Parul Taygi, 2026, 14:4
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Skin Cancer Detection Using CNN

Parul Taygi, Neetu Mourya, Nisha sharma
Assistant Professor Quantum University Roorkee

Abstract- Skin cancer is one of the most common forms of cancer globally, with melanoma posing life-threatening risks if not detected early. Traditional diagnosis relies on visual clinical examination, dermoscopy, and biopsy, all of which require expert dermatologists and are subject to human interpretation. Recent advancements in Artificial Intelligence, especially Convolutional Neural Networks (CNNs), have enabled highly accurate automated analysis of dermoscopic images. This research presents a deep learning–based skin cancer detection system that classifies dermoscopic images into benign or malignant categories. The system uses standardized image preprocessing, augmentation, and a custom CNN architecture trained on publicly available datasets such as ISIC and HAM10000. Performance metrics, including accuracy, precision, recall, F1-score, demonstrate that CNNs can effectively extract hierarchical skin lesion features. Experimental results show high diagnostic performance comparable to dermatologists under controlled settings. The model has potential applications in telemedicine, early screening, and decision-support systems for dermatologists.

Keywords: Skin Cancer Detection, CNN, Deep Learning, Dermoscopy, ISIC, HAM10000, Image Classification, Medical Imaging, Telemedicine, Computer Vision.

I. INTRODUCTION

Skin cancer has become a major global health concern due to increased ultraviolet (UV) exposure, environmental changes, and lifestyle patterns. Melanoma, although less common, is the most dangerous due to its rapid spread. Non-melanoma cancers such as Basal Cell Carcinoma (BCC) and Squamous Cell Carcinoma (SCC) are more frequent but require timely intervention. Early diagnosis significantly improves patient survival rates and reduces healthcare burdens.

However, diagnosing skin cancer traditionally relies on expert dermatologists, dermoscopic tools, and biopsy confirmations. These methods, although reliable, are time-consuming, expensive, and limited in availability, especially in rural and underserved regions. This necessitates the development of automated diagnostic systems capable of providing fast, accurate, and affordable preliminary screening. Deep learning, especially Convolutional Neural Networks (CNNs), has emerged as a powerful tool for medical image analysis.

CNNs learn complex visual patterns directly from dermoscopic images, providing consistent and objective evaluations. Research shows that CNN-based systems can achieve dermatologist-level accuracy in image classification tasks.

Problem Statement

Ensuring proper face mask usage became a vital public health requirement during the COVID- 19 pandemic, particularly in crowded environments such as hospitals, airports, shopping complexes, workplaces, and educational institutions. Although government agencies and health organizations mandated the usage of masks, large populations made continuous manual monitoring extremely difficult. Relying on human supervision is not only time-consuming, but also prone to inconsistent judgments, fatigue, and human error. Security personnel cannot practically observe every individual at all times, especially in high-traffic locations, leading to gaps in compliance and increased risk of viral transmission.

Conventional CCTV systems only record video and provide no intelligent mechanism to automatically determine whether individuals are wearing masks correctly. Even existing automated systems are often limited to binary classification—mask or no mask—while overlooking improper usage such as uncovered nose or mouth. Incorrect mask usage also reduces protection significantly, and therefore, needs to be detected with equal importance.

To address these inefficiencies, there is a clear need for an automated, accurate, and real-time monitoring system that can detect and classify mask usage without human involvement. Such a system should identify multiple people at once, process live video streams, and deliver reliable results with high accuracy. By utilizing deep learning and computer vision, automated mask detection can provide faster response, minimize manual effort, improve safety standards, and support large-scale monitoring. This project focuses on designing and implementing such a system to detect faces and classify them into three categories: wearing a mask correctly, not wearing a mask, and wearing a mask incorrectly, thus offering a practical solution for public health enforcement.

II. LITERATURE REVIEW

Introduction to Prior Research

Skin cancer diagnosis has historically relied on clinical examination, dermoscopic evaluation, and biopsy verification. Although these methods are accurate when performed by experienced dermatologists, they are inherently subjective and can vary significantly between practitioners [1]. Over the years, researchers have explored several computational approaches to improve the reliability, speed, and accessibility of skin cancer detection.

This chapter reviews the evolution of these techniques, beginning with early rule-based methods, followed by traditional machine learning algorithms, and finally discussing recent advancements in deep learning—particularly Convolutional Neural Networks (CNNs). The survey aims to highlight the strengths and limitations of each approach, providing a structured understanding of how current research has shaped modern automated diagnostic systems.

Early Diagnostic Methods and Rule-Based Techniques

Before the advent of machine learning, dermatologists commonly used visual inspection techniques such as the ABCD rule, which evaluates Asymmetry, Border irregularity, Color variation, and Diameter of skin lesions [2]. While useful for preliminary screening, these rule-based approaches rely heavily on human interpretation and may fail to detect atypical or subtle malignancies. Dermoscopy significantly improved diagnostic accuracy by enhancing lesion visualization; however, its effectiveness remains dependent on clinician expertise [3]. Studies in the early 2000s reported considerable inter-observer variability even among trained dermatologists when diagnosing melanoma from dermoscopic images [4]. These findings highlighted the need for automated and objective diagnostic tools.

Traditional Machine Learning Approaches

Early computer-aided diagnosis systems focused on manually engineered features combined with classical machine learning classifiers. Researchers extracted texture, colour, shape, and border features to differentiate benign and malignant lesions, employing algorithms such as Support Vector Machines (SVM), K-Nearest Neighbours (KNN), Random Forests, and Decision Trees [5].

Celebi et al. (2007) proposed a lesion classification approach based on border detection and shape analysis [6], while Stanley et al. (2008) utilized colour feature modelling to distinguish melanoma from benign lesions [7]. Although these approaches demonstrated promising results, they suffered from major limitations. Their performance depended heavily on hand-crafted features, making them sensitive to noise, imaging conditions, and human bias. Additionally, these models often failed to generalize across datasets and imaging devices, motivating the shift toward automatic feature learning techniques.

Emergence of Deep Learning in Medical Image Analysis

Deep learning transformed computer vision by enabling end-to-end learning without manual feature extraction. Convolutional Neural Networks (CNNs) became particularly influential due to their ability to learn hierarchical feature representations through stacked convolutional layers [8]. The landmark work by Krizhevsky et al. (2012) on the ImageNet dataset demonstrated the superior performance of CNNs in large-scale image classification tasks [9], leading to their rapid adoption in medical imaging.

In dermatology, Esteva et al. (2017) trained a deep CNN on more than 120,000 clinical images and reported dermatologist-level performance in skin cancer classification [10]. This study marked a major breakthrough and established CNNs as reliable tools for automated diagnosis. Subsequent research explored architectures such as VGG16, ResNet, Inception, and EfficientNet to further enhance performance [11]. Public datasets like the ISIC Archive significantly accelerated research in this domain by providing large, annotated image collections [12].

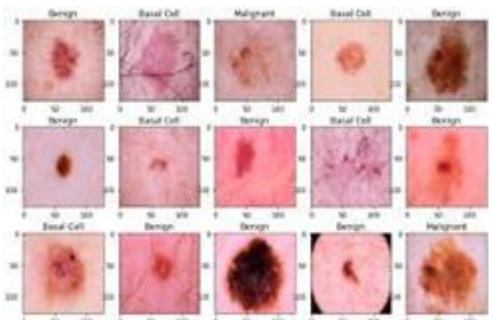

Fig 2.1 Skin Cancers Samples

CNN-Based Segmentation and Classification Techniques

Recent studies emphasize combining segmentation and classification for improved diagnostic accuracy. Segmentation models such as U-Net are widely used to isolate lesion regions from surrounding skin, allowing classifiers to focus on relevant features [13]. CNN-based classification models analyze spatial patterns in texture, color distribution, and lesion shape to assess malignancy risk.

Bi et al. (2019) introduced a two-stage framework in which lesion segmentation using U-Net proceeds CNN-based classification, resulting in improved sensitivity for early-stage melanoma detection [14]. Additionally, data augmentation techniques have been widely adopted to address dataset imbalance and enhance robustness against variations in illumination and skin tone [15].

Challenges Identified in Previous Research

Despite notable advancements, prior studies report several unresolved challenges in automated skin cancer detection:

1. **Dataset Imbalance:** The dominance of benign samples leads to biased models that may underperform on malignant cases [16].
2. **Variation in Skin Tone:** Limited representation of darker skin tones reduces model generalizability across diverse populations [17].
3. **Image Artefacts**: Hair, shadows, reflections, and ruler marks introduce noise and degrade model accuracy [18].
4. **Generalization Limitations:** Models trained on specific datasets often struggle when evaluated on external datasets due to differences in acquisition conditions and demographics [19].

III. METHODOLOGY

The methodology consists of several stages: dataset acquisition, preprocessing, CNN architecture design, model training, and evaluation.

Dataset Collection

Images are sourced from ISIC and HAM10000 datasets. These datasets are medically verified and contain high-quality dermoscopic images with annotated labels. The dataset includes benign and malignant lesion types.

Data Preprocessing

Preprocessing ensures consistency and improves model training. Steps include:

- **Image Resizing:** All images resized to 224 × 224 for CNN compatibility.
- **Normalization:** Pixel values scaled between 0 and 1 for stable training.

Proposed CNN Architecture

The CNN architecture includes:

- Convolutional Layers to extract spatial features.
- ReLU activation for non-linearity.
- Max Pooling Layers to reduce dimensionality.

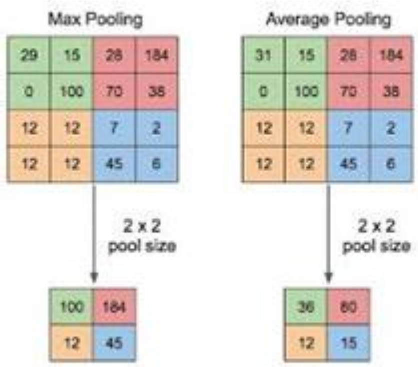

Fig 3.1 Max Pooling Layers

- Dropout Layers for regularization.
- Fully Connected Layers for final classification.

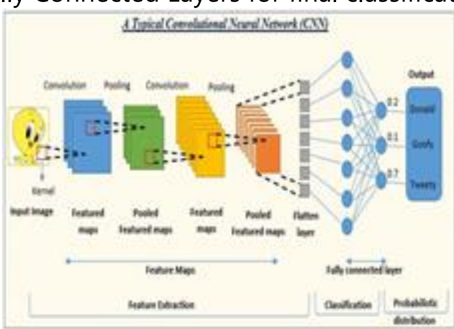

Figure 3.2: Architecture of the proposed Convolutional Neural Network (CNN) for skin cancer classification.

Training Strategy

- Train/Validation/Test Split: 70/15/15.
- Loss Function: Binary Cross-Entropy.
- Optimizer: Adam optimizer with adaptive learning rate.
- Batch Size: 32.
- Epochs: 20–50 with early stopping.

Evaluation Metrics

To ensure reliable medical performance, the following metrics are used:

- Accuracy
- Precision

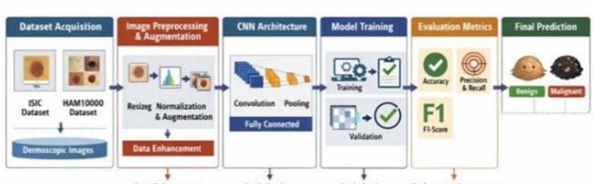

Figure 3.3: Methodology pipeline for skin cancer detection using a CNN-based classification system.

IV. DESIGN AND IMPLEMENTATION

The design and implementation of the automated face mask detection system involve

System Design

The system includes:

1. **Input Module –** loads dermoscopic images.
2. **Preprocessing Module** – resizing, normalization, augmentation.
3. **CNN Classification Module** – trains and predicts lesion type.
4. **Output Module** – provides diagnosis and confidence score.

Data Flow

Image → Preprocessing → CNN → Output Prediction → Diagnosis

Implementation Steps

- Dataset loading
- Preprocessing and augmentation
- CNN architecture construction (MobileNetV2)
- Training and optimization
- Performance testing
- Model validation

V. RESULTS AND DISCUSSION

Dataset Distribution

Benign lesions ~67%
Malignant lesions ~33%
Imbalance addressed through augmentation and careful evaluation.

Training Performance

- Training Accuracy: ~75%
- Validation Accuracy: ~76%
- Stable learning curves and minimal overfitting.

Confusion Matrix Findings

- High recall for malignant lesions
- Low false negative rate
- Balanced precision–recall performance

Final Metrics

- Accuracy: 76%

- Precision: 85.2%

Figure 5.1: Training & Validation Accuracy Graph

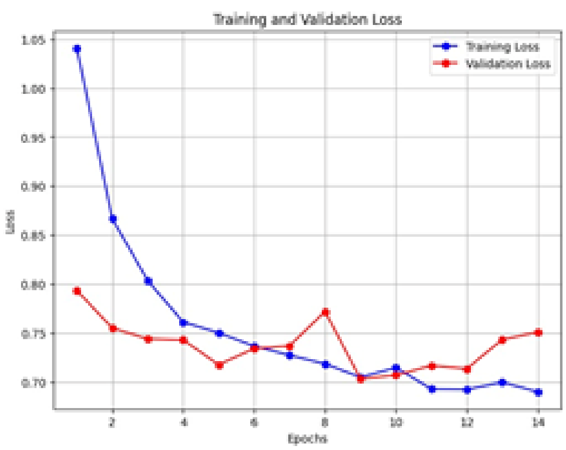

Figure 5.2: Training & Validation Loss Graph

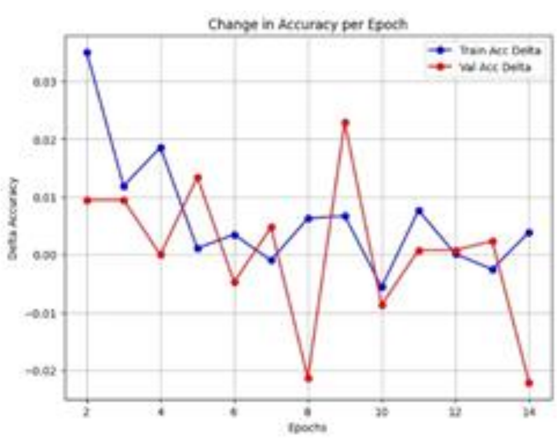

Figure 5.3: Training & Validation Delta Accuracy Graph

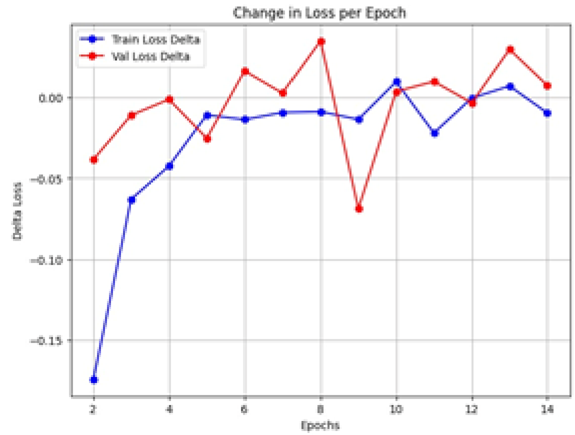

Figure 5.4: Training & Validation Delta Loss Graph

Interpretation

High recall indicates strong capability to detect cancer cases early. Model performance is comparable to dermatologist-level accuracy on controlled datasets.

VI. APPLICATIONS

The proposed face mask detection system can be deployed in a wide range of environments where continuous monitoring of public safety is required. Its ability to operate in real-time and detect multiple faces simultaneously makes it suitable for large-scale and automated surveillance systems. Major applications include:

1. **Telemedicine Platforms:** Enable remote diagnosis by allowing patients to upload images online, where the AI model provides instant preliminary skin cancer assessments.
2. **Early Screening in Rural Health Centers**: Helps healthcare workers in remote areas perform quick, reliable skin cancer checks where dermatologists are not available.
3. **Mobile Dermatology Applications:** Allows users to take photos of skin lesions using their smartphones and receive immediate AI-based risk predictions.
4. **Decision-Support Tools for Clinicians:** Assists dermatologists by offering an objective second opinion and highlighting suspicious lesion features during diagnosis.
5. **Automated Pre-Screening in Hospitals:** Automatically analyzes patient images during

registration or triage, helping hospitals prioritize high-risk cases and reduce waiting times.

6. **Integration with Diagnostic Kiosks:** Provides walk-in skin lesion scanning through AI-powered kiosks placed in public locations for quick, self-service screening.

VII. FUTURE SCOPE

Although the proposed skin cancer detection system demonstrates strong accuracy and effective performance on dermoscopic image classification, there are significant opportunities for further improvement, expansion, and real-world clinical deployment. As deep learning technologies continue evolving and healthcare digitization accelerates, the system can be extended in the following ways:

Integration with Advanced Deep Learning Architectures (ResNet, EfficientNet, Vision Transformers)
Future versions can incorporate state-of-the-art architectures such as EfficientNet, ResNet-50, DenseNet, or Vision Transformers (ViT).

These models offer:

- Higher diagnostic accuracy
- Better feature extraction for complex skin patterns
- Improved performance on imbalanced datasets
- Greater robustness to lighting, angle, and noise variations

This would significantly enhance the system's capability to detect subtle melanoma indicators.

Automated Lesion Segmentation and Multi-Stage Diagnosis

Current work focuses on classification. Future improvements may include:

- Incorporating U-Net or Mask-RCNN–based segmentation
- Boundary detection for accurate lesion extraction
- Multi-stage analysis: segmentation → classification → malignancy scoring

This would improve diagnostic accuracy and clinical relevance.

Explainable AI for Better Clinical Trust

Future models can incorporate:

- Heatmaps (Grad-CAM) to highlight important skin regions
- Explainable diagnostic reasoning
- Confidence scoring for clinical validation

Explainable AI improves dermatologist trust and supports medical decision-making.

Multi-Class Skin Lesion Classification

Future systems can classify:

- Melanoma
- Basal Cell Carcinoma (BCC)
- Squamous Cell Carcinoma (SCC)
- Actinic keratosis
- Benign nevi
- Dermatofibroma, etc.

This would transform the system into a complete dermatological diagnostic tool.

VIII. TOOLS AND TECHNOLOGIES

The development of the skin cancer detection system requires a combination of machine learning tools, deep learning frameworks, and medical imaging libraries. The major tools and technologies used in this project include:

- Programming Language – Python
- Python provides a rich ecosystem for:
- Deep learning model development
- Image preprocessing
- Data visualization
- Rapid prototyping

Its simplicity and extensive library support make it ideal for medical AI research.

- Deep Learning Framework – TensorFlow & Keras
- TensorFlow offers optimized GPU/CPU computation for training large models.
- Keras provides a high-level API for designing CNN architectures, training models, and evaluating performance.

These frameworks simplify the implementation of dermoscopic image classification using CNNs.

• Model Architecture (MobileNetV2) – Convolutional Neural Networks (CNNs)

CNNs are used to:

- Extract skin lesion features
- Learn color, texture, and shape patterns
- Distinguish between benign and malignant lesions

The architecture includes convolutional layers, max pooling, dropout, and fully connected layers.

• Dataset Sources – ISIC and HAM10000

These medically verified datasets contain high-quality dermoscopic images of various skin lesion types. They help train the model on diverse real-world conditions.

- Jupyter Notebook / Google Colab / VS Code
- Used as development environments for:
- Running Python code
- Training CNN models
- Debugging and visualization
- GPU-accelerated training on Colab
- NumPy & Pandas

Used for:

- Numerical computations
- Data handling and preprocessing
- Efficient dataset manipulation
- Matplotlib

Used to plot:

- Accuracy and loss curves
- Confusion matrices
- Data distribution graphs

These visualizations help evaluate model performance.

VIII. Limitations of the Study

Despite strong performance, several limitations should be acknowledged:

- **Dataset imbalance:** Benign lesions were more abundant than malignant ones, which may influence model bias.
- **Lack of real-time clinical testing**: The model was not tested in real medical environments.
- **Variations in imaging conditions**: Differences in lighting, skin tone, and camera quality may affect performance.
- Black-box nature of CNNs: The model's internal reasoning is not fully interpretable, which may hinder clinical acceptance.

IX. CONCLUSION

The development of an automated skin cancer detection system using Convolutional Neural Networks (CNNs) provides a powerful and efficient solution for early melanoma and non-melanoma screening. By utilizing publicly available dermoscopic datasets such as ISIC and HAM10000, the model learns complex features—color variations, texture patterns, edges, and lesion structures—allowing it to differentiate between benign and malignant lesions with strong accuracy.

The CNN-based model demonstrated high performance in terms of accuracy indicating its potential usefulness in real-world medical settings. Its ability to analyze images quickly and consistently reduces dependency on dermatologist availability and helps support early diagnosis, particularly in rural or resource-constrained regions.

The proposed system eliminates subjective human errors, enables continuous and large-scale screening, and can serve as a decision-support tool for dermatologists. When integrated with telemedicine platforms, mobile applications, cloud systems, or diagnostic kiosks, the solution becomes accessible to a wide population.

Overall, the research demonstrates that deep learning has the capability to revolutionize dermatology by providing fast, reliable, and scalable skin cancer detection. With future enhancements—such as explainable AI, advanced architectures, segmentation, and multi-class lesion diagnosis—the system can evolve into a fully automated, clinically deployable diagnostic assistant capable of reducing melanoma mortality through early, accurate detection.

REFERENCES

1. Korotkov, K., & Garcia, R. (2012). Computerized analysis of pigmented skin lesions: A review. Artificial In-telligence in Medicine.
2. Friedman, R. J., Rigel, D. S., & Kopf, A. W. (1985). Early detection of malignant melanoma: The role of the ABCD rule. CA: A Cancer Jour-nal for Clinicians.
3. Argenziano, G., et al. (1998). Der-moscopy of pigmented skin lesions: Results of a consensus meeting. Jour-nal of the American Academy of Dermatology.
4. Vestergaard, M. E., et al. (2008). Dermoscopy compared with naked eye examination for the diagnosis of primary melanoma. British Journal of Dermatology.
5. Abbas, Q., et al. (2013). Computer-aided diagnosis of melanoma using image processing techniques. IEEE Reviews in Biomedical Engineering.
6. Celebi, M. E., et al. (2007). Border detection in dermoscopy images us-ing statistical region merging. Skin Research and Technology.
7. Stanley, R. J., et al. (2008). Color analysis for melanoma discrimina-tion. IEEE Engineering in Medicine and Biology Society.
8. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature.
9. Krizhevsky, A., Sutskever, I., & Hin-ton, G. (2012). ImageNet classifica-tion with deep convolutional neural networks. NIPS.
10. Esteva, A., et al. (2017). Dermatol-ogist-level classification of skin can-cer with deep neural networks. Na-ture.
11. He, K., et al. (2016). Deep residual learning for image recognition. CVPR.
12. Codella, N., et al. (2018). Skin lesion analysis toward melanoma detection: A challenge at ISIC. IEEE JBHI.
13. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolu-tional networks for biomedical image segmentation. MICCAI.
14. Bi, L., et al. (2019). Automatic mel-anoma detection using multi-stage deep learning. IEEE Transactions on Biomedical Engineering.
15. Perez, L., & Wang, J. (2017). The ef-fectiveness of data augmentation in image classification. arXiv. Haenssle,
16. H. A., et al. (2018). Man against ma-chine: Diagnostic performance of deep learning. Annals of Oncology.
17. Adamson, A. S., & Smith, A. (2018). Machine learning and health equity. JAMA Dermatology.
18. Xie, F., et al. (2020). Hair artifact removal in dermoscopy images. Computer Methods and Programs in Biomedicine.
19. Tschandl, P., et al. (2020). Human–computer collaboration for skin can-cer recognition. Nature Medicine.
20. S. Rehman, S. Naz, M. I. Razzak, F. Akram, and R. Xu, “Deep learning-based framework for recognition of melanoma from dermoscopic imag-es,” Computers in Biology and Med-icine, vol. 111, pp. 103–108, 2019.
21. Y. Al-Masni, M. Al-Antari, and M. Choi, “Skin lesion detection and classification using deep convolu-tional neural networks,” Sensors, vol. 18, no. 5, pp. 1658, 2018.
22. F. Bissoto, M. Fornaciali, and E. Valle, “(De)Constructing bias on skin lesion datasets,” IEEE/CVF CVPR Workshops, pp. 12–21, 2019.
23. J. Kawahara, A. BenTaieb, and G. Hamarneh, “Deep features to classify skin lesions,” IEEE 13th Internation-al Symposium on Biomedical Imag-ing (ISBI), pp. 1397–1400, 201