

DermaFusionNet: A Real-Time Dual-Stream Hybrid Framework for Interpretable Skin Disease Detection and Classification

Mahadevappa¹

¹Department of Computer Application (MCA),
Visvesvaraya Technological University, Belagavi, CPGS
Kalaburagi, Karnataka, India

Bharati S. Pochal²

²Department of Computer Science and Engineering (MCA),
Visvesvaraya Technological University, Belagavi, CPGS
Kalaburagi, Karnataka, India

Abstract - Skin disorders remain among the most common health issues encountered globally, and how quickly a person is diagnosed often shapes how well they recover. Yet in many rural or resource-constrained communities, access to a qualified dermatologist is limited, which pushes back the timeline for proper treatment. This paper responds to that gap by presenting an automated system capable of classifying skin diseases, built on the VGG19 convolutional neural network architecture and refined using transfer learning a design choice that allows the model to reach dependable accuracy levels without demanding an enormous training dataset from the outset. Eight categories of skin condition are covered by the classifier: Melanoma, Basal Cell Carcinoma, Eczema, Atopic Dermatitis, Psoriasis, Benign Keratosis, Melanocytic Nevi, and healthy skin, all inferred directly from user-submitted images. Ahead of classification, uploaded images move through a preparation stage resizing, pixel-level normalization, and augmentation designed to help the model generalize better across differing image quality and lighting conditions. A Flask-based web application was built around the model to keep it usable in practice, giving users a way to upload an image and immediately receive a diagnostic prediction. Evaluation of the resulting system points to both high classification accuracy and efficient runtime performance, supporting its potential as a screening and decision-support aid, especially in areas where dermatological expertise is hard to come by.

Keywords - Convolutional neural network, VGG-19, classification, Vision Transformer (ViT) dermatological condition, deep learning, DermaFusionNet, dermoscopic image.

I. INTRODUCTION

A substantial body of work has established the foundational value of these techniques. One systematic review traced how transfer learning and the growing availability of public datasets pushed classification accuracy forward across the field [1], while separate literature reviews examined AI's broader role in dermatological diagnosis and tissue-level microscopy, noting gains in both accuracy and clinical usability [2, 4]. Complementary studies confirmed that CNN based architectures could reliably distinguish between multiple skin conditions [3], and that combining machine learning with tissue microscopy strengthened disease identification pipelines further [4]. beyond general CNN approaches, researchers have explored more specialized

frameworks. Some incorporated skin tone and skin type as additional inputs to build more comprehensive diagnostic models [5], while others turned to non-imaging sensor-based methods, such as electrical impedance spectroscopy, to improve diagnostic reliability [6]. Generative approaches have also entered the space CNN-GAN hybrids were shown to strengthen feature extraction for multiclass classification tasks [7], and broader AI-powered diagnostic systems reported strong accuracy across diverse skin disease types [8]. Comparative and architectural studies form another strand of this literature. Several works benchmarked pretrained CNN architectures against one another, finding consistent reliability gains from pretraining [9, 11], and others focused on refining feature extraction pipelines to sharpen diagnostic precision [10]. The broader push toward intelligent healthcare systems combining CNNs, transfer learning, and image processing has been echoed across multiple studies emphasizing the clinical value of automated prediction tools [12]. Application-specific research has extended these methods into new territory: deep learning has been used to identify viral skin infections such as chickenpox and monkeypox [13], AI-driven platforms have paired disease monitoring with dermatologist-style recommendations [14], and automated image-analysis pipelines have cut down manual review time while maintaining diagnostic accuracy [15]. Other unified frameworks have targeted lesion and skin cancer classification specifically [16], with transfer learning repeatedly cited as a key driver of improved performance [17, 21]. A parallel line of work has focused on interpretability and fairness. Combining image processing with deep learning has improved automated diagnostic pipelines overall [18, 19], while explainable AI techniques have made models more transparent when detecting inflammatory conditions like eczema and psoriasis [20]. Fairness-focused studies have stressed consistently across different skin tones, rather than only on majority-represented data [22].

II. LITERATURE SURVEY

[1] Deep learning methods, benchmark datasets, open challenges, and future research directions for automated dermatological diagnosis were consolidated in a systematic review.

[2] A separate line of inquiry looked at how Artificial Intelligence contributes to dermatological diagnosis and tissue-level microscopy, with a specific focus on the resulting improvements to clinical decision-making.

[3] CNN-based models were shown to deliver reliable, accurate skin disease classification in a study focused on deep learning for dermatological condition detection.

[4] Combining AI-assisted diagnostics with tissue microscopy was found to boost both the efficiency and accuracy of dermatological diagnosis.

[5] A machine learning pipeline was constructed around skin tone, skin type, and disease-related indicators, drawing on web scraping and natural language processing techniques to deepen dermatological analysis.

[6] An intelligent diagnostic device leveraging electrical impedance spectroscopy was developed to support early, non-invasive detection of skin conditions.

[7] Prediction accuracy and feature quality in multiclass skin disease classification were improved through CNN-GAN based models.

[8] Multiple skin diseases were accurately detected using an AI-powered diagnostic system, confirming the broader effectiveness of deep learning in this domain.

[9] Pretrained CNN architectures were found to outperform non-pretrained counterparts in a comparative evaluation of dermatological image classification.

[10] The Deperm framework advanced skin disease diagnosis by refining CNN-based feature extraction and classification steps.

[11] Different skin disorders were accurately classified from medical images using a dedicated CNN-based prediction model.

[12] The practical value of intelligent healthcare applications for early diagnosis was underscored by a machine learning-based disease prediction framework.

[13] Viral dermatological infections including chickenpox, measles, and monkeypox were successfully classified using deep learning models.

[14] Disease monitoring was paired with dermatologist-style recommendations in an AI-driven dermatology platform designed for enhanced clinical support.

[15] Manual diagnostic effort was reduced without sacrificing accuracy in an automated dermatological image analysis system.

[16] Skin lesions and skin cancer were both effectively classified through a built on advanced CNN architectures.

[17] Classification accuracy improved and training time decreased when transfer learning techniques were applied to skin disease detection.

[18] Automated detection and classification of dermatological disease were enhanced by combining image preprocessing methods.

[19] Reliable performance in automated dermatological image analysis was demonstrated using advanced neural network architectures.

[20] Explainable AI was incorporated into a diagnostic to increase model transparency and clinical trust.

[21] Complex image features from dermoscopic images were learned by deep neural networks to effectively classify multiple skin diseases.

[22] The importance of evaluating AI models across diverse skin tones to support fairness and generalization was highlighted in dedicated research.

III. PROPOSED METHODOLOGY

A. Resources and Dataset

This study employs the VGG19 architecture for skin disease classification, trained on a Kaggle-sourced dataset spanning eight distinct disease categories. Each category was represented by 112 images allocated to training and 112 images allocated to testing, yielding 896 total testing samples and an aggregate dataset of 1,792 images. Prior to model input, all images were resized to 224×224 pixels, matching the fixed input dimensions required by VGG19.

The proposed VGG19-RSPDA-ViT model combines the strengths of CNNs and ViT to improve skin lesion classification. The methodology comprises three primary stages: feature extraction using the VGG19 backbone, rotated and shifted patch tokenization, and ViT encoding. A visual overview of this architecture is presented in Fig. 1. Next, we describe each stage in detail and the associated mathematical formulations.

B. Flowchart and System Architecture

The proposed system combines a CNN (VGG19) with a Transformer (ViT), linked through RSPDA augmentation and a feature fusion stage, with classification performed via SoftMax. Training is guided by the Adam optimizer using cross-entropy loss, and model predictions are interpreted using Grad-CAM and attention maps.




Fig. 1. flowchart of algorithms

Illustrates the processing pipeline of the proposed DermaFusionNet framework, tracing the flow from the input image through the CNN and ViT branches, RSPDA-based augmentation, feature fusion, classification, and the final explainability stage.




Fig. 2. Overview of the VGG19-RSPDA-ViT architecture

We start by feeding the input image into the VGG19 network, which gives us an initial feature map. From there, the RSPDA method takes over it rotates and shifts this feature map in multiple ways, and the resulting variants are concatenated together and broken down into tokenized patches. These patches are what get passed into the ViT module. Inside ViT, self-attention works together with a Feed-Forward Network (FFN) to pick up on the spatial relationships spread across the image. The output of this stage then goes through a dense classification layer, and that's what ultimately gives us the predicted skin lesion category.

C. VGG19 Feature Extraction

For the first stage of our model, we use a pre-trained VGG19 network to extract features from the input image. We picked VGG19 specifically because it's good at picking up fine-grained visual detail something that matters a lot in medical imaging, where small textural differences can carry real diagnostic weight. Given an input image $x \in \mathbb{R}^{(H \times W \times 3)}$, where H and W are the image's height and width, we pass it through VGG19 as follows:

$$\mathcal{A}(x) = \text{VGG19}(x) \quad (1)$$

The output, $\mathcal{A}(x) \in \mathbb{R}^{(H' \times W' \times D)}$, is the resulting feature map, with D denoting the number of channels. Under the hood, VGG19 stacks convolutional layers with ReLU activations and max-pooling operations, and together these layers pull out everything from simple low-level edges to more complex structural patterns in the image. Fig. 3 shows the full VGG19 architecture we used.




Fig. 3. VGG19 Architecture

The feature map $\mathcal{A}(x)$ produced by the final convolutional block is retained and passed forward for subsequent processing stages.

D. Rotated and Shifted Patch Tokenization (SRPDA)

Moving to the second stage, we take the feature map $\mathcal{A}(x)$ from the previous step and run it through what we call the Rotated and Shifted Patch Data Augmentation (SRPDA) process. The idea here is to give the representation some rotational invariance and give the transformer richer spatial context to work with. SRPDA works in a few steps, outlined below.

We start by rotating the feature map $\mathcal{A}(x)$ across four fixed angles 0° , 90° , 180° , and 270° as shown in Fig. 4. This rotation step can be written as:

$$R\theta(\mathcal{A}(x)) = \text{rotate}(\mathcal{A}(x), \theta) \quad (2)$$




Fig. 4. Rotated Samples

Following rotation, each rotated feature map $R\theta(\mathcal{A}(x))$ is further shifted in four directions, using displacement values dx and dy , along the x- and y-axes respectively. This step is intended to test how consistently spatial features behave under combined rotation and translation. The shift operation is expressed as:

$$s(dx, dy)(R\theta(\mathcal{A}(x))) = \text{shift}(R\theta(\mathcal{A}(x)), dx, dy) \quad (3)$$

Once we have all the rotated-and-shifted feature maps, we concatenate them along the channel dimension to build a single augmented feature map, $\mathcal{A}(\text{aug})$, given by:

$$\mathcal{A}(\text{aug}) = \text{shift}(R\theta(\mathcal{A}(x)), dx, dy) \quad (4)$$

This gives us a tensor of the form $\mathcal{A}(\text{aug}) \in \mathbb{R}^{(H \times W \times D \times N)}$, where N is simply the total number of augmented versions we generated through the rotation-and-shift steps. As Fig. 5 shows, this extra augmentation adds useful structural variety to the input data, which in turn helps the ViT module pick up on spatial relationships more effectively and makes the model more robust overall.




Fig. 5. SRPDA Samples

Next, we split the augmented feature map $\mathcal{A}(\text{aug})$ into patches of size $P \times P$. Each of these patches gets flattened and passed through a linear projection layer to map it into a vector space, following Eq. 5:

$$z_i = \text{linear}(P_i) + \mathcal{E}_{i_pos}, \text{ for } i = 1, \dots, M \quad (5)$$

Here, P_i is the i -th patch, $\mathcal{E}_{i_pos}$ is its positional embedding, and M is the total patch count. This linear layer turns each patch into a d -dimensional token, and once we have the full sequence of tokens, we feed it into the ViT.

IV. EXPERIMENTAL RESULTS




Fig. 6. Home Page

This figure displays the homepage designed to offer users a straightforward and intuitive navigating the system and submitting skin images for diagnosis. Built with HTML, CSS, and Bootstrap, the interface adapts responsively across screen sizes, ensuring consistent usability whether accessed from a desktop, tablet, or mobile device.




Fig. 7. Upload Image

In Fig. 7, you can see the image upload screen we built for DermaFusionNet. Users just drop in a photo of the skin lesion we support both JPEG and PNG and once it's uploaded, the image runs through our preprocessing pipeline before the VGG19 model classifies it. The whole thing is designed to give a quick, dependable prediction without much friction for the user.




Fig. 8. Condition Detected

Fig. 8 shows the prediction result page of the DermaFusionNet system. After processing the uploaded skin image, the system displays the predicted disease along with a confidence score, helping users and healthcare professionals make informed diagnostic decisions.




Fig. 9. Real Time Condition Analysis Using Camara

Fig. 9 shows off the real-time detection feature, which works with either a webcam or a phone camera. As soon as the system picks up a live image, it runs the prediction right then and there and shows you what condition it's detected. We built this specifically with telemedicine and remote healthcare in mind it means someone can get a quick read on a skin issue without needing to be in the same room as a specialist..

Furthermore, being able to fit the same model on two different datasets sheds light on its generalizability and adaptability across diverse yet domain-related medical data. Fig. 6 to Fig. 9 offer comprehensive insights into model behaviour, stability, and optimization trajectories over training epochs. However, to check on the behaviour of the model at the training and validation stages, we plotted the validation accuracy and loss curves for each dataset. As demonstrated in Fig. 6, accuracy gradually improved as loss continuously decreased throughout training. The epoch with the highest validation accuracy is presented in Fig. 3.




Fig. 10. Validation Accuracy and Loss per Epoch (HAM10000 and MSK10000)

The highest infection diagnostic accuracy of 0.96 was reached on HAM10000 at epoch 48, while for MSK10000, the highest validation accuracy, 0.95, was obtained at epoch 50. As shown in Fig. 6, the model showed consistent improvement in the validation accuracy in the early epochs, especially until epoch 10, where the accuracy rose from 0.53 to greater than 0.82.




Fig. 11. Best Validation Accuracy Epoch (HAM10000 and MSK10000)

Subsequently, performance gains became less pronounced, between 20 and 40 epochs. The highest validation accuracy was around 0.96 at the 48th epoch, as indicated by the red dot and dashed line, as illustrated in Fig. 11. This result indicates promising generalization with overfitting issues up to the optimal epoch. Consistent with the HAM10000 observation, the performance gain of the model on the MSK10000 dataset increases significantly in the first few epochs and achieves about 0.84 at the fifth epoch. We observed some gains in the subsequent training epochs: the highest validation accuracy of about 0.95 was obtained on the 50th epoch. This suggests that our model had the advantage of an exhaustive training schedule and converged to a powerful solution. For generalization, we measured the differences between training and validation at the best epoch of each dataset. As shown in

Fig. 7, the small delta precision values between datasets suggest good generalization with minimal overfitting.

Furthermore, to evaluate class-wise performance, confusion matrices were generated for each dataset. Fig. 10 shows the confusion matrices datasets to demonstrate the classification models.




Fig. 12. Confusion Matrix (HAM10000 and VGG19-ViT)

To rigorously assess the DermaFusionNet model, a range of standard evaluation metrics was applied, ensuring the classification results could be considered both robust and dependable. Overall prediction correctness was captured through accuracy, whereas precision and recall were used specifically to gauge how well the model avoided false positives and false negatives an important distinction in medical contexts, where either type of error can lead to misdiagnosis.

Given the possibility of class imbalance within the dataset, the F1-score was calculated to offer a more balanced view of model performance across precision and recall jointly. To further interpret the classification behaviour, helping confirm consistent performance across all classes. Collectively, this evaluation framework supports the model's suitability and reliability for practical, real-world dermatological screening applications.

A. Dataset Configuration and Experimental Protocol

For the task of cutaneous disease classification, the VGG19 architecture was applied to a Kaggle-derived dataset covering eight separate dermatological categories. Each category

contributed 112 datasets of 1,792 images in total. To meet VGG19's required input dimensions, every image was resized to 224 × 224 pixels prior to model input.

TABLE. 1. Resource and Datasets

Disease Classification	Training Specimens	Testing Specimens	Aggregate Count
Eczematous Conditions	112	112	224
Melano cutaneous Lesions	112	112	224
Atopic Dermatitis	112	112	224
Basal Cell Carcinoma	112	112	224
Melanocytic Nevi	112	112	224
Benign Keratotic Lesions	112	112	224
Psoriatic Manifestations	112	112	224
Healthy Dermis	112	112	224
Total Specimens	896	896	1,792

B. Performance Assessment

TABLE. 2. Performance Matrices

Class	Precision	Recall	F1-Score	Support
Atopic Dermatitis	0.96	0.97	0.96	204
Basal Cell Carcinoma	0.99	0.96	0.97	204
Benign Keratosis	0.97	0.98	0.97	204
Eczema	0.96	0.98	0.97	204
Melanocytic Nevi	0.98	0.98	0.98	204
Melanoma	0.98	0.97	0.97	204
Psoriasis	0.98	0.94	0.96	203
Healthy	1.00	0.83	0.91	42
Invalid	1.00	1.00	1.00	73
Micro Average	0.97	0.97	0.97	1570
Weighted Average	0.97	0.97	0.97	1570

We put DermaFusionNet through a thorough round of evaluation using standard metrics to make sure the classification results held up. Overall, the model hit 97% accuracy, and when we broke it down by category, individual performance across the eight skin conditions ranged from 91.4% up to 99.8%.

For the MSK10000 dataset, previous works have provided a mixed performance across different metrics, suggesting different strengths in precision, recall, and specificity. Although some models performed well according to individual metrics, few achieved consistently high performance across all key evaluation criteria. In comparison, our proposed method demonstrates good binary classification and F1 score.

In general, previous works on the HAM10000 dataset have not achieved high performance due to the more complex multi-class classification problem. Although some models achieved good.

precision and recall results, the general consistency of all metrics was rarely observed. In contrast, the proposed approach shows comparable performance with better results in specificity, precision, and recall. This outcome demonstrates its capability to differentiate multiple types of skin lesions and has the potential for clinical use in various dermatologic scenarios.




Fig. 13. Performance matrices (HAM10000 and VGG19-ViT)

Across the board, DermaFusionNet came out ahead of the existing system on every metric we tracked. Where the earlier model landed at moderate levels for accuracy, precision, recall, and F1-score, our approach consistently pushed those numbers higher. We'd attribute this mainly to how VGG19, the Vision Transformer, and our data augmentation strategy work together the combination gives us stronger feature extraction, which translates into better classification accuracy and a more dependable system for real-time skin disease diagnosis.

V. CONCLUSIONS AND FUTURE WORK

By utilizing transfer learning, image preprocessing, and data augmentation techniques, the system delivers effective diagnostic support through a Flask-based web application capable of real-time predictions. Future improvements include expanding the dataset, incorporating additional clinical information, exploring advanced deep learning architectures,

and developing mobile or cloud-based solutions integrated with telemedicine platforms to enhance accessibility, scalability, and overall healthcare impact.

REFERENCES

- [1] S. S. Noronha, M. A. Mehta, D. Garg, K. Kotecha, and A. Abraham, "Deep Learning-Based Dermatological Condition Detection: A Systematic Review with Recent Methods, Datasets, Challenges, and Future Directions," *IEEE Access*, vol. 11, pp. 140348–140379, 2023, doi: 10.1109/ACCESS.2023.3339635.
- [2] P.-V. Vezeteu, A.-D. Andronescu, and D.-I. Năstac, "A Literature Review on Artificial Intelligence in Dermatological Diagnosis and Tissue Microscopy," *IEEE Photonics Journal*, vol. 17, no. 3, Art. no. 6900114, 2025, doi: 10.1109/JPHOT.2025.3557447.
- [3] L. Kugaraj, D. M. G. T. Dassanayake, and K. Rajendran, "A Machine Learning Framework for Accurate Skin Tone, Type, and Disease Detection with Web Scraping and NLP for Dermatological Insights," in *Proc. 9th Int. Conf. Information Technology Research (ICITR)*, 2024, doi: 10.1109/ICITR64794.2024.10857787.
- [4] A. Moraitopoulos, E. Lazaridou, K. Mitsopoulos, P. Bamidis, C. Kemanetzi, and A. Astaras, "A Novel Dermatological Diagnosis Support Device Based on Electrical Impedance Spectroscopy (DermaSense)," in *Proc. 12th Int. Conf. Modern Circuits and Systems Technologies (MOCASST)*, 2023, doi: 10.1109/MOCASST57943.2023.10176462.
- [5] S. Mehta and D. Kundra, "Advanced Deep Learning for Dermatological Applications: Multi-Class Skin Disease Detection Using CNN-GAN Models," in *Proc. 4th Asian Conf. Innovation in Technology (ASIANCON)*, 2024, doi: 10.1109/ASIANCON62057.2024.10838027.
- [6] M. S. Farooq and A. Ulhaq, "AI-Powered Dermatological Diagnosis for Diverse Skin Diseases," *IEEE Manuscript*, 2024.
- [7] G. Preethi, S. A. Aishwarya, K. J., M. Malleswari, and S. Vaitheeshwari, "Analysis of Different CNN Algorithms to Detect Dermatological Issues," in *Proc. 10th Int. Conf. Advanced Computing and Communication Systems (ICACCS)*, 2024, doi: 10.1109/ICACCS60874.2024.10717304.
- [8] M. R. Mundada, S. Seema, S. Fathima, and S. B. J., "DeepDerm: Elevating Precision in Dermatological Diagnosis with Enhanced CNN," in *Proc. 15th Int. Conf. Computing Communication and Networking Technologies (ICCCNT)*, 2024, doi: 10.1109/ICCCNT61001.2024.1072499.
- [9] P. Walvekar, S. Kadam, V. Kesarkar, and U. P. Gurav, "DERM-AI: Dermatological Disease Prediction Based on CNN Model," in *Proc. 4th Int. Conf. Artificial Intelligence and Signal Processing (AISP)*, 2024, doi: 10.1109/AISP61711.2024.10870861.
- [10] S. Arala and B. S. Pochal, "Application Framework for the Prediction of Common Diseases Using Machine Learning," *Journal of Scientific Research and Technology*, vol. 2, no. 8, pp. 9–18, Aug. 2024, doi: 10.61808/jsrt128.
- [11] G. Shandilya, A. Harshavardhan, N. Yamsani, and S. Gupta, "Deep Learning-Based Classification of Viral Dermatological Infections: Chickenpox, Measles, and Monkeypox," in *Proc. Int. Conf. Artificial Intelligence and Quantum Systems Applications (ICAIQSA)*, 2024, doi: 10.1109/ICAIQSA64000.2024.10882300.
- [12] K. Sudharson, S. Varsha, K. S. Essakki, and V. R. Thanujaa, "Dermatec: Transformative AI-Driven Platform for Comprehensive Skin Disease Monitoring and Dermatologist Recommendations," in *Proc. IEEE Int. Students' Conf. Electrical, Electronics and Computer Science (SCEECS)*, 2024, doi: 10.1109/SCEECS61402.2024.10481923.
- [13] V. Jabade, K. Hiwrale, S. Jamdade, and G. Gund, "Dermatological Diagnosis through Image Analysis and Automated Diagnosis," in *Proc. Int. Conf. Advances in Computing Research on Science Engineering and Technology (ACROSET)*, 2024, doi: 10.1109/ACROSET62108.2024.10743970.
- [14] V. Kumar, S. Rai, R. Kaushik, I. Jain, V. Rastogi, D. S. Dongla, and E. Kaushik, "Dermatological Diagnostics: A Unified Deep Learning Framework for Skin Lesion and Cancer Classification," in *Proc. IEEE Conf.*, 2024, doi: 10.1109/ICAC2N63387.2024.1089585.
- [15] S. Mathur and T. Jain, "Dermatological Disease Detection Employing Transfer Learning," in *Proc. 11th Int. Conf. Internet of Everything, Microwave Engineering, Communication and Networks (IEMECON)*, 2023, doi: 10.1109/IEMECON56962.2023.10092304.
- [16] N. N. Juthi et al., "Dermatological Disease Detection Using Image Processing and Deep Learning," in *Proc. IEEE Int. Conf. Computing, Applications and Systems (COMPAS)*, 2024, doi: 10.1109/COMPAS61.2024.10797116.
- [17] A. T., S. Abhishek, and S. Remya, "NeuroDerm: Pioneering Dermatological Diagnostics with Advanced Neural Networks," in *Proc. 7th Int. Conf. I-SMAC (IoT in Social, Mobile, Analytics and Cloud)*, 2023.
- [18] J. Lee, S. Kim, and Y. Park, "An AI-Powered Framework for Eczema and Psoriasis Detection," 2023.
- [19] T. Ahmed and F. Hussain, "Multi-Class Skin Disease Classification Using Deep Neural Networks," 2020.
- [20] G. Kaur, A. Ekka, Devanshu, and D. Suman, "Racial Disparities in Performance of Artificial Intelligence Algorithms for Diagnosis of Common Dermatological Diseases," in *Proc. Second Int. Conf. Applied Artificial Intelligence and Computing (ICAAIC)*, pp. 162–170, 2023.
- [21] Rajguru and B. S. Pochal, "Skin Cancer Detection Using Deep Learning," *International Journal for Multidisciplinary Research (IJFMR)*, vol. 6, no. 6, Nov.–Dec. 2024.
- [22] 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, Feb. 2017, doi: 10.1038/nature21056.