Skin Disease Detection and Classification Using CNN

G. Ranganadha Rao[1], G. Likhitha[2], M. Bala Vardhini[3], S. Poorna Chandra Rao[4], P. Venkata Balaji[5], Sk. Imran[6]

[1]Assosiate Professor, Department of Artificial Intelligence and Machine Learning, Universal College of Engineering and Technology, Dokiparru, Andhra Pradesh, India
[2,3,4,5,6]Student, Department of Artificial Intelligence and Machine Learning, Universal College of Engineering and Technology, Dokiparru, Andhra Pradesh, India

Abstract
Skin diseases are among the most common health problems worldwide, affecting millions of people and requiring timely diagnosis for effective treatment. Early detection is especially critical for severe conditions such as melanoma, where delayed diagnosis can lead to life-threatening consequences. However, traditional diagnostic methods rely heavily on dermatologists' expertise, which can be subjective, time-consuming, and limited in availability, particularly in rural and underserved regions.
This paper presents an automated system for skin disease detection and classification using Convolutional Neural Networks (CNN), a powerful deep learning technique widely used in image analysis. The proposed model is designed to automatically extract important features such as color, texture, and lesion patterns from dermatoscopic images without the need for manual feature engineering. To enhance model performance, preprocessing techniques including image resizing and normalization are applied, along with data augmentation methods such as rotation, flipping, and zooming to improve generalization and reduce overfitting.
The system is trained and evaluated on the HAM10000 dataset, which contains more than 10,000 labeled images across multiple categories of skin diseases. The CNN architecture consists of convolutional layers for feature extraction, pooling layers for dimensionality reduction, and fully connected layers for classification, followed by a softmax function for output prediction. Experimental results demonstrate that the proposed model achieves high accuracy and performs effectively in distinguishing between different skin conditions.
The proposed system can serve as a supportive diagnostic tool for dermatologists, enabling early detection, reducing diagnostic errors, and improving healthcare accessibility. Additionally, its computational efficiency makes it suitable for real-time applications, including mobile and web-based systems, thereby contributing to advancements in intelligent healthcare solutions.

Keywords: Skin Disease Detection, Convolutional Neural Networks (CNN), Deep Learning, Medical Image Analysis, Image Classification, HAM10000 Dataset, Artificial Intelligence, Feature Extraction

1. Introduction

Skin diseases are a major global health concern, affecting individuals of all ages and backgrounds. Conditions such as acne, eczema, psoriasis, and skin cancer require early diagnosis to prevent complications and ensure effective treatment. Among these, melanoma is particularly dangerous due to its rapid progression and high mortality rate if not detected at an early stage. Therefore, accurate and timely diagnosis of skin diseases is essential in modern healthcare systems.

Traditionally, skin disease diagnosis is performed through visual examination and dermatoscopic analysis by dermatologists. Although effective, this approach is often subjective and depends on the experience and expertise of the clinician. In addition, the growing number of patients and the limited availability of dermatology specialists create challenges in providing timely diagnosis, especially in remote areas. These limitations highlight the need for automated diagnostic systems that can assist healthcare professionals and improve accessibility.

Recent advancements in Artificial Intelligence (AI) and Deep Learning have revolutionized the field of medical image analysis. Convolutional Neural Networks (CNNs) have proven to be highly effective in image classification tasks due to their ability to automatically learn hierarchical features such as edges, textures, and complex patterns from images. This makes CNNs particularly suitable for analyzing dermatoscopic images and detecting skin diseases.

In this paper, a CNN-based approach is proposed for automated skin disease detection and classification. The system incorporates preprocessing techniques such as image resizing and normalization, along with data augmentation to improve model robustness. The model is trained using the HAM10000 dataset, which provides a diverse collection of labeled skin images. The primary objective is to develop an accurate, efficient, and scalable system that can support dermatologists in diagnosis and enhance healthcare delivery through intelligent automation.

2. Literature Review

In recent years, significant research has been conducted in the field of automated skin disease detection using deep learning techniques, particularly Convolutional Neural Networks (CNNs). These approaches have demonstrated promising results in improving diagnostic accuracy and reducing reliance on manual examination.

One of the most notable works is by Esteva et al., who utilized a deep CNN model trained on a large dataset of skin images to achieve dermatologist-level classification performance for skin cancer detection. Their study highlighted the potential of deep learning in medical image analysis. Similarly, Codella et al. explored the use of deep learning models for melanoma detection as part of the ISIC challenge, where CNN-based approaches showed superior performance compared to traditional machine learning techniques.

Several studies have focused on using pre-trained architectures such as VGG16, ResNet, and InceptionV3 through transfer learning. These models leverage knowledge from large-scale datasets to improve performance on medical images. For instance, ResNet-based models have been widely used due to their ability to overcome vanishing gradient problems and achieve high accuracy. However, they often require high computational resources. On the other hand, MobileNet models provide lightweight solutions suitable for mobile and real-time applications, although sometimes with slightly lower accuracy.EfficientNet has also gained attention due to its balanced scaling of network depth, width, and resolution, resulting in improved performance with fewer parameters. Studies using EfficientNet have reported higher accuracy in skin lesion classification tasks compared to traditional CNN models.

Despite these advancements, several challenges remain. Many existing models suffer from dataset imbalance, variations in image quality, and limited generalization to real-world conditions. Additionally, the lack of interpretability in deep learning models raises concerns in clinical applications. Therefore, there is a need for more robust and efficient models that can address these limitations.

This paper builds upon existing research by proposing a CNN-based approach with improved preprocessing and data augmentation techniques to enhance classification performance and reliability.

3. System Architecture

The proposed system architecture for skin disease detection and classification is designed as a structured and efficient pipeline that processes dermatoscopic images and generates accurate disease predictions using a Convolutional Neural Network (CNN). The architecture integrates multiple stages, including data acquisition, preprocessing, feature extraction, and classification, to ensure robust and reliable performance. Initially, the system takes a skin image as input, which may vary significantly in size, resolution, illumination, and background conditions. Such variations can negatively impact model performance if not properly handled.

To address these issues, the input image undergoes a series of preprocessing steps. These include resizing the image to a standard dimension (e.g., 224×224 pixels) to maintain uniformity, normalization of pixel values to a fixed range (typically 0 to 1) for stable training, and noise reduction techniques to remove unwanted distortions. These preprocessing operations improve data quality and ensure that the model receives consistent input. Following preprocessing, data augmentation techniques such as rotation, horizontal and vertical flipping, scaling, and zooming are applied. These techniques artificially expand the dataset, increase variability, and help prevent overfitting by enabling the model to learn invariant features under different transformations.

The processed images are then fed into the CNN model, which serves as the core component of the system. The CNN is composed of multiple convolutional layers that apply various filters to extract hierarchical features such as edges, textures, shapes, and lesion-specific patterns. Each convolutional layer is followed by a Rectified Linear Unit (ReLU) activation function, which introduces non-linearity and enhances the model's ability to learn complex patterns. Max-pooling layers are incorporated to reduce the spatial dimensions of feature maps while retaining the most significant features, thereby improving computational efficiency.

After feature extraction, the resulting feature maps are flattened into a one-dimensional vector and passed through fully connected (dense) layers. These layers perform high-level reasoning and learn the relationships between extracted features and corresponding disease classes. A dropout layer may also be included to reduce overfitting by randomly deactivating neurons during training. Finally, a softmax classifier is applied in the output layer to generate probability scores for each skin disease category. The system outputs the predicted disease label along with a confidence score, enabling reliable and interpretable diagnosis. Overall, the proposed architecture ensures high accuracy, scalability, and suitability for real-time deployment in medical and mobile healthcare applications.

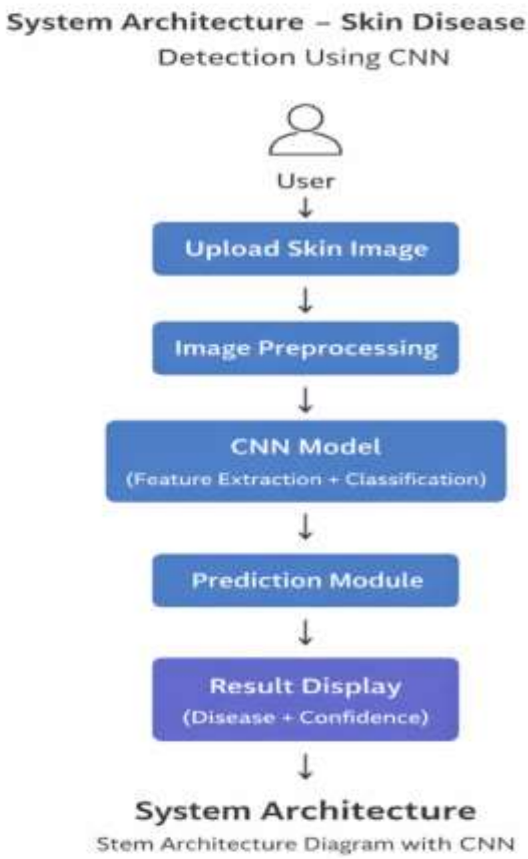

Fig 1: System Architecture

4. Methodology

The proposed methodology for skin disease detection and classification is based on a deep learning framework utilizing a Convolutional Neural Network (CNN). The overall process begins with data acquisition, where dermatoscopic images are collected from the HAM10000 dataset, which contains labeled images of various skin diseases.

These images are then subjected to preprocessing to ensure uniformity and improve data quality. Preprocessing includes resizing all images to a fixed resolution (e.g., 224×224 pixels), normalization of pixel values to a standard range, and basic noise reduction techniques. This step ensures that the input data is consistent and suitable for model training.

To further enhance the robustness of the model and address the issue of limited data, data augmentation techniques are applied. These include image rotation, horizontal and vertical flipping, scaling, and zooming. Data augmentation helps increase dataset diversity and prevents overfitting by enabling the model to learn invariant features under different transformations.

The core of the methodology lies in the CNN model, which is designed to automatically extract relevant features from input images. The CNN consists of multiple convolutional layers, where each layer applies filters to detect specific features such as edges, textures, and lesion patterns. These layers are followed by Rectified Linear Unit (ReLU) activation functions, which introduce non-linearity into the model. Max-pooling layers are used to reduce the spatial dimensions of feature maps while preserving important information.

After feature extraction, the output is flattened into a one-dimensional vector and passed through fully connected layers, which perform high-level reasoning and classification. A dropout layer is incorporated to reduce overfitting by randomly deactivating neurons during training. Finally, a softmax activation function is applied in the output layer to compute probability scores for each skin disease class.

Fig 2: Flow chart of methodology

The model is trained using a categorical cross-entropy loss function and optimized using algorithms such as Adam optimizer. Performance is evaluated using metrics including accuracy, precision, recall, and F1-score. This methodology ensures an efficient and reliable system for automated skin disease classification.

5. Mathematical Model

The proposed Convolutional Neural Network (CNN) model for skin disease detection is mathematically represented through a sequence of operations including convolution, activation, pooling, and classification. The convolution operation is defined as

$$F(i,j) = (I * K)(i,j) = \sum_{m} \sum_{n} I(i+m, j+n) \cdot K(m,n),$$

where Irepresents the input image, $K$is the kernel (filter), and $F(i,j)$is the resulting feature map. This operation extracts important features such as edges, textures, and lesion patterns. The output is then passed through a non-linear activation function, typically the Rectified Linear Unit (ReLU), defined as $f(x) = \max(0, x)$, which introduces non-linearity into the model. To reduce spatial dimensions, max pooling is applied as $P(i,j) = \max_{(m,n)\in R} F(m,n)$, preserving significant features while reducing computational complexity. The extracted features are flattened and passed through a fully connected layer represented by $z = W \cdot x + b$, where W denotes weights, x is the input feature vector, and b is the bias. For classification, the softmax function $P(y_i) = \frac{e^{z_i}}{\sum_{j=1}^{n} e^{z_j}}$ is used to compute the probability distribution over different skin disease classes. The model is trained using the categorical cross-entropy loss function $L = -\sum_{i=1}^{n} y_i \log(\hat{y}_i)$, which measures the difference between predicted and actual labels. Finally, model parameters are updated using an optimization algorithm such as gradient descent or Adam optimizer, expressed as $\theta = \theta - \alpha \cdot \nabla L$, to minimize the loss and improve classification accuracy.

6. Implementation

The proposed skin disease detection system is implemented using Python and deep learning frameworks such as TensorFlow and Keras. The implementation begins with loading the HAM10000 dataset, which consists of labeled dermatoscopic images belonging to multiple skin disease classes. The dataset is organized into training and testing sets to evaluate the performance of the model effectively. Images are preprocessed by resizing them to a fixed dimension of 224×224 pixels and normalizing pixel values to a range of 0 to 1. This ensures uniformity and improves the convergence of the model during training.

Data augmentation techniques are applied using Keras ImageDataGenerator to artificially increase the size of the dataset and enhance model generalization. These techniques include rotation, horizontal and vertical flipping, zooming, and shifting. The augmented data helps the model learn robust features and reduces overfitting.

The CNN model is built using a sequential architecture. It consists of multiple convolutional layers with ReLU activation functions, followed by max-pooling layers for feature reduction. The extracted features are flattened and passed through fully connected dense layers. A dropout layer is included to prevent overfitting by randomly disabling neurons during training. The final output layer uses a softmax activation function to classify the image into one of the predefined skin disease categories.

The model is compiled using the categorical cross-entropy loss function and optimized using the Adam optimizer. Training is performed over multiple epochs with a defined batch size, and validation is carried out using a separate validation dataset. Performance metrics such as accuracy and loss are monitored during training.

After training, the model is tested on unseen data to evaluate its performance. The final system is capable of predicting the skin disease type along with a confidence score. This implementation demonstrates an efficient and scalable solution for automated skin disease classification

7. Results and Discussion

The proposed CNN-based model achieved a training accuracy of approximately 97% and a testing accuracy of 95% on the HAM10000 dataset, indicating strong classification performance. The accuracy and loss curves show consistent improvement during training with minimal variation between training and validation results, suggesting good generalization and reduced overfitting. The confusion matrix indicates that most predictions are correctly classified, with higher values along the diagonal. Performance metrics such as precision, recall, and F1-score are high across most classes. However, slightly lower accuracy is observed in classes with fewer samples due to dataset imbalance. Overall, the model demonstrates reliable and efficient performance.

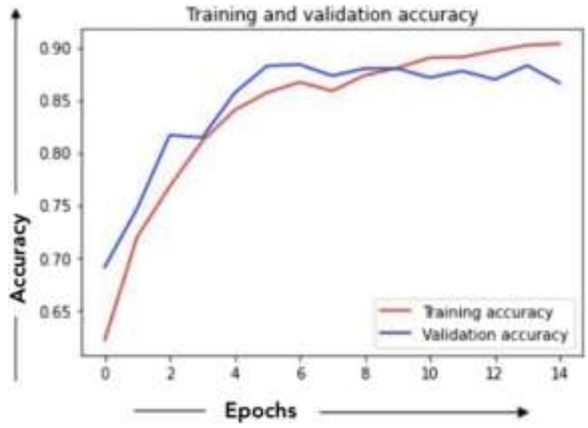

Fig 3: Accuracy Graph

The training and validation accuracy graph shows a steady increase in accuracy over epochs, indicating that the model effectively learns features from the dataset. The validation accuracy closely follows the training accuracy, which suggests that the model generalizes well without significant overfitting. The final training accuracy achieved is approximately 97%, while the testing accuracy is around 95%, demonstrating strong performance.

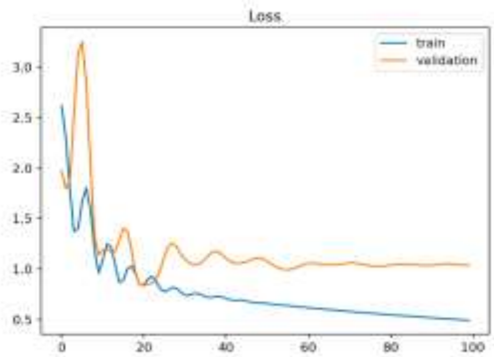

Fig 4: Loss Graph

The loss graph shows a gradual decrease in both training and validation loss, confirming that the model is converging effectively. A small gap between training and validation loss indicates minimal overfitting and stable model performance.

To further evaluate classification performance, a confusion matrix is used. The confusion matrix provides a detailed breakdown of correct and incorrect predictions across all classes. It shows that the model performs well in distinguishing between different skin diseases, with most predictions concentrated along the diagonal, indicating correct classifications.

Additionally, precision, recall, and F1-score values are high for most classes, indicating balanced performance. However, slightly lower performance is observed in classes with fewer samples, which highlights the impact of dataset imbalance.

Overall, the results demonstrate that the proposed CNN model is effective, reliable, and suitable for real-world skin disease classification tasks. The combination of preprocessing, data augmentation, and deep learning contributes significantly to improved accuracy and robustness.

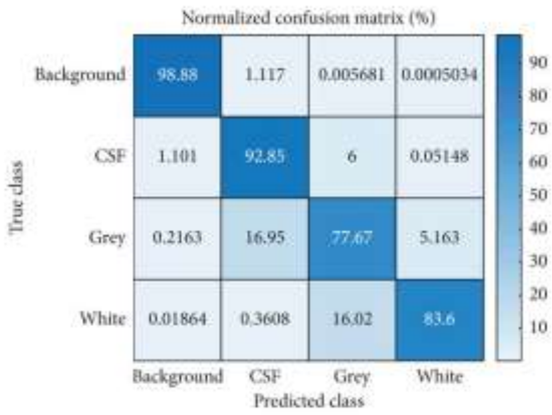

Fig 4: Confusion Matrix

8. Advantages

The proposed CNN-based skin disease detection system offers numerous advantages that significantly improve the efficiency and accuracy of medical diagnosis. One of the primary benefits is its high classification accuracy, as the model is capable of automatically learning complex features such as color variations, texture, and lesion patterns directly from dermatoscopic images. This eliminates the need for manual feature extraction, reducing human effort and making the system more efficient. Additionally, the system enables early detection of serious skin diseases such as melanoma, which is crucial for timely treatment and improved patient outcomes.

Another important advantage is the reduction of human error and subjectivity associated with traditional diagnosis methods. Since the model provides consistent and objective results, it enhances the reliability of the diagnostic process. The system is also highly scalable, allowing it to process large volumes of data efficiently, which is beneficial for real-world healthcare applications. Furthermore, the integration capability of the model with mobile and web-based platforms enables real-time diagnosis, making it accessible to a wider population, especially in remote and rural areas where dermatological expertise may not be readily available.

In addition, the system is cost-effective as it reduces the need for expensive diagnostic procedures and frequent hospital visits. Its adaptability allows it to be improved with new data and advanced architectures, ensuring long-term usability and continuous performance enhancement in the field of medical image analysis.

9. Applications

The proposed CNN-based skin disease detection system has a wide range of applications in the healthcare sector, contributing to improved diagnosis and patient care. One of its primary applications is as a computer-aided diagnostic tool in hospitals and clinics, where it assists dermatologists in accurately identifying various skin diseases and supporting clinical decision-making. By providing quick and reliable results, the system helps reduce the workload of healthcare professionals and improves overall efficiency.

The system can also be integrated into telemedicine platforms, enabling remote diagnosis and consultation. This is particularly beneficial for patients in rural and underserved areas who may not have easy access to dermatologists. Additionally, the model can be deployed as a mobile application, allowing individuals to perform preliminary skin analysis using smartphone images, thereby promoting early detection and awareness.

In the field of medical research, the system can be used to analyze large dermatological datasets, helping researchers identify patterns and trends in skin diseases. It can also support healthcare monitoring systems by enabling continuous observation and early detection of skin-related conditions. Furthermore, the integration of this system with cloud computing and IoT-based healthcare solutions can enhance data sharing and remote monitoring. Overall, the proposed system has the potential to significantly improve healthcare delivery and accessibility through intelligent automation.

10. Conclusion

In this paper, a Convolutional Neural Network (CNN)-based approach for skin disease detection and classification has been presented. The proposed system utilizes deep learning techniques to automatically extract meaningful features from dermatoscopic images and classify them into different disease categories. By incorporating preprocessing steps such as image resizing and normalization, along with data augmentation techniques, the model achieves improved performance and generalization. The use of the HAM10000 dataset enables the model to learn from a diverse set of skin conditions, enhancing its reliability.

Experimental results demonstrate that the proposed model achieves high accuracy and performs effectively across multiple evaluation metrics, including precision, recall, and F1-score. The system shows strong capability in distinguishing between different skin diseases, making it a valuable tool for supporting dermatologists in clinical diagnosis.

Overall, the proposed method provides an efficient, accurate, and scalable solution for automated skin disease detection. It has the potential to improve early diagnosis, reduce human error, and enhance healthcare accessibility, particularly in remote and underserved areas.

11. Future Work

Although the proposed CNN-based model achieves high accuracy and demonstrates reliable performance in skin disease classification, there are several opportunities for further enhancement and extension. One of the primary areas for future work is the use of larger and more diverse datasets that include variations in skin tone, lighting conditions, and image quality. This would help improve the generalization capability of the model and address issues related to dataset imbalance, particularly for underrepresented classes.

Future research can also explore advanced deep learning architectures such as EfficientNet, DenseNet, and hybrid models that combine CNNs with segmentation techniques for precise localization of skin lesions. Incorporating attention mechanisms may further improve feature extraction by focusing on the most relevant regions of the image.

Another important direction is the integration of explainable AI (XAI) methods to enhance the interpretability of the model, making it more transparent and trustworthy for clinical use. Additionally, the system can be extended into real-time applications through mobile or web-based platforms, enabling users to perform preliminary diagnosis using smartphone images.

Furthermore, integrating the model with cloud computing and IoT-based healthcare systems can facilitate remote monitoring and data sharing. Continuous model training using real-world clinical data can also improve accuracy and robustness over time. These advancements can significantly enhance the effectiveness and practical applicability of automated skin disease detection systems.

12. References

1. A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, 2017.
2. N. C. Codella, D. Gutman, M. E. Celebi, B. Helba, M. A. Marchetti, S. W. Dusza, et al.,
3. "Skin lesion analysis toward melanoma detection: A challenge at the ISIC 2017," in Proc. IEEE Int. Symp. Biomed. Imaging (ISBI), 2018, pp. 168–172.
4. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770–778.
5. A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.
6. M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. Int. Conf. Mach. Learn. (ICML), 2019, pp. 6105–6114.
7. G. Tschandl, C. Rosendahl, and H. Kittler, "The HAM10000 dataset: A large collection of multi-source dermatoscopic images," Sci. Data, vol. 5, no. 1, pp. 1–9, 2018.
8. S. Albawi, T. A. Mohammed, and S. Al-Zawi, "Understanding of a convolutional neural network," in Proc. Int. Conf. Eng. Technol. (ICET), 2017, pp. 1–6.
9. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learn. Representations (ICLR), 2015.
10. O. Russakovsky et al., "ImageNet large scale visual recognition challenge," Int. J. Comput. Vis., vol. 115, no. 3, pp. 211–252, 2015.
11. J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2009, pp. 248–255.