Tuijin Jishu/Journal of Propulsion Technology
ISSN: 1001-4055
Vol. 47 No. 02 (2026)

A Cloud-Based Hybrid Deep Learning Framework for Automated Skin Cancer Classification Using the ISIC Dataset

Dr. Sangeetha K[1], Sindhuja R[2], Akshya M[3], Abinaya S[4+,] Dr. Vinmathi S[5], Glory Madona J[6]

[1] Professor, Department of CSE, Panimalar Engineering College Chennai, India
[2,4,6] Department of CSE, Panimalar Engineering College Chennai, India
[3] Assistant Professor, Department of CSE, Panimalar Engineering College Chennai, India
[5]Professor, Department of CSE, Panimalar Engineering College Chennai, India

Abstract: Early identification of malignant skin lesions is critical for reducing mortality associated with skin cancer. Manual examination through dermoscopy followed by biopsy, although clinically reliable, is time-consuming and dependent on specialist availability. This research designs a hybrid neural network architecture aimed at enabling automated large-scale detection of skin cancer through multi-class analysis of dermoscopic images collected from the ISIC archive. The proposed model integrates two complementary convolutional neural network backbones— Xception and ResNet50—initialized with ImageNet weights and configured in a parallel feature extraction architecture. Preprocessed images are resized and normalized before being passed through both networks, where high-level feature representations are extracted and concatenated into a unified descriptor. The combined feature representations are subsequently passed through a series of dense layers integrated with batch normalization, weight regularization, and dropout mechanisms to improve model stability and prevent overfitting. The final classifier assigns each input image to one of four diagnostic categories: Actinic Keratosis, Basal Cell Carcinoma, Melanoma, or Squamous Cell Carcinoma. Performance assessment conducted on 1,866 validation samples indicates stable and reliable predictive outcomes across all lesion classes. A web-based interface developed using Streamlit enables real-time inference, facilitating practical deployment as a clinical decision-support tool. The results indicate that combining heterogeneous CNN architectures improves robustness and discriminative capability in automated dermatological image analysis.

Keywords: *Skin Cancer Classification; Hybrid Deep Learning; Transfer Learning; Xception; ResNet50; ISIC Dataset; Dermoscopic Image Analysis; Feature Fusion; Medical Image Processing; Clinical Decision Support System*

1. Introduction

Skin cancer has emerged as a major worldwide public health issue, with cases rising progressively as a result of extended ultraviolet radiation exposure, environmental conditions, and behavioral risk factors. The disease arises from abnormal proliferation of skin cells and is broadly

categorized into melanoma and non-melanoma types, including Basal Cell Carcinoma (BCC), Squamous Cell Carcinoma (SCC), and Actinic Keratosis (AK). While early-stage lesions are often treatable, delayed identification may lead to severe clinical complications and increased mortality risk. Consequently, reliable early detection mechanisms are essential for effective disease management.

Dermoscopic examination followed by histopathological analysis remains the standard clinical procedure for confirming skin malignancies. Although biopsy provides definitive diagnosis, it is invasive and requires expert dermatological assessment. In many regions, limited access to specialists and increasing patient volume create delays in screening and treatment. These practical limitations highlight the need for automated diagnostic assistance capable of supporting clinicians with rapid and consistent evaluations.

Progress in artificial intelligence has facilitated the creation of intelligent diagnostic systems that rely on data-centric learning approaches for medical image analysis. Among these methods, Convolutional Neural Networks (CNNs) have proven highly effective in automatically learning multi-level visual representations from raw image inputs, eliminating the need for handcrafted feature extraction. Transfer learning strategies further enhance performance by adapting models pretrained on large-scale datasets, thereby reducing training time and improving convergence when working with comparatively smaller medical datasets.

The International Skin Imaging Collaboration (ISIC) repository offers a well-curated and publicly accessible set of dermoscopic images that supports the development and evaluation of automated skin lesion classification systems. However, variability in lesion texture, pigmentation, border irregularity, and structural complexity presents significant challenges for single-network architectures. A

model capable of capturing both fine-grained textural information and high-level structural characteristics is therefore desirable.

To overcome these limitations, this research introduces an integrated hybrid deep learning model that employs Xception and ResNet50 networks in a parallel configuration for simultaneous feature extraction. Xception contributes efficient depthwise separable convolutions that emphasize detailed texture representation, whereas ResNet50 leverages residual connections to enable deeper semantic feature learning. By fusing complementary feature representations from both networks, the proposed approach aims to enhance multi-class discrimination across four lesion categories: AK, BCC, MEL, and SCC.

In addition to model development, this work integrates the trained system into a web-based application to enable real- time inference. This deployment-oriented design ensures that the framework is not limited to experimental validation but can serve as a practical decision-support tool in dermatological screening environments. The main aim of this work is to design a reliable, scalable, and clinically applicable artificial intelligence–based system for automated classification of skin cancer lesions.

2. Related Work

The field of automated skin lesion assessment has progressed considerably alongside the growth of machine learning and deep learning methodologies. Initial computational strategies depended on manually engineered features such as color distribution patterns, texture characteristics, border shape analysis, and morphological measurements to perform lesion classification. The extracted manual features were commonly integrated with traditional machine learning algorithms, including Support Vector Machines (SVM), k- Nearest Neighbors (k-NN), and Random Forest classifiers, to perform lesion categorization. Although such methods provided foundational insights into computer-aided dermatological diagnosis, their performance was limited by feature selection dependency and difficulty in capturing complex lesion variability.

The emergence of deep Convolutional Neural Networks (CNNs) significantly advanced medical image analysis by allowing models to learn discriminative features directly from raw pixel inputs in an end-to-end manner. Pretrained architectures such as VGGNet, Inception, DenseNet, and ResNet have been extensively fine-tuned for dermoscopic image classification tasks. Transfer learning, particularly from large-scale datasets like ImageNet, has proven effective in improving convergence speed and reducing overfitting when dealing with limited annotated medical datasets. Studies utilizing ResNet-based models have demonstrated strong capability in learning hierarchical lesion features, including edge irregularities and structural patterns associated with malignant growth.

More recent investigations have explored ensemble and hybrid modeling strategies to enhance classification robustness. Rather than relying on a single backbone network, ensemble techniques aggregate predictions from multiple architectures to reduce variance and improve generalization. Feature-level fusion approaches have gained attention, where intermediate feature representations from distinct networks are combined before final classification. Such strategies aim to capture complementary information, particularly when differentiating visually similar lesion types.

Data imbalance and inter-class similarity remain persistent challenges in skin cancer classification. To mitigate these issues, researchers have incorporated data augmentation, class weighting, and regularization strategies to stabilize training and improve sensitivity across minority classes. Despite reported improvements in accuracy metrics, many existing studies focus primarily on model performance within controlled experimental settings and provide limited emphasis on deployment or real-time usability.

The present work extends prior research by integrating two architecturally distinct CNN models—Xception and ResNet50—within a unified hybrid framework. Unlike traditional ensemble methods that combine prediction outputs, this approach performs feature-level concatenation to construct a richer representation space. By leveraging depthwise separable convolutions from Xception and residual learning from ResNet50, the proposed model is designed to capture both fine-grained texture information and high-level structural semantics. Additionally, this study emphasizes practical deployment through a web-based interface, addressing the gap between research experimentation and clinical applicability.

3. System Architecture

The developed framework employs a hybrid deep learning architecture to perform automated classification of skin cancer across multiple lesion categories. The architecture integrates structured data preprocessing, parallel feature extraction using two pretrained convolutional neural networks, feature-level fusion, and a web-based deployment layer. The overall workflow consists of five primary stages: data acquisition, preprocessing, dual feature extraction, hybrid classification, and real-time inference.

A. Overall Architecture Overview

The model follows a parallel processing strategy in which two pretrained deep learning backbones—Xception and ResNet50—receive the same preprocessed dermoscopic image as input. Instead of relying on a single network, both architectures independently extract high-dimensional feature representations. These feature vectors are concatenated to form a unified descriptor, which is then passed through fully connected layers for final classification.

This design enables the system to capture complementary information from distinct architectural mechanisms, improving discriminative performance across multiple lesion categories.

B. Data Acquisition and Preprocessing Module

Dermoscopic image samples are obtained from the ISIC repository and subsequently divided into separate training and validation groups for model development and evaluation. To ensure uniformity and model stability, each image undergoes a standardized preprocessing pipeline that includes:

- Validation of image format (JPEG/PNG)
- Conversion to RGB color space
- Noise reduction via Gaussian filtering
- Center cropping to maintain aspect ratio
- Resizing to $128 \times 128 \times 3$ resolution
- Pixel normalization for stable gradient updates.

This preprocessing strategy enhances lesion visibility, reduces illumination inconsistencies, and ensures compatibility with transfer learning models pretrained on ImageNet.

C. Parallel Feature Extraction Layer

The two pretrained convolutional neural Backbone:

Xception Backbone: The Xception architecture employs depthwise separable convolution operations that independently model spatial patterns and inter-channel relationships within feature maps. This mechanism efficiently captures subtle color variations, fine textures, and localized dermoscopic patterns critical for lesion differentiation.

Tuijin Jishu/Journal of Propulsion Technology
ISSN: 1001-4055
Vol. 47 No. 02 (2026)

ResNet50 Backbone: ResNet50 employs residual learning through skip connections, allowing deeper network training without gradient degradation. This structure facilitates hierarchical feature learning, capturing broader lesion structures such as borders, asymmetry, and morphological characteristics.

Both networks are initialized with ImageNet pretrained weights. The final classification layers of each architecture are removed, and Global Average Pooling is applied to convert convolutional outputs into fixed-length feature vectors (2048 dimensions each).

D. Feature Fusion and Classification Module

The extracted 2048-dimensional feature vectors from Xception and ResNet50 are concatenated to form a combined 4096-dimensional representation. This fused feature vector is then processed through a deep classification head consisting of:

- Dense layer (1024 units) with L2 regularization
- Batch Normalization
- ReLU activation
- Dropout for regularization
- Dense layer (512 units)
- Additional Batch Normalization, Dropout
- Final Dense layer with Softmax activation

The Softmax layer outputs probability distributions across four classes:

- Actinic Keratosis (AK)
- Basal Cell Carcinoma (BCC)
- Melanoma (MEL)
- Squamous Cell Carcinoma (SCC)

This hybrid feature-level fusion approach reduces model bias and enhances robustness compared to single-backbone architectures.

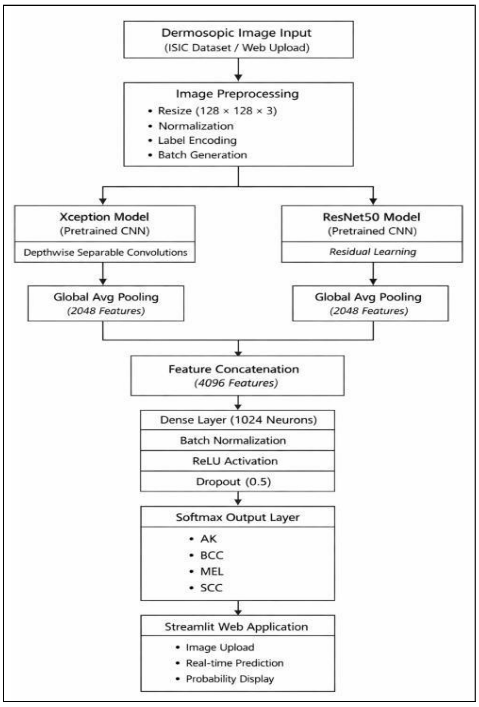

Fig. 1. Overall System Architecture of the Proposed Skin Cancer Classification Framework

4. Proposed Methodology

The developed system employs a hybrid transfer learning approach to perform multi-class classification of skin cancer using dermoscopic image data's. The methodology integrates systematic data preparation, advanced image enhancement, dual-network feature learning, feature-level integration, and optimized multi-class classification. The overall pipeline is designed to enhance lesion representation while maintaining computational efficiency and generalization capability.

A. Dataset Organization and Splitting Strategy

Dermoscopic image samples are obtained from the publicly available ISIC repository and organized into four medically significant categories: Actinic Keratosis (AK), Basal Cell Carcinoma (BCC), Melanoma (MEL), and Squamous Cell Carcinoma (SCC). For dependable performance assessment, the dataset undergoes randomization before being divided into training and validation sets following an 80:20 split configuration. This randomized stratified division preserves class distribution while minimizing sampling bias. The organized directory structure allows efficient batch loading during training.

B. Data Preprocessing

Input preprocessing is performed to enhance the reliability and consistency of the skin cancer classification model. Dermoscopic images are first validated and converted into a standardized RGB format to maintain uniform input structure. Contrast enhancement using CLAHE is applied to improve lesion visibility under varying lighting conditions, followed by Gaussian filtering to reduce noise while preserving important details. Images are center-cropped to maintain aspect ratio and resized to 128 × 128 pixels for computational efficiency. Pixel values are normalized to stabilize gradient updates during training. The processed dataset is randomly shuffled to minimize sampling bias. Finally, the images are divided into training and validation sets using an 80:20 split for balanced performance evaluation.

C. Feature Extraction Using Transfer Learning

Feature extraction in the proposed framework is performed using transfer learning to leverage knowledge from large-scale visual datasets. Two pretrained convolutional neural networks, Xception and ResNet50, are employed as parallel backbone architectures. Both architectures are pretrained on the ImageNet dataset, allowing them to leverage generalized visual representations acquired from large-scale image learning, including edge detection, texture patterns, and structural features. The final classification layers of each network are excluded so that the remaining convolutional layers operate solely as deep feature extractors for dermoscopic image analysis. Global Average Pooling is applied to convert convolutional feature maps into compact fixed- length feature vectors. Selective fine-tuning of higher layers is performed to adapt the pretrained filters to dermoscopic image characteristics while preserving general visual features. This transfer learning strategy reduces training time, improves convergence stability, and enhances the discriminative capability of the model for multi-class skin cancer classification.

D. Hybrid Classification Strategy

To improve multi-class prediction performance, a feature- level hybrid fusion strategy is implemented. The high-dimensional feature vectors obtained from Xception and ResNet50 are concatenated to create a unified representation. This combined feature space captures both fine-grained textural patterns and high-level structural characteristics of skin lesions. The fused vector is processed through fully connected layers incorporating batch normalization, L2 regularization, and dropout to enhance generalization. A Softmax layer produces probability scores for the four target classes: AK, BCC, MEL, and SCC. This hybrid integration approach improves class discrimination and reduces bias compared to single-architecture models.

E. Web-Based Deployment

For real-world implementation, the trained hybrid network is deployed through a web application built on the Streamlit framework. The user interface permits submission of dermoscopic images and provides immediate diagnostic predictions along with associated probability values. Upon image submission, the system automatically performs preprocessing and forwards the input through the trained network for inference. The predicted class label and probability distribution are displayed in real time. This deployment framework transforms the research model into an accessible clinical decision-support tool. The web integration ensures scalability, ease of use, and applicability in dermatological screening environments.

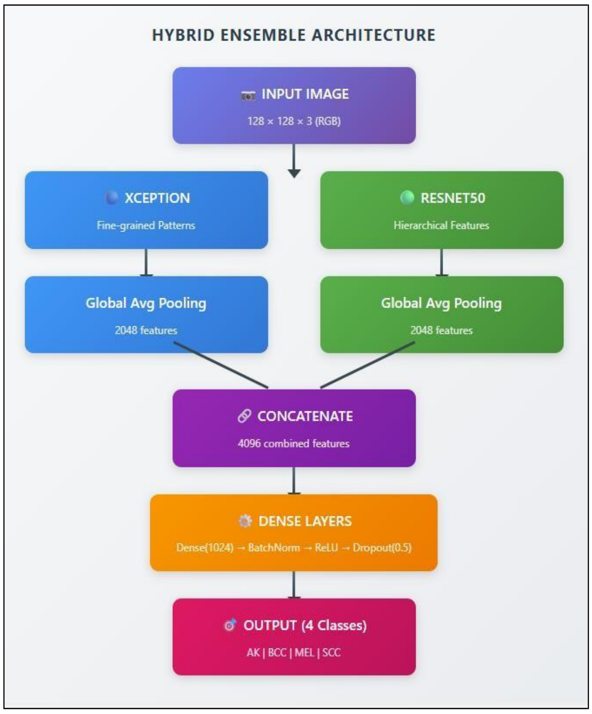

Fig. 2. Proposed Hybrid Ensemble Architecture Integrating Xception and ResNet50 for Skin Cancer Classification.

5. Results And Discussion

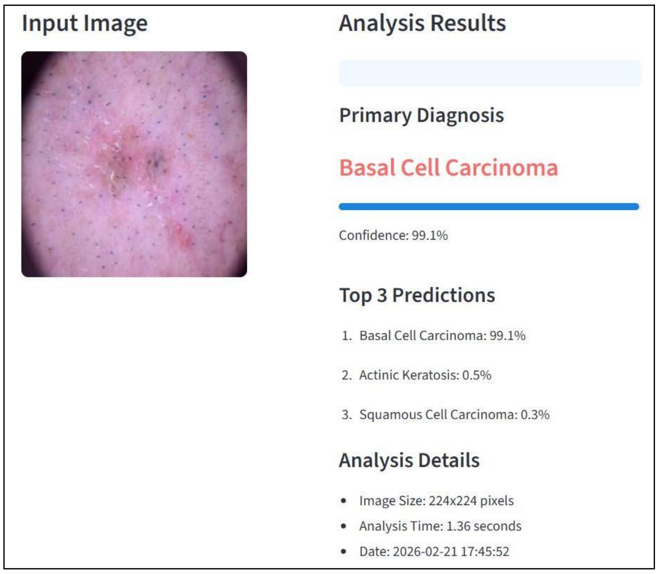

The experimental results demonstrate that the developed hybrid model achieves improved predictive performance on dermoscopic images obtained from the ISIC dataset. In comparison with the standalone pretrained networks, Xception and ResNet50, the combined framework produces higher overall accuracy and maintains more balanced predictions across all lesion classes. The performance improvement is primarily due to the integration of complementary feature representations, which strengthens the model's capacity to differentiate subtle visual differences among various skin cancer categories.

A. Accuracy and Performance Graph:

The comparative accuracy graph illustrates a clear performance advantage of the hybrid configuration over standalone architectures. The visual comparison highlights a consistent improvement in predictive reliability, confirming that feature-level integration strengthens the model's learning capability. The results suggest that combining heterogeneous CNN features enhances overall classification stability.

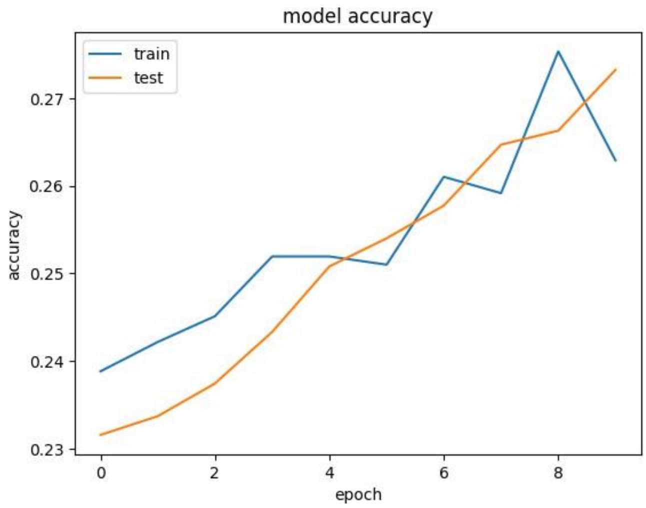

B. Confusion Matrix:

The confusion matrix offers a comprehensive breakdown of class-level prediction results. Larger values along the diagonal represent correctly classified instances for each corresponding category. Lower values in the off-diagonal cells indicate fewer misclassifications across different classes, especially among visually and clinically related lesions such as Basal Cell Carcinoma and Squamous Cell Carcinoma. This distribution demonstrates that the hybrid fusion strategy effectively captures discriminative lesion characteristics and maintains balanced sensitivity across all classes.

C. Precision, Recall, and F1-Score Analysis:

The per-class performance indicators further validate the stability and reliability of the developed framework. The consistently strong precision and recall values across all lesion categories demonstrate the model's ability to effectively minimize both false positive detections and missed classifications. The corresponding F1-scores

reflect a strong balance between sensitivity and specificity, reinforcing the reliability of the hybrid framework in multi-class prediction scenarios.

Overall, the experimental analysis validates that the ensemble-based hybrid architecture enhances discriminative capability, reduces classification bias, and improves generalization performance. These characteristics support its applicability in automated skin cancer screening and clinical decision-support systems.

Class	Name	Total Images	Training	Validation
AK	Actinic Keratosis	867	693	174
SCC	Squamous Cell Carcinoma	628	502	126
BCC	Basal Cell Carcinoma	3,323	2,658	665
MEL	Melanoma	4,522	3,617	905
TOTAL		**9,340**	**7,470**	**1,870**

6. Conclusion

This research presented an integrated hybrid deep learning approach for automatic multi-class classification of skin cancer using dermoscopic images obtained from the ISIC repository. The framework combines two structurally different pretrained convolutional networks, Xception and ResNet50, to exploit diverse feature extraction capabilities. By merging the extracted representations at the feature level, the system captures detailed textural information along with broader structural attributes, resulting in improved discrimination across four major skin lesion categories.

The use of systematic preprocessing, transfer learning adaptation, and regularized dense layers enhances training stability and ensures balanced prediction performance across classes. Comparative analysis confirms that the hybrid configuration achieves superior accuracy and improved class-level consistency relative to the individual backbone models. Confusion matrix evaluation further indicates reduced misclassification among visually similar lesion types and strengthened generalization performance.

Beyond model training and evaluation, the implementation of the network within a web-based platform demonstrates its suitability for practical, real- time dermatological screening applications. Overall, the findings indicate that integrating multiple CNN architectures within a single framework offers a dependable and scalable artificial intelligence solution for automated skin cancer diagnosis.

7. Future Work

Although the proposed hybrid framework demonstrates strong performance in multi-class skin cancer classification, several research directions can further enhance its clinical reliability and scalability. Future investigations may involve extending the dataset to include additional skin lesion categories and incorporating more diverse demographic and imaging conditions to improve model generalization across varied populations. Expanding the dataset size would also support deeper fine-tuning and reduce potential sampling bias.

Advanced architectural enhancements such as attention mechanisms, weighted feature fusion strategies, or transformer-based vision models may be explored to further improve feature representation and class discrimination. Incorporating explainable artificial intelligence techniques, such as saliency mapping or gradient-based visualization methods, could enhance interpretability and assist clinicians in understanding model predictions.

Additionally, optimizing the framework for mobile or cloud-based deployment would improve accessibility in resource-constrained healthcare environments. Future work may also focus on real-time clinical validation studies to evaluate performance in practical diagnostic settings. By integrating these advancements, the system can evolve into a more interpretable, scalable, and clinically adaptable decision-support solution for early skin cancer detection.

Acknowledgment

This research was conducted under the academic guidance of **Dr. Sangeetha K, Professor, Department of Computer Science and Engineering, Panimalar Engineering College, Chennai**. The authors sincerely thank her for her valuable mentorship, constructive feedback, and continuous support throughout the development of the hybrid deep learning framework for skin cancer classification.

The authors also express their gratitude to the faculty members and peers of Panimalar Engineering College for their technical assistance and institutional support. We further acknowledge the International Skin Imaging Collaboration (ISIC) for providing access to dermoscopic image datasets used in this study. Any remaining errors are the sole responsibility of the authors.

References

[1] Y. Wu, H. Zhang, B. Alshaghdali, H. Xiong, Y. Zhao, and L. Wen, "Skin cancer classification with deep learning: A systematic review," *IEEE Access*, vol. 8, pp. 1–20, 2020.

[2] A. Panthakkan, S. S. Raj, and A. P. J. *et al.*, "A novel hybrid approach for accurate skin cancer prediction," *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–12, 2021.

[3] A. S. Al Waisy, A. R. Al Zubaidi, M. A. *et al.*, “A deep learning framework for automated early diagnosis of skin cancer,” *Neural Computing and Applications*, vol. 32, no. 14, pp. 1–15, 2020.
[4] S. N. Hasan, “Accurate deep learning algorithms for skin lesion classification,” *International Journal of Advanced Computer Science and Applications*, vol. 11, no. 6, pp. 1–8, 2020.
[5] K. M. Hosny, M. A. Kassem, and M. M. Foaud, “Classification of skin lesions using transfer learning and deep convolutional neural networks,” *Applied Soft Computing*, vol. 90, pp. 1–12, 2020.
[6] M. J. Hossain, F. Afroz, and M. A. Hossain, “Deep learning approaches for the identification and classification of skin cancer,” *International Journal of Computer Applications*, vol. 177, no. 25, pp. 1–7, 2019.
[7] M. A. Kassem, K. M. Hosny, and M. M. Fouad, “Skin lesions classification into eight classes for ISIC 2019 challenge,” *IEEE Access*, vol. 8, pp. 1–14, 2020.
[8] S. Salekin, M. H. Kabir, and M. Khaliluzzaman *et al.*, “Classification of skin disease with fine-tuned ResNet50 network,” *International Journal of Computer Vision and Signal Processing*, vol. 10, no. 1, pp. 1–6, 2020.
[9] A. Khraisat and A. Yassine, “An integrated deep learning model for skin cancer classification,” *Journal of Biomedical Informatics*, vol. 113, pp. 1–10, 2021.
[10] S. Mahajan, A. Singh, and R. Gupta *et al.*, “A comprehensive analysis of deep learning and transfer learning techniques for skin cancer classification,” *Multimedia Tools and Applications*, vol. 80, no. 15, pp. 1– 25, 2021.