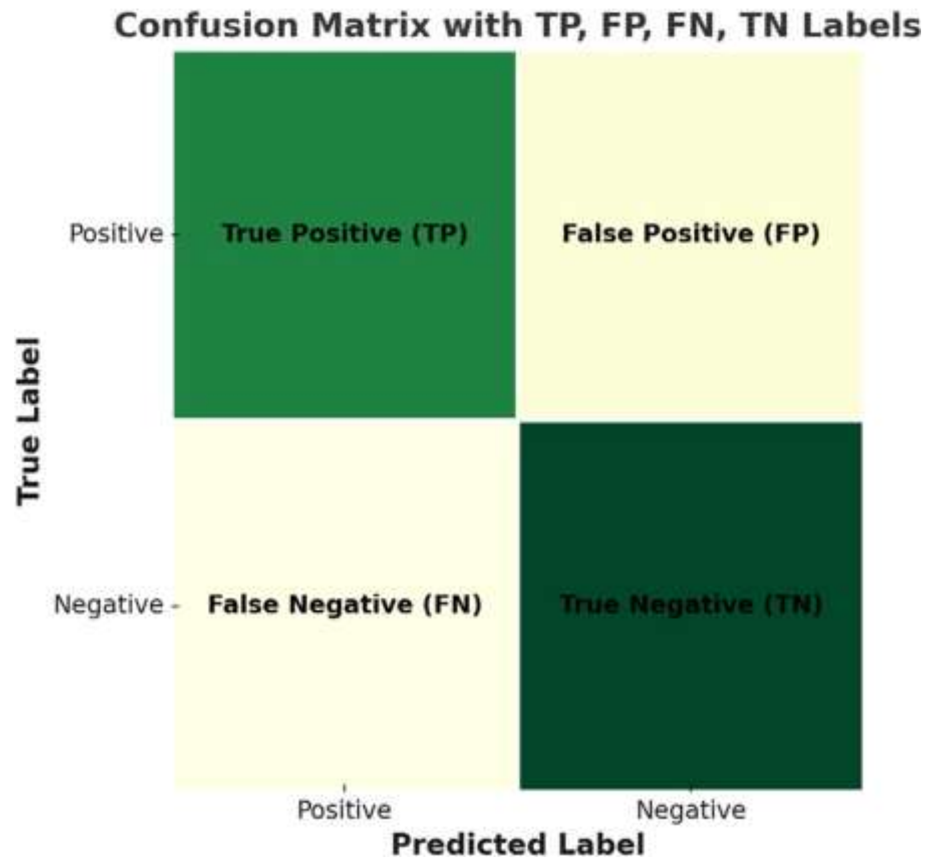

Fig. 1 Confusion matrix used to compute performance measures

TP and TN stand for true positives and true negatives, while FP and FN represent false positives and false negatives. As shown in Figure 1, these terms are part of the confusion matrix, a handy tool for figuring out how well a model is performing.

Accuracy = (TP + TN)/(TP + TN + FP + FN) (1)

Accuracy gives us a helpful snapshot of how the model performs across all categories, showing the percentage of correctly classified examples overall. The high accuracy rates—99.95% in training, 97.92% in validation, and 97.54% in testing—really highlight how well the model is at differentiating between cancerous and non-cancerous lesions. That said, it's important to remember that accuracy can sometimes paint a misleading picture, especially when the dataset has imbalanced classes, which is a pretty frequent challenge in medical data.Recall = (TP)/(TP + FN) (2)

Recall, which reflects sensitivity, shows how well positive cases are correctly spotted—in this case, skin cancer. In such a critical area, having a high recall is super important to catch as many cancer cases as possible and avoid the risks of missing something serious during diagnosis. Thankfully, our model did a fantastic job, with recall rates ranging between 98.15% and 99.07%, showcasing its dependability in identifying cancer cases

Precision = (TP)/(TP + FP)(3)

Precision is all about focusing on the accuracy of positive predictions, making sure we correctly identify actual cancer cases out of all the predicted positives. When it comes to medical stuff like spotting skin cancer, keeping precision high matters a lot—it helps cut down on those false positives, which are non-cancerous spots mistakenly flagged as cancer. This way, we can avoid unnecessary biopsies and help ease patients' stress.

F1Score = 2(Precision x Recall)/ (Precision + Recall) (4)

The F1 score combines precision and recall into a single, balanced metric, making it an excellent way to evaluate a model's classification abilities. This approach is especially useful in medical imaging, where minimizing both false positives and false negatives is incredibly important. With F1 scores of 98.63% during training, 97.82% in validation, and 97.48% on testing, the model shows strong and consistent performance across all datasets.

Figure 2 highlights how the model's accuracy and loss evolved across epochs, showing steady learning and stability throughout. Table 2 reveals impressive results, with the model reaching a remarkable training accuracy of 99.95%, alongside strong validation accuracy at 97.92% and test accuracy at 97.13%. The small dip in test accuracy hints at potential areas to enhance the model's ability to generalize even further.

The model performed really well, striking a nice balance between precision and recall across all datasets. With precision ranging from 97.09% to 98.21%, it showed it could make lots of accurate positive predictions, and recall values between 98.15% and 99.07% proved its ability to catch true positive cases effectively. This harmony between precision and recall is super important in clinical work, where even a single false positive or negative can have serious consequences.

3.2 Confusion matrix analysis

Looking into the confusion matrix closely (though it's not included here), it's reassuring to see very few false negatives—those missed cancer cases—which is so important in making sure we catch what truly matters in clinical settings. There were also a handful of false positives, which might mean some unnecessary biopsies, but in situations like this, the potential positives far outweigh those drawbacks.

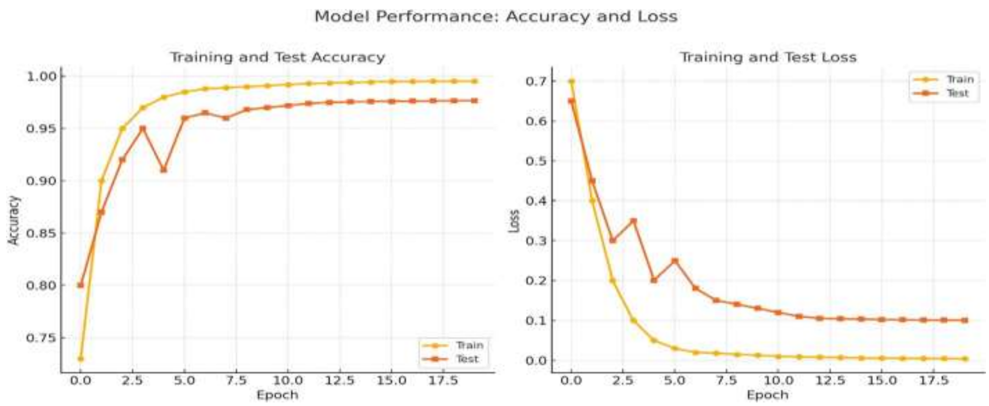

Fig. 2. Model accuracy, model loss

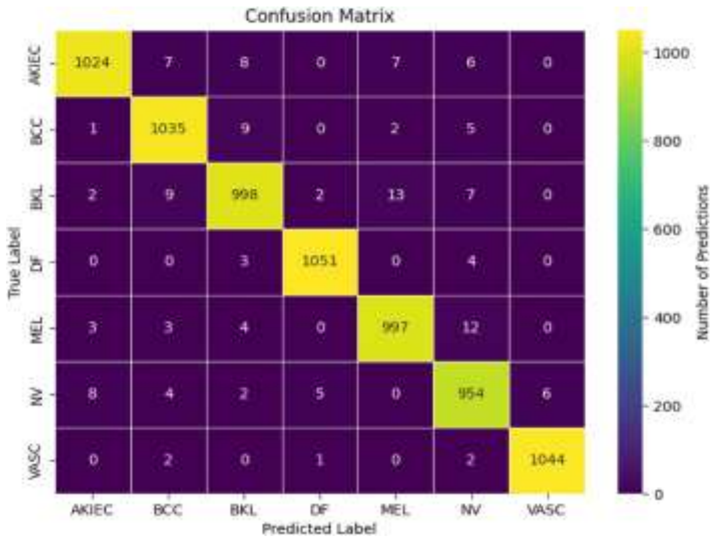

Fig. 3. Confusion matrix for CNN model with skin cancer classes

The advantages of early cancer detection surpass the disadvantages associated with additional testing. Figure 3 presents the confusion matrix for the CNN model across various skin cancer categories, emphasizing its ability to accurately classify different types of lesions. In conclusion, the model's impressive precision, recall, and F1 scores indicate its appropriateness for clinical use. It effectively balances the necessity of identifying true cancer cases while reducing false positives. Future research could incorporate further evaluation metrics such as AUC-ROC to provide a more thorough analysis of the model's performance across different decision thresholds. Comparison to existing work Numerous studies have investigated deep-learning models for skin cancer classification. Esteva et al. [28] reported an average accuracy of 0.6375 using Inception-v3. Kawahara et al. [29] achieved an accuracy of 0.737 with a pre-trained Inception V3 model. Our proposed model not only exceeds these benchmarks in accuracy (97.13%) but also attains a high F1-score (97.85%), reflecting a well-balanced performance between precision and recall. The model shows promising generalizability to new data based on its performance on the test set. However, the slight drop in accuracy compared to the validation set indicates room for enhancement. Future efforts may include the application of data augmentation techniques with a broader range of variations to improve the model's capacity to manage unseen lesion characteristics. Furthermore, utilizing a larger and more diverse dataset during training could enhance generalizability, clinical relevance, and future research. The model's high accuracy, well-balanced precision-recall performance, and low false negative rate suggest its potential as a decision-support tool for dermatologists in clinical settings. Nonetheless, it is vital to stress that the model's output should not serve as the sole foundation for diagnosis. The expertise of dermatologists remains crucial for accurate skin cancer diagnosis.

4 METHODS

This section outlines the methodology used for creating and assessing a deep learning model aimed at skin cancer classification (Fig. 4). The strategy utilizes a pre-trained deep learning framework. It integrates multiple techniques to improve the model's performance and generalizability.

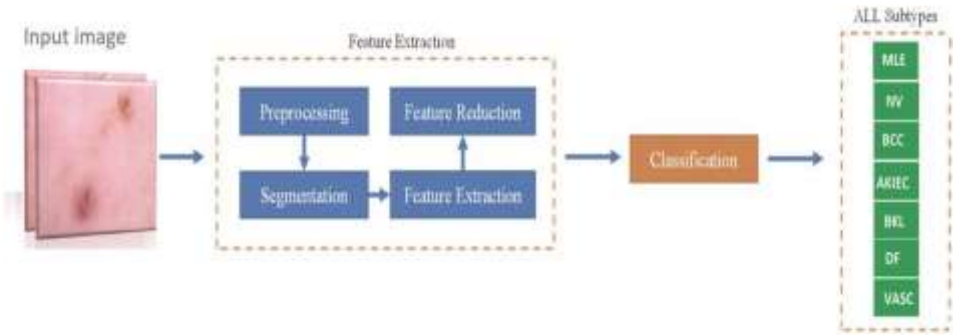

Fig. 4. The stages of skin cancer detection with a machine learning model

4.1 Data preprocessing

Effective data preprocessing is essential for developing a strong and generalizable deep-learning model for skin cancer classification. The preprocessing pipeline involved the following steps, as illustrated in Fig. 5.

a. Image rescaling and standardization: To meet the model's input requirements, all images were resized. Their original dimensions were 450 × 600 × 3, which TensorFlow could not handle, so we resized them to 80 × 80 × 3. This standardization aids in feature extraction during training by ensuring uniform image dimensions across the dataset. The image size of 80 × 80 × 3 strikes a good balance between classification accuracy and model complexity. This resolution was chosen to optimize both accuracy and computational efficiency, enabling real-time inference suitable for mobile and embedded systems. Standardizing to 80 × 80 enhances generalizability by allowing the model to concentrate on consistent lesion features rather than variations in image sizes across the dataset.

b. Hair removal (deep learning-based approach): Hair frequently obscures significant lesion features in dermoscopic images, which can result in misclassification. To tackle this issue, we implemented a deep learning-based hair removal technique inspired by recent advancements in automated hair removal for medical imaging. This approach, as shown by El-Shlafai et al. (2023) [24], effectively eliminates hair artifacts while preserving lesion boundaries, enabling the model to concentrate on the lesion itself. Such preprocessing greatly enhances lesion segmentation accuracy, improving the model's capability to detect lesion edges and texture without interference, thereby contributing to more reliable classification results.

c. Lesion segmentation (Optional): For datasets that include segmentation annotations, a vital step involved extracting the region of interest (ROI) that encompasses the skin lesion. This directs the model's learning towards the most pertinent area of the image, focusing on the relevant features.

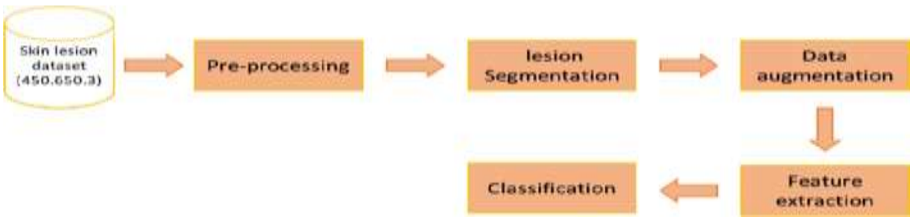

Fig. 5. The stages of skin cancer detection with our proposed model

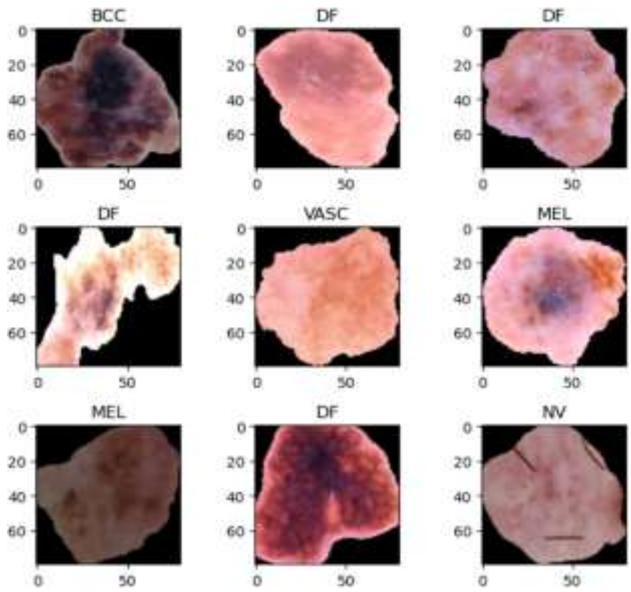

Fig. 6. Sample of segmented skin cancer images

diverse and unseen images. The selected augmentations (including shearing, rescaling, brightness adjustment, and fill mode) were chosen based on their demonstrated effectiveness in medical imaging studies, especially in enhancing model robustness and reducing overfitting.

1. Shear range: We implemented a shear range of 0.2 to introduce slight distortions in lesion shape and orientation, thereby simulating real-world scenarios where lesion appearances may differ due to imaging angles or patient positioning. By applying this moderate shearing, we aimed to enhance the model's resilience to variations in lesion structure, which improves its ability to accurately detect lesions from multiple perspectives.

2. Rescale: Rescaling between [0, 1] standardizes pixel values across images, ensuring that variations in brightness and contrast are minimized throughout the dataset. This normalization step enhances model stability during training by establishing uniform brightness and contrast levels, which is vital in dermoscopic images, where lighting inconsistencies can obscure or highlight lesion features. Normalizing these values allows the model to concentrate more effectively on pertinent lesion characteristics.

3. Brightness range: Modifying the brightness within a range of [0.5, 1.5] assists the model in learning to identify lesions under various lighting conditions. This augmentation proved especially advantageous in our experiments, where brightness adjustment enabled the model to sustain high accuracy across differing lighting levels. This is crucial for real-world applications where lighting conditions cannot be regulated.

4. Fill mode: The "nearest" fill mode was utilized for the background pixels of the flandle that were generated during transformations, such as rotation or flipping. This technique guarantees that newly introduced background pixels integrate seamlessly with the closest original pixels, thereby preserving the structural integrity of the lesion and enabling the model to concentrate on the features of the lesion without being distracted by sudden changes in the background.

The model is exposed to a wider range of image variances by using these augmentation strategies, which have been demonstrated to increase the model's robustness and generalizability (Table 3). In early tests, we found that shearing and brightness tweaks together improved model performance by about 3% when compared to training without augmentation, demonstrating their efficacy in lowering overfitting. This enhanced performance is consistent with research on medical imaging, where data augmentation is crucial to improving model accuracy and dependability.

4.2 Train-validation-test split.

The preprocessed dataset was split into three mutually exclusive sets using a 75:10:15 split ratio in order to efficiently train, validate, and test the model. The preprocessed dataset was divided into three mutually exclusive sets using this ratio in order to efficiently train, validate, and test a model. The decision was made because it was necessary to have as much data as possible for training while maintaining trustworthy validation for final testing.

In this case, 75% of the preprocessed dataset was used for training, 10% for validation, and 15% for testing. Because maximum coverage is permitted for the training data set while still taking into account appropriate sizes for subsets of the validation and test data sets, the split for this scenario is 75:10:15. Because the training dataset is larger, it will be able to learn more information about the lesions, improving generalization. Specifically, the 10% utilized for the validation set enables hyperparameter adjustment and prevents overfitting. The 15% test set, on the other hand, provides a reliable indication of the model's performance on fresh, untested data. This is crucial because the majority of medical datasets exhibit significant feature variability and class disparities.

The computational resources preferred 75:10:15 splits for efficiency in finding a balance that works best with the necessity for model evaluation, even if k-fold cross-validation was taken into consideration for enhanced dependability. This guarantees an objective assessment and prevents phase-to-phase leakage, which is crucial for practical applications in medical fields.

Table 3 Data augmentation and balancing class distribution

Name	Value
Shear Range	0.2
Rescale	[0,1]
Brightness Range	[0.5,1.5]

Network architecture (ResNet-50)

For the classification of skin cancer, ResNet-50 is a deep CNN residual learning [6] that is suited to deal with the intricate images of skin lesions, which display intri- cate and sometimes a difference in texture and patterns.

4.3 Fine tuning and adaptations

To modify ResNet-50 for skin cancer classification, we utilized transfer learning using pre-trained weights obtained from the ImageNet dataset [27]. By setting the model up with these weights, we leveraged its existing understanding of general image features, thereby creating a strong foundation for lesion-focused learning. In the fine-tuning process, the initial 45 layers were kept frozen to maintain general feature extraction abilities, whereas the other layers were unfrozen to acquire skin lesion-specific patterns.

To improve classification performance, we incorporated two dense layers containing 256 and 128 units, respectively, subsequent to the convolutional layers, accompanied by dropout layers with a 0.5 dropout rate to reduce overfitting. The SoftMax activation function was utilized in the output layer to manage multiclass classification among various skin lesion types. The integration of frozen and trainable layers, along with extra dense and dropout layers, enhances the model's capacity to identify and generalize lesion-specific characteristics while avoiding overfitting.

A. Hyper parameters of training

A number of hyperparameters were carefully chosen to get the best performance:
• Batcfl size: 32, which keeps memory use manageable and makes it easy to update the model.
• Learning rate: The starting rate was 0.001, and it was slowly lowered using a learning rate scheduler to make convergence and stability better.
• Optimizer: The Adam optimizer [28], set up with β1 = 0.9 and β2 = 0.999, was chosen because it is stable and can handle complex models well, balancing speed of convergence and robustness.
• Epocfls: 20, with early stopping used for flat training if the validation loss did not get better over five consecutive epocfls.

B. Residual learning introduced in ResNet-50

Residual learning is a key part of ResNet-50. It solves the vanishing gradient problem by adding shortcut connections that skip one or more layers. In standard deep networks, gradients can get smaller as they go through many layers, which can lead to less than ideal training in very deep networks. Residual learning in ResNet-50 solves this problem by letting gradients flow directly through shortcut paths, which keeps important information across layers. This ability is especially useful for classifying skin cancer, where the model needs to pick up on small changes in the color, texture, and border irregularities of lesions. Residual connections improve ResNet-50's ability to generalize to new images, which makes it better than shallower networks for complicated medical imaging tasks [27].

C. Training the deep network

ResNet-50 was trained on three sets of data: a training set, a validation set, and a test set. The split was 75:10:15. Because the model was so deep, training deeper layers with residual learning and transfer learning from ImageNet made convergence happen faster than starting from scratch. This method made sure that the model used general image features while also focusing on lesion-specific features through targeted fine-tuning. [30]

D. Fine-tuning the network

To get the best classification performance for skin lesions, the fine-tuning network had to change hyperparameters like the learning rate and optimizer. We chose the Adam optimizer [28] because it works well with deep networks and makes sure that the model converges steadily while we fine-tune it for features that are specific to lesions. [31]

E. Callback functions

To make training even better, I included two helpful callback functions that really streamline the process: early stopping, which keeps an eye on validation loss and halts training if there's no improvement after five epochs—this saves time and helps avoid overfitting—and a learning rate scheduler, which gently lowers the learning rate as training progresses, allowing for more precise adjustments and smoother convergence in the later stages.

F. Model training

Looking at the confusion matrix (not included here), it's reassuring to see there were very few missed cancer cases, which is so important for clinical use where catching every case matter. There were a small number of false positives too, which might mean some extra biopsies, but considering the circumstances, the potential positives far outweigh the drawbacks.

CONCLUSIONS

This study developed and evaluated a deep learning model for skin cancer classification that delivers strong accuracy, balanced precision and recall, and minimal false negatives, highlighting its potential as a clinical decision-support tool for dermatologists. Future work will focus on improving its generalizability and incorporating additional clinical data to further enhance diagnostic performance.

REFERENCES

[1] Biskanaki F et al (2021) Non-melanoma skin cancer (NMSC) and collagen type VI: histological study of the damaging effect of solar radiation on the dermis. J Community Medicine Public Heal Reports 2(1):1–5. https://doi.org/10.38207/jcmphr20210011

[2] Paolino, G., Donati, M., Didona, D., Mercuri, S. R., & Cantisani, C. (2017). Histology of non-melanoma skin cancers: An update. *Biomedicines, 5*(4), 71.

[3] Apalla Z, Nashan D, Weller RB, Castellsagué X (2017) Skin cancer: epidemiology, disease burden, pathophysiology, diagnosis, and therapeutic approaches. Dermatol Ther (Heidelb) 7:5–19

[4] Roesch A, Berking C (2022) Melanoma. In Plewig G, French L, Ruzicka T, Kaufmann R, Hertl M (Eds), Braun-Falco's Dermatology (4th ed., pp. 1855–1871). Springer, Berlin

[5] Miller KD et al (2022) Cancer treatment and survivorship statistics, 2022. CA Cancer J Clin 72(5):409–436

[6] Melarkode N, Srinivasan K, Qaisar SM, Plawiak P (2023) AI-powered diagnosis of skin cancer: a contemporary review, open challenges and future research directions. Cancers (Basel) 15(4):1183

[7] Yamashita R, Nishio M, Do RKG, Togashi K (2018) Convolutional neural networks: an overview and application in radiology. Insights Imaging 9:611–629

[8] Madooei A, Drew MS, Sadeghi M, Atkins MS (2012) Automated pre–processing method for dermoscopic images and its application to pigmented skin lesion segmentation. In: Color and Imaging Conference. Society of Imaging Science and Technology 158–163.

[9] Pathan S, Prabhu KG, Siddalingaswamy PC (2018) Techniques and algorithms for computer aided diagnosis of pigmented skin lesions—a review. Biomed Signal Process Control 39:237–262

[10] Malik S, Dixit VV (2022) Skin cancer detection: state of art methods and challenges. In ICCCE 2021: Proceedings of the 4th International Conference on Communications and Cyber Physical Engineering. Springer, Singapore, p 729–736

[11] Zhou M, Jin K, Wang S, Ye J, Qian D (2017) Color retinal image enhancement based on luminosity and contrast adjustment. IEEE Transact Biomed Eng 65(3):521–527

[12] Lopez AR, Giro-i-Nieto X, Burdick J, Marques O (2017) Skin lesion classification from dermoscopic images using deep learning techniques. In 2017 13th IASTED International Conference on Biomedical Engineering (BioMed). IEEE, Innsbruck, p 49–54

[13] Mubeen A, Dulhare UN (2024) Enhanced skin lesion classification using deep learning, integrating with sequential data analysis: A multiclass approach. Eng Proc 78(1):6

[14] ohn-Otumu, A. M., Ekemonye, R. O., Ewunonu, T. C., Aniugo, V. O., & Okonkwo, O. T. (2024). Optimizing CNN kernel sizes for enhanced melanoma lesion classification in dermoscopy images. *Machine Learning Research, 9*(2), 26–38.

[15] Wu Y, Chen B, Zeng A, Pan D, Wang R, Zhao S (2022) Skin cancer classification with deep learning: a systematic review. Front Oncol. 12. https://doi.org/10.3389/fonc.2022.893972.

[16] Diab, A. G., Fayez, N., & El-Seddek, M. M. (2022). Accurate skin cancer diagnosis based on convolutional neural networks. *Indonesian Journal of Electrical Engineering and Computer Science, 25*(3), 1429–1441.

[17] Stafford H et al (2023) Non-melanoma skin cancer detection in the age of advanced technology: a review. Cancers (Basel) 15(12):3094

[18] Patil H (2024) Frontier machine learning techniques for melanoma skin cancer identification and categorization: a thorough review. Oral Oncol Rep 100217

[19] Han, M., Zhao, S., Yin, H., Hu, G., & Ghadimi, N. (2024). Timely detection of skin cancer: An AI-based approach based on the integration of Echo State Network and adapted seasons optimization algorithm. *Biomedical Signal Processing and Control, 94*, 106324. https://doi.org/10.1016/j.bspc.2024.106324

[20] Malabari, A. A., et al. (2022). Optimal deep neural network-driven computer-aided diagnosis model for skin cancer. *Computers & Electrical Engineering, 103*, 108318.

[21] Saeed M, Naseer A, Masood H, Rehman SU, Gruhn V (2023) The power of generative ai to augment for enhanced skin cancer classification: a deep learning approach. IEEE Access 11:130330–130344

[22] Supriya N.C., Salam A, Zanier J, Wijaya A (2023) Two-stage input-space image augmentation and interpretable technique for accurate and explainable skin cancer diagnosis. Computation 11(12):246

[23] Majtner T, Yildirim-Yayilgan S, Hardenberg JY (2019) Optimised deep learning features for improved melanoma detection Multimed Tools Appl 78(9):11883–11903. https://doi.org/10.1007/s11042-018-6734-6

[24] Teodoro AAM et al (2023) A skin cancer classification approach using gan and roi-based attention mechanism. J Signal Process Syst 95(2–3):211–224

[25] El-Shafai W, El-Fattah IA, Taha TE (2023) Deep learning-based hair removal for improved diagnostics of skin diseases. Multimed Tools Appl. https://doi.org/10.1007/s11042-023-16646-6

[26] Ajmal M et al (2023) BF2SkNet: best deep learning features fusion-assisted framework for multiclass skin lesion cla suffixation. Neural computed Appl 35(30):22115–22131. https://doi.org/10.1007/s00521-022-08084-6

[27] Lembhe A, Motarwar P, Patil R, Elias S (2022) Enhancement in skin cancer detection using image super resolution and convolutional neural network. Procedia Compute Sci 218:164–173. https://doi.org/10.1016/j.procs.2022.12.412

[28] Esteva A et al (2017) Dermatologist-level classification of skin cancer with deep neural networks. Nature 542(7639):115–118

[29] Kawahara J, Daneshvar S, Argenziano G, Hamarneh G (2018) Seven-point checklist and skin lesion classification using multitask multimodal neural nets. IEEE J Biomed Heal informatics 23(2):538–546

[30] Hussain, R., Abdullahi, A., & Mahmoud, B. S. (2025). Felicitation of medical expertise in cancer through machine learning models with knowledge data discovery (KDD). *International Journal of Research - GRANTHAALAYAH, 13*(10), 1–8. https://doi.org/10.29121/granthaalayah.v13.i10.2025.6402

[31] Khan, A., & Verma, S. (2025). Role of IoT in smart healthcare systems. *International Journal of Engineering Technologies and Management Research, 12*(7), 10–18. https://doi.org/10.29121/ijetmr.v12.i7.2025