SKIN CANCER DETECTION USING DEEP LEARNINIG

[1]CH. Roshini
M.Tech,
computer science and engineering,
SR University, Hasanparthy,
Hanamkoda, Telangana
Email: roshini35891@gmail.com

[2]DR.N Venkatesh
Assoc. Prof. computer science and engineering,
SR University, Hasanparthy,
Hanamkoda, Telangana
Email: venkatesh.naramula@sru.edu.in

Abstract

Skin cancer is one of the most common and life-threatening forms of cancer worldwide, with its incidence increasing due to factors such as ozone depletion and prolonged ultraviolet (UV) exposure. Early detection plays a critical role in improving survival rates; however, conventional diagnostic methods, such as dermatological examinations and biopsies, are time-consuming, expensive, and prone to subjectivity. Recent advancements in artificial intelligence, particularly deep learning, have shown significant potential in addressing these challenges by enabling automated, accurate, and efficient skin cancer detection.This research proposes a deep learning–based skin cancer detection system using Convolutional Neural Networks (CNNs) for the classification of dermoscopic images. The system leverages the HAM10000 dataset to classify seven skin cancer types, including melanoma and non-melanoma variants. Through preprocessing techniques such as resizing, normalization, and data augmentation, the proposed model ensures improved robustness and generalization. Furthermore, advanced architectures like ResNet50 are explored to enhance classification accuracy.The experimental results demonstrate that CNN-based approaches achieve high precision and reliability, significantly outperforming traditional methods. By integrating such models into mobile and cloud-based healthcare platforms, the system offers scalability, accessibility, and real-time diagnostic support, making it particularly useful in remote and resource-limited regions. The proposed work highlights the potential of deep learning to revolutionize dermatological screening, improve diagnostic accuracy, reduce reliance on invasive procedures, and ultimately contribute to better patient outcomes and global healthcare standards.

Keywords: Skin Cancer, Deep Learning, Convolutional Neural Networks, Medical Image Analysis, Computer-Aided Diagnosis.

1. Introduction

Skin cancer is one of the most prevalent and life-threatening diseases worldwide, with the number of cases rising steadily due to environmental and lifestyle factors. Prolonged exposure to ultraviolet (UV) radiation, genetic predisposition, and depletion of the ozone layer caused by pollution have all contributed to the global increase in incidence rates. According to the World Health Organization (WHO), approximately three million cases of non-melanoma skin cancer and more than one hundred thousand cases of melanoma are diagnosed each year. Statistics further indicate that one in five individuals is likely to develop some form of skin cancer during their lifetime. This makes skin cancer a critical public health concern that requires early detection and effective diagnostic methods.

Skin cancers are generally classified into two main categories. The first is melanoma, an aggressive and malignant form of skin cancer that spreads rapidly if not diagnosed early. The second is non-melanoma skin cancer, which includes basal cell carcinoma, actinic keratoses, and other benign lesions. While melanoma is considered the most dangerous form, both categories pose significant risks if left untreated. Early detection is therefore vital for improving patient survival rates, reducing treatment costs, and preventing complications such as metastasis or permanent tissue damage.

Conventional methods for diagnosing skin cancer, including dermatological visual examinations and biopsy procedures, are reliable but often time-consuming, invasive, and expensive. Moreover, these methods are heavily dependent on the experience and expertise of dermatologists, which can lead to subjective interpretations and variations in diagnosis. In resource-constrained or remote regions where access to specialized dermatologists is limited, the challenge of timely diagnosis becomes even more pronounced. This situation underscores the urgent need for automated, efficient, and widely accessible systems that can assist healthcare professionals in the early detection of skin cancer.

Recent advances in Artificial Intelligence (AI) and Deep Learning (DL) have shown remarkable potential in the field of medical imaging. Convolutional Neural Networks (CNNs), in particular, have revolutionized image classification tasks due to their ability to automatically extract meaningful features from raw image data without manual intervention. In dermatology, CNNs can analyze dermoscopic images of skin lesions and distinguish between benign and malignant conditions with high accuracy. This capability not only enhances the efficiency of diagnosis but also reduces dependence on invasive procedures such as biopsies.

The objective of this research is to design and implement a CNN-based system capable of classifying dermoscopic images into different types of skin cancer. The proposed system uses the HAM10000 dataset, which contains a diverse collection of annotated skin lesion images. By applying preprocessing steps such as resizing, normalization, and augmentation, the dataset is prepared for effective model training. The CNN architecture is optimized for robust feature extraction, and advanced transfer learning methods such as ResNet50 are also explored to further enhance performance.

This research envisions the integration of the proposed model into mobile and cloud-based healthcare platforms, making skin cancer screening more scalable and accessible, particularly in underserved regions. By combining deep learning with medical imaging, the system provides a non-invasive, accurate, and efficient solution to support dermatologists in clinical decision-making and improve patient outcomes. The overarching goal is to demonstrate that AI-powered systems can bridge the gap between technological innovation and practical healthcare applications, ultimately contributing to better survival rates and improved global healthcare standards.

2. Literature Survey

Skin cancer detection has been an active area of research in recent years, with several studies highlighting the effectiveness of deep learning techniques for accurate classification of skin lesions. Researchers have explored different approaches ranging from traditional image processing and machine learning methods to advanced convolutional neural network architectures. Early studies focused on handcrafted feature extraction methods, where lesion characteristics such as color, shape, texture, and symmetry were analyzed using classical classifiers like Support Vector Machines (SVM) and Decision Trees. While these approaches provided a foundation, they often struggled with scalability, high-dimensional data, and variations in lesion appearance across different patients and demographics.

Md Shahin Ali et al. proposed an enhanced technique of skin cancer classification using deep convolutional neural networks with transfer learning. Their model achieved a training accuracy of 93.16% and testing accuracy of 91.93%, demonstrating that transfer learning significantly improves classification performance on limited datasets. Mahamudul Hasan and colleagues presented a CNN-based approach for early diagnosis of skin cancer, reporting a training accuracy of 93.7% and a testing accuracy of 89.5%. Their work emphasized the importance of automated detection systems to support dermatologists in clinical practice.

Another notable contribution is from Mehwish Dildar et al., who conducted a systematic review of deep learning techniques for early detection of skin cancer. Their findings suggested that CNNs, due to their ability to learn hierarchical features from images, outperform traditional machine learning methods that rely heavily on feature engineering. Yunendah Nur Fu'adah and collaborators developed a CNN model with three hidden layers and multiple optimizers, achieving an accuracy of 99% on a publicly available dataset. This result highlights the potential of carefully designed deep architectures to achieve near-perfect performance in controlled experimental settings.

Hybrid approaches have also been explored to improve classification accuracy. Jinen Daghrir and colleagues proposed a system that combined CNNs with classical machine learning classifiers, using parameters such as lesion borders, texture, and color to perform majority voting for final predictions. Their hybrid approach demonstrated that combining the strengths of deep learning with traditional classifiers can enhance robustness in diagnosis. Similarly, Viswanatha Reddy Allugunti constructed a CNN-based diagnostic model for melanoma using a well-known dataset and achieved an overall accuracy of 88.83%, establishing a strong benchmark for CNN-based systems.

Other researchers have focused on exploring different architectures and datasets. Tanzina Afroz Rimi and colleagues developed Derm-NN, a CNN model with four convolutional layers trained on 2400 images, which achieved an accuracy of 73%. Though lower compared to other models, their work demonstrated the challenges posed by limited and imbalanced datasets. T. Shanthi et al.

utilized the AlexNet architecture to categorize skin diseases, reporting competitive results on smaller datasets. Their research underscores the role of pre-trained networks in boosting classification performance when data availability is limited.

From the literature, it is evident that CNN-based approaches consistently outperform traditional methods in terms of accuracy and robustness. However, challenges such as dataset imbalance, generalization to diverse populations, and the need for clinical validation remain significant barriers to widespread adoption. Moreover, while many studies report high accuracy on benchmark datasets, real-world applications require systems to handle noisy data, variations in imaging conditions, and differences in skin types. These gaps in existing research provide a strong motivation for the development of improved CNN-based models that are not only accurate but also scalable and adaptable for practical healthcare deployment.

3. System Analysis

3.1 Existing System

The existing systems for skin cancer detection rely heavily on manual diagnosis by dermatologists, which typically involves visual inspection of the affected area using dermoscopy, followed by histopathological analysis of biopsy samples when abnormalities are suspected. While this process remains the gold standard for diagnosis, it is time-consuming, invasive, and costly. Furthermore, the accuracy of manual diagnosis depends greatly on the expertise of the dermatologist, which introduces subjectivity and variability in clinical outcomes. In regions where access to specialized dermatologists and advanced diagnostic equipment is limited, patients often face delays in diagnosis and treatment, leading to poor prognoses.

In recent years, Computer-Aided Diagnosis (CAD) systems have been introduced as an aid to dermatologists. These systems generally employ machine learning algorithms to analyze digital images of skin lesions and provide classification outputs that support clinical decision-making. CAD systems, however, face several limitations. Their performance depends heavily on the quality of the dataset used for training, and they often struggle to generalize well when applied to diverse populations. Additionally, many existing CAD systems require extensive preprocessing and feature engineering, making them less flexible for real-world clinical use. As the complexity of skin lesions increases with variations in color, size, shape, and texture, traditional approaches struggle to capture the subtle differences that separate benign from malignant lesions. Consequently, the diagnosis process remains labor-intensive and prone to errors, highlighting the need for more advanced and reliable solutions.

3.2 Proposed System

To overcome the limitations of traditional diagnosis and existing CAD systems, this research proposes a deep learning–based system utilizing Convolutional Neural Networks (CNNs) for the automated detection and classification of skin cancer. CNNs have the unique ability to extract and learn features directly from raw dermoscopic images without the need for manual feature engineering. This enables the system to identify intricate visual patterns that may be imperceptible to the human eye, thereby improving diagnostic accuracy and consistency. The proposed model focuses on classifying seven distinct types of skin lesions, including melanocytic nevi, melanoma, benign keratosis-like lesions, basal cell carcinoma, actinic keratoses, vascular lesions, and dermatofibroma. The system employs the publicly available HAM10000 dataset, which contains a diverse set of dermoscopic images suitable for training and testing.

The design of the system incorporates preprocessing techniques such as image resizing, normalization, and augmentation to address the challenges of dataset imbalance and variability. Augmentation strategies including rotation, flipping, and zooming ensure that the model generalizes well across unseen data. The CNN architecture is carefully optimized with multiple convolutional and pooling layers, followed by dense layers for classification. Additionally, transfer learning approaches using advanced architectures such as ResNet50 are explored to further enhance accuracy. By leveraging these techniques, the proposed system significantly improves upon traditional CAD systems in terms of precision, scalability, and adaptability.

The system is designed not only to achieve high accuracy in controlled experiments but also to be practical for real-world applications. By integrating the model into mobile or cloud-based platforms, dermatological screening can be made more accessible to patients in remote and underserved regions. Real-time classification of skin lesions enables timely medical intervention, potentially reducing the need for invasive biopsies and lowering diagnostic costs. The proposed system, therefore, not only addresses the limitations of existing methods but also envisions a scalable, efficient, and patient-centric approach to skin cancer detection.

4. System Design

The design of the proposed skin cancer detection system is structured to ensure efficient preprocessing of dermoscopic images, robust feature extraction using deep learning models, and accurate classification of skin lesions. The system architecture is designed in multiple stages, beginning with image acquisition and preprocessing, followed by feature learning through Convolutional Neural Networks (CNNs), and concluding with the classification of lesions into predefined categories.

At the core of the system is the CNN architecture, which consists of a sequence of convolutional layers responsible for extracting hierarchical features from input images. These layers are interspersed with pooling layers to reduce dimensionality and prevent overfitting. The learned features are then passed through fully connected layers, which map the high-level representations to output classes. Activation functions such as ReLU (Rectified Linear Unit) are used in hidden layers to introduce non-linearity, while a softmax or sigmoid activation is applied in the output layer to generate probabilistic classification outputs. To further improve model robustness, dropout regularization is employed to reduce the risk of overfitting during training.

The system also incorporates preprocessing techniques to prepare dermoscopic images for analysis. Images are resized to a fixed dimension, normalized for consistent pixel intensity distribution, and augmented through rotation, flipping, and scaling to ensure variability in training data. These steps not only enhance the model's generalization capability but also address the imbalance in the dataset by artificially increasing the number of samples for underrepresented classes.

The architecture of the proposed system can be represented through a block diagram that outlines the flow from data input to classification output. The design begins with the acquisition of dermoscopic images, which undergo preprocessing before being fed into the CNN. The CNN then performs feature extraction, followed by classification into one of the seven cancer categories.

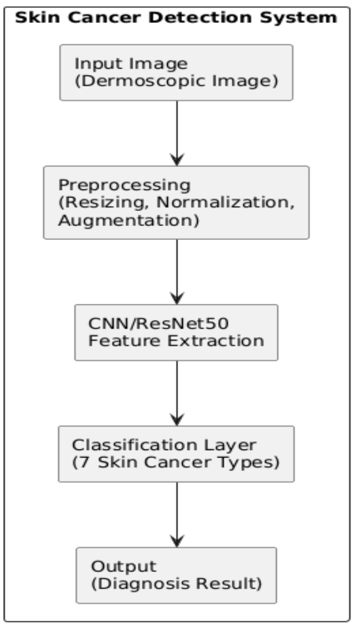

To better capture the interaction between different system components, Unified Modeling Language (UML) diagrams are utilized. A use case diagram provides an overview of how users, such as dermatologists or patients, interact with the system to upload images and receive diagnostic outputs

.

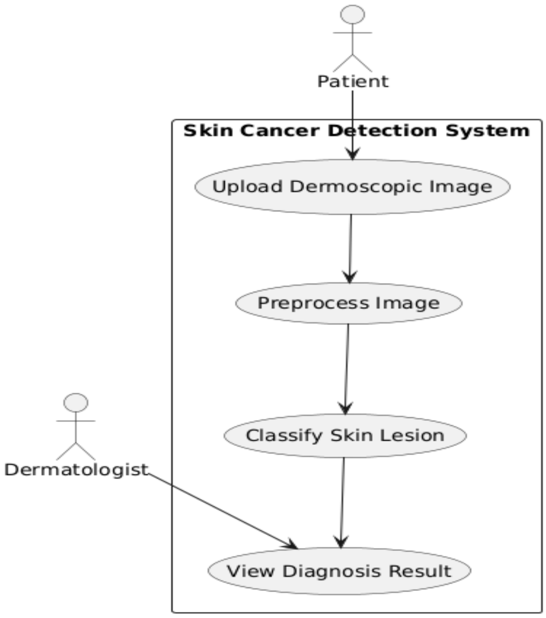

A class diagram illustrates the internal structure of the system, showing the relationships among data preprocessing modules, CNN model components, and classification outputs.The sequence diagram demonstrates the order of interactions between system modules, including the input of dermoscopic images, preprocessing steps, feature extraction, classification, and the display of results.Finally, an activity diagram represents the workflow of the system, capturing the step-by-step process from image acquisition to result interpretation, including decision points for classification outcomes.The system design ensures modularity, scalability, and efficiency, making it adaptable for deployment in real-world healthcare environments. By combining effective preprocessing, a robust CNN model, and structured workflow design, the system achieves a balance between accuracy and usability.

5. Implementation

The implementation of the proposed skin cancer detection system follows a structured pipeline consisting of data collection, preprocessing, model development, training, evaluation, and prediction. Each stage is designed to ensure that the Convolutional Neural Network (CNN) effectively learns the features of skin lesions and achieves reliable classification performance.

The first step in the implementation is data collection. For this research, the HAM10000 dataset, a widely used collection of dermoscopic images, was employed. This dataset contains diverse examples of seven types of skin lesions, providing an adequate foundation for training and testing the proposed CNN model. However, the dataset is imbalanced, with certain classes having significantly more samples than others. To address this imbalance, preprocessing and data augmentation techniques were applied to artificially expand the dataset and improve generalization.

Preprocessing plays a vital role in preparing images for analysis. Each image was resized to a uniform dimension (for example, 224 × 224 pixels) to ensure consistency across the dataset. Normalization was applied to scale pixel values between 0 and 1, which improves training efficiency and stability. Data augmentation techniques such as random rotation, flipping, zooming, and shifting were employed to increase variability in the dataset and reduce the risk of overfitting. These preprocessing strategies help the CNN adapt to real-world conditions where images may differ in orientation, brightness, or resolution.

The core of the implementation is the CNN model. The network architecture consists of multiple convolutional layers, each followed by pooling operations to capture spatial hierarchies in the dermoscopic images. Deeper layers extract increasingly complex features, which are subsequently fed into fully connected dense layers responsible for final classification. The ReLU activation function was used in convolutional and dense layers to introduce non-linearity, while the output layer used a softmax activation function for multi-class classification. To prevent overfitting, dropout layers were incorporated, randomly deactivating neurons during training.

Training the model involved feeding preprocessed images into the CNN and optimizing its weights using backpropagation. The Adam optimizer was initially employed, with binary and categorical cross-entropy serving as the loss functions depending on the classification type. The dataset was split into training, validation, and testing subsets to evaluate performance at each stage. Early stopping and learning rate scheduling techniques were introduced to prevent overfitting and improve convergence speed.

Evaluation was conducted using metrics such as accuracy, precision, recall, and F1-score to assess the classification performance. The CNN achieved promising results, initially producing an accuracy of 70% with the baseline architecture. To enhance performance, transfer learning was implemented using ResNet50, a pre-trained deep neural network that has demonstrated strong performance in image classification tasks. With this modification, the system achieved a significant improvement, reaching an accuracy of 92%. This demonstrates the effectiveness of using advanced architectures in medical image analysis.

Finally, the trained model was applied for prediction on unseen dermoscopic images. New images were preprocessed in the same manner as the training data and passed through the model to generate classification outputs. The system successfully identified whether a lesion was malignant or benign, along with its specific category. Predictions were generated in real time, making the system suitable for practical integration into clinical workflows and mobile health applications.

The implementation demonstrates that CNN-based architectures, when combined with appropriate preprocessing and augmentation strategies, can significantly improve the accuracy of skin cancer detection. Moreover, the use of advanced transfer learning models such as ResNet50 highlights the potential of leveraging pre-trained architectures for achieving state-of-the-art performance in medical diagnostics.

6. Results

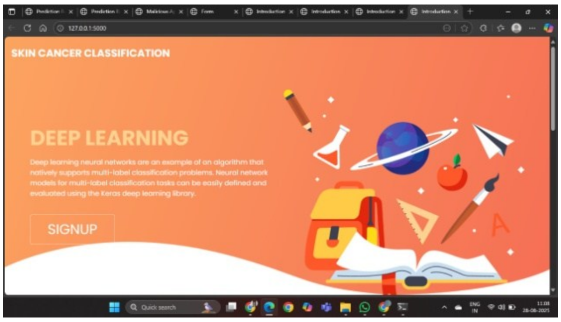

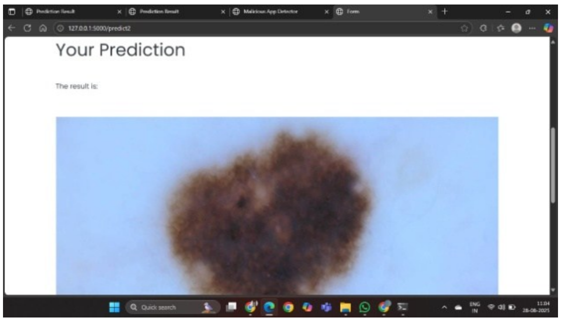

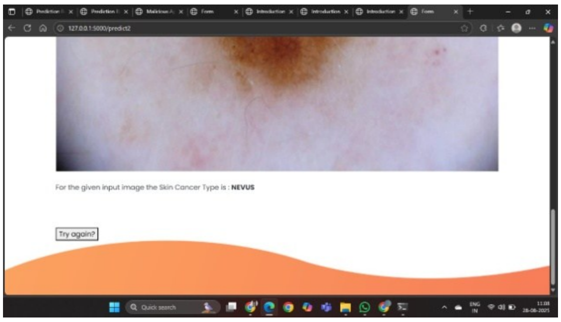

7. Conclusion and Future Work

This research presented a deep learning–based system for skin cancer detection using Convolutional Neural Networks (CNNs), with an emphasis on automating the classification of dermoscopic images into seven distinct categories of skin lesions. The study demonstrated that CNNs, when combined with appropriate preprocessing, augmentation, and transfer learning techniques, can achieve a high level of accuracy in detecting and classifying skin cancer. The baseline CNN model provided an accuracy of 70%, while the implementation of ResNet50 significantly improved performance, reaching an accuracy of 92%. These results highlight the potential of deep learning to outperform traditional diagnostic methods and existing Computer-Aided Diagnosis (CAD) systems in both accuracy and efficiency.

The proposed system addresses key limitations of current diagnostic practices by offering a non-invasive, scalable, and real-time approach to skin cancer detection. It reduces dependency on biopsies and minimizes the subjectivity associated with manual examination. Furthermore, by leveraging cloud and mobile integration, the system has the potential to extend dermatological screening to remote and underserved regions where access to expert care is limited. This aligns with the broader vision of using artificial intelligence to democratize healthcare and improve global patient outcomes.

While the system achieved promising results, certain challenges remain. Dataset imbalance continues to affect classification accuracy for rare lesion types, and further work is needed to collect larger, more balanced datasets that represent diverse skin tones and demographics. Additionally, the “black-box” nature of CNNs poses challenges in clinical acceptance, as medical practitioners often seek interpretable and transparent diagnostic tools. Incorporating explainable AI techniques into the system could enhance trust and adoption among healthcare professionals.

Future work will focus on several key directions. First, advanced architectures such as EfficientNet and Vision Transformers (ViTs) could be explored to further boost classification performance. Second, lesion segmentation can be integrated into the system to highlight regions of interest, improving both accuracy and interpretability. Third, clinical validation through collaboration with dermatologists and medical institutions is essential to ensure the reliability of the system in real-world scenarios. Finally, the system could be extended into a patient-facing mobile application, enabling individuals to perform preliminary self-screening and seek timely medical advice.

In conclusion, this research demonstrates that deep learning offers a transformative solution for skin cancer detection. By bridging the gap between medical imaging and artificial intelligence, the proposed system represents a step toward early diagnosis, better patient care, and the integration of AI-driven tools in modern healthcare.

8. References

[1] J. Ferlay, et al., “Estimating the global cancer incidence and mortality in 2018: GLOBOCAN sources and methods,” International Journal of Cancer, vol. 144, no. 8, pp. 1941–1953, 2019.
[2] G. Kasinathan, et al., “Automated 3-D lung tumor detection and classification by an active contour model and CNN classifier,” Expert Systems with Applications, vol. 134, pp. 112–119, 2019.
[3] Z. Gao, et al., “HEp-2 cell image classification with deep convolutional neural networks,” IEEE Journal of Biomedical and Health Informatics, vol. 21, no. 2, pp. 416–428, 2016.
[4] P. Wang, et al., “Automatic cell nuclei segmentation and classification of cervical Pap smear images,” Biomedical Signal

Processing and Control, vol. 48, pp. 93–103, 2019.

[5] S. Sharma, S. Maheshwari, and A. Shukla, "An intelligible deep convolution neural network-based approach for classification of diabetic retinopathy," BioAlgorithms and Med-Systems, vol. 14, no. 2, 2018.

[6] K. M. Hosny, M. A. Kassem, and M. M. Foaud, "Classification of skin lesions using transfer learning and augmentation with Alex-Net," PLOS ONE, vol. 14, no. 5, e0217293, 2019.

[7] X. He, et al., "Dense deconvolution net: Multi-path fusion and dense deconvolution for high-resolution skin lesion segmentation," Technology and Health Care, vol. 26, no. S1, pp. 307–316, 2018.

[8] B. Harangi, "Skin lesion classification with ensembles of deep convolutional neural networks," Journal of Biomedical Informatics, vol. 86, pp. 25–32, 2018.

[9] T. J. Brinker, et al., "Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task," European Journal of Cancer, vol. 113, pp. 47–54, 2019.

[10] T. J. Brinker, et al., "Skin cancer classification using convolutional neural networks: systematic review," Journal of Medical Internet Research, vol. 20, no. 10, e11936, 2018.