International Journal of Advanced Trends in Multidisciplinary Research
ISSN: 3139-0986, Volume 2 Issue 7
July 2026, PP 1-7

A Deep Learning Framework for Clinical Facial Skin Disease Identification Using CNN Models

M.Swathi[1], B. Ranjana[2]
[1]Assistant Professor, Department of MCA, Megha Institute of Engineering and Technology for Women, Edulabad (Village), Ghatkesar (Mandal), Medchal District, Telangana – 501301, India.
[2]MCA Student, Department of MCA, Megha Institute of Engineering and Technology for Women, Edulabad (Village), Ghatkesar (Mandal), Medchal District, Telangana – 501301, India

***Abstract* – The impact of facial skin problems on one's look and general health makes them a significant dermatological issue. Timely treatment relies on early detection of many illnesses, but manual diagnosis may be time-consuming and frequently requires clinical skill. Convolutional Neural Networks (CNNs) and other deep learning advancements have shown exceptional promise in reliably analysing medical pictures. This study examines the efficacy of convolutional neural network (CNN) models for automated facial skin disease categorisation using clinical face photos. The illnesses under investigation include Rosacea, Actinic Keratosis, Seborrhoeic Keratosis, Lupus Erythematosus, Basal Cell Carcinoma, and Squamous Cell Carcinoma. One of the biggest publicly disclosed collections of face dermatological photos in China, the Xiangya-Derm dataset is used in the experiments. A subset of 2,656 facial photographs was created from this dataset. We took a look at the classification performance of five popular CNN architectures: ResNet-50, Inception-v3, DenseNet121, Xception, and Inception-ResNet-v2. In addition, pretrained models were able to enhance their feature extraction and classification capabilities via the use of transfer learning, which included drawing on an independent dataset that had the same illness categories gathered from other body areas. Experiments show that transfer learning always improves model performance, with Inception-ResNet-v2 doing the best overall. The model achieved recall values of 92.9% for Lupus Erythematosus, 89.2% for Basal Cell Carcinoma, and 84.3% for Seborrhoeic Keratosis on the facial-image test set of 388 clinical images. On average, the model achieved 77.0% recall and 70.8% precision across all six disease categories. The results show that transfer learning has the ability to improve automated skin disease detection in clinical practice and validate that CNN-based facial skin analysis may provide trustworthy assistance for computer-aided dermatological diagnosis.**

***Keywords* – Facial Skin Disease Classification, Convolutional Neural Networks (CNN), Deep Learning, Transfer Learning, Medical Image Analysis, Dermatological Diagnosis, Xiangya-Derm Dataset, Clinical Facial Images, Skin Lesion Detection, Computer-Aided Diagnosis (CAD).**

I. INTRODUCTION

Every year, millions of people throughout the globe are impacted by skin illnesses, which place heavy physical, psychological, and social demands on society as a whole. Delays in diagnosis or insufficient treatment may cause cosmetic deformity, persistent pain, diminished quality of life, and disease progression, even though many dermatological conditions are not immediately fatal. Diseases of the face are especially noteworthy because they manifest on the most noticeable area of the body and may have a profound impact on one's sense of self-worth, relationships with others, and mental health. Since then, one of the main goals of contemporary dermatological research has been the creation of accurate and effective methods for diagnosing skin problems affecting the face.

In the past, dermatologists' clinical examinations were the gold standard for diagnosing skin problems. Observing the colour, texture, form, and distribution of the lesion, as well as comparing it to past clinical instances, is an important part of the diagnostic procedure. Although this method has been successful in the past, it is quite susceptible to inter-observer variability and relies heavily on specialised knowledge. Accurate and rapid diagnosis is still a major obstacle in many areas, especially those with a scarcity of dermatological experts. Opportunities for the development of automated computer-aided diagnostic systems that may help healthcare professionals and improve access to early screening have arisen due to the increased availability of high-resolution digital imaging equipment and smartphones.

The field of medical image analysis has been greatly impacted by the recent developments in artificial intelligence. In the realm of deep learning, Convolutional Neural Networks (CNNs) stand out for their remarkable capacity to extract hierarchical visual features from pictures directly, without the need for manual feature engineering. Countless healthcare applications have found great success using CNN-based models. These include dermatological lesion categorisation, pathological picture interpretation, and illness diagnosis from radiological scans. Because of their remarkable capacity to acquire intricate spatial representations, they are well suited to the task of enhancing diagnostic consistency and decreasing human error by detecting the small visual distinctions among different skin conditions.

In spite of all these improvements, there is still a dearth of studies that focus on face skin disease categorisation in particular. It is challenging to create models tailored for face diagnosis since the majority of publicly accessible dermatological image archives include pictures acquired

from other parts of the body without highlighting facial abnormalities. Automated face categorisation is already challenging enough without adding in the unique anatomical and visual features of facial skin, such as lighting, contours, pigmentation, and expression. Therefore, developing very accurate diagnosis algorithms has been severely hindered by the lack of well annotated datasets in the field of face dermatology.

In light of this shortcoming, the current investigation employs the Xiangya-Derm clinical image dataset to examine six prevalent skin diseases of the face: Rosacea (ROS), Actinic Keratosis (AK), Seborrhoeic Keratosis (SK), Lupus Erythematosus (LE), Basal Cell Carcinoma (BCC), and Squamous Cell Carcinoma (SCC). In order to test the viability of convolutional neural network (CNN) facial skin disease detection, a subset of the whole dataset was created, including 2,656 clinical photographs of the face. Under controlled conditions, the experimental framework compares the efficacy of five popular deep learning architectures: ResNet-50, Inception-v3, DenseNet121, Xception, and Inception-ResNet-v2. In addition, while keeping the original experimental technique, transfer learning is used to improve feature learning and classification performance utilising photos acquired from various parts of the body.

Evaluating the efficacy of well-established CNN architectures for automated face skin disease categorisation using clinical photos while preserving high diagnostic reliability is the major goal of this study. The work sheds light on the topic of choosing suitable architectures for computer-aided dermatological diagnostics by comparing several deep learning models and examining the effects of transfer learning. The results of the experiment show that with the help of deep learning algorithms, dermatologists can identify skin illnesses on the face more accurately. They also show how intelligent diagnostic systems might help with early diagnosis, clinical decision-making, and other future uses in dermatology.

II. LITERATURE SURVEY

The use of AI in dermatology has greatly improved the ability to automatically identify and categorise skin conditions using clinical and dermoscopy pictures. Convolutional Neural Networks (CNNs) are one kind of artificial intelligence that has surprised many with their ability to learn discriminative visual features from medical pictures automatically, without the need for human feature engineers. Researchers have been motivated to create computer-aided diagnostic systems that may assist dermatologists in recognising skin problems more accurately and consistently due to the fast advancements in deep learning.

Hay et al., who highlighted that skin illnesses are one of the most common non-fatal health issues globally, performed one of the first thorough examinations of the global effect of dermatological disorders. The study's findings underscore the significant financial, emotional, and physical costs linked to skin conditions, highlighting the need for effective diagnostic methods that may aid in the early detection of diseases. The importance of creating automated diagnostic tools for the treatment of skin diseases was shown in this study.

The causes of many skin diseases have been the subject of several investigations. According to research by Huang et al., who looked at the correlation between acne severity and food habits in Chinese teenagers, drinking too many soft drinks can lead to the development of moderate to severe acne. In a similar vein, Deng et al. created and validated a rosacea-specific quality-of-life evaluation tool, giving doctors a better way to gauge the emotional toll of dealing with persistent skin conditions on the face. A person's mental, emotional, and social health are all impacted by dermatological illnesses, as these studies show.

Important discoveries in dermatological diagnosis have also come from studies of inflammatory and autoimmune skin disorders. A better knowledge of the pathogenic underpinnings of systemic lupus erythematosus was achieved when Chen et al. developed an experimental lupus model using leptin-transgenic pigs. An further study by Xie et al. examined the role of intracellular ATP and ATP synthase in keratinocyte differentiation, which shed light on the molecular processes that govern skin growth and disease progression. In addition, Bewley stressed the underappreciated psychosocial aspects of chronic skin illnesses and advocated for more comprehensive treatment plans that include mental health services alongside dermatological ones.

With the advent of deep learning, skin disease detection by computer has been greatly enhanced. The automated extraction of hierarchical visual representations using CNN architectures has led to exceptional results in the identification of different skin lesions. The promise of CNN-based models for clinical decision support was highlighted by Haenssle et al., who showed that these models could categorise skin lesions as well as professional dermatologists. In a similar vein, Avdeenko et al. looked at deep learning methods for skin cancer detection and found encouraging results in terms of classification accuracy, suggesting that AI may help doctors spot cancerous skin lesions in medical photos. Also, a meta-analysis and thorough evaluation of deep learning techniques for skin cancer diagnosis was carried out by Zhang et al., who found that convolutional neural network (CNN) based systems routinely beat traditional image processing techniques for dermatological picture categorisation.

While convolutional neural networks (CNNs) have shown efficient for general dermatological image analysis in earlier research, face skin disease categorisation using specialised facial clinical image datasets has received very less attention. Developing models specialised for face lesion detection is made more challenging by the fact that most publicly accessible skin disease datasets incorporate photos gathered from numerous body areas. Automatic

diagnosis becomes more complicated when dealing with face skin because of its distinct visual properties, such as differences in lighting, facial features, skin texture, and anatomical structures. Therefore, when it comes to face dermatology, models that were trained on datasets of generic skin images could not work as well.

To overcome this shortcoming, the current research makes use of the Xiangya-Derm dataset, which contains annotated clinical pictures of the face that depict six prevalent skin disorders of the face. While most prior studies have concentrated on skin lesion classification in general, this study uses transfer learning to assess how well ResNet-50, Inception-v3, DenseNet121, Xception, and Inception-ResNet-v2 perform in recognising facial skin diseases. This work adds to the development of trustworthy computer-aided clinical decision-support systems by evaluating deep learning approaches for face dermatological diagnosis and by comparing several CNN architectures under similar experimental settings.

III. RESEARCH BACKGROUND

Because of deep learning's capacity to learn intricate visual patterns straight from clinical pictures, its application to dermatological image analysis has garnered a lot of interest. In the past, most CAD systems used manually-created picture features including texture descriptors, colour histograms, and shape-based characteristics. These features generally needed a lot of domain knowledge and didn't work consistently for diverse skin disorders. Automatic hierarchical feature extraction, made possible by Convolutional Neural Networks (CNNs), greatly improved classification accuracy and generalisability, and thereby revolutionised medical image analysis.

In the field of dermatology, a number of CNN architectures have shown encouraging outcomes. ResNet-50 improved feature propagation during training by including residual learning to circumvent the degradation issue that deeper neural networks face. Utilising multi-scale convolutional processes inside a single module, Inception-v3 enhanced computational efficiency. This enabled the network to collect both local and global picture features. DenseNet121 improved feature reuse even further by decreasing duplicate learning and boosting information flow throughout the network via dense connections established across subsequent levels. Xception improves computing efficiency and feature representation by expanding the Inception principle and using depthwise separable convolutions. In addition, Inception-ResNet-v2 integrated residual learning into the Inception framework, enabling deeper structures to get better convergence without sacrificing recognition performance. Due to their proficiency in detecting minute visual variations among clinically comparable disorders, these structures have seen extensive usage in medical picture categorisation.

Facial skin disorders have received less attention than general skin lesion identification, despite the fact that several experiments have shown good classification performance for both. There are extra obstacles to overcome when dealing with face clinical photos due to the wide range of lighting, facial orientation, skin pigmentation, lesion size, facial emotions, and surrounding anatomical features. Furthermore, there is a lack of publicly accessible dermatological datasets that include annotated photos of face skin in comparison to databases that include lesions acquired from other parts of the body. As a result, thorough dataset curation and feature-learning algorithms that can handle these variances are necessary for constructing specialised CNN models for face dermatology. Xiangya-Derm is a clinical image repository that includes professionally annotated photographs of six common skin disorders affecting the face.

This work uses a subset of facial images from this repository to solve these problems. Using the same experimental settings, the study thoroughly evaluates five existing CNN models rather than developing a whole new deep learning architecture. Furthermore, networks are trained utilising transfer learning by first being trained on clinical pictures acquired from various parts of the body, and then they are fine-tuned using photos of the face. By using this approach, models may improve their classification performance and diagnosis reliability by learning generalised dermatological aspects and successfully adjusting to face skin characteristics.

This study's comparative analysis sheds light on the efficacy of several CNN architectures for skin disease identification in faces. The study finds the architecture that provides the most dependable diagnostic performance for computer-aided facial dermatology by maintaining identical datasets, preprocessing procedures, training methodology, and evaluation metrics across all experiments. This allows for a fair assessment of each network's capability.

IV. FACIAL DERMATOLOGY IMAGE COLLECTION

Automated skin disease detection using deep learning models relies on a high-quality, well-annotated clinical picture collection. It is strongly related to the availability of relevant clinical pictures that the model's learning power and prediction accuracy are affected, as CNN-based classification algorithms rely substantially on various training examples. A specialised dataset for face skin diseases was built using the Xiangya-Derm clinical picture repository. This dataset was used to train and evaluate several CNN architectures in this work. Dermatologists with extensive expertise in the field hand-picked the dataset to guarantee reliable illness annotations and excellent clinical validity.

With 150,223 clinical images illustrating 543 distinct skin illnesses, the Xiangya-Derm repository is widely acknowledged as one of China's most extensive clinical dermatological image collections. Comprehensive

International Journal of Advanced Trends in Multidisciplinary Research
ISSN: 3139-0986, Volume 2 Issue 7
July 2026, PP 1-7

computer-aided diagnostic research is made possible by the inclusion of extensive clinical information with each picture, including medical history and pathology diagnosis. This study relies on a subset of 2,656 face clinical pictures taken from this archive. The six face skin illnesses shown in the photographs include basal cell carcinoma (BCC), actinic keratosis (AK), rosacea (ROS), seborrhoeic keratosis (SK), and squamous cell carcinoma (SCC). The research is able to assess CNN performance for disorders affecting one of the most visually sensitive and clinically significant parts of the human body by limiting its focus to pictures of the face.

Separate training and testing subsets were created using patients' names instead of pictures' names in order to keep the experimental assessment honest. By avoiding duplicate patient photos in both datasets, this technique reduces the risk of data leakage and guarantees a more accurate evaluation of the generalisability of the model. When used in a clinical setting, when diagnostic algorithms need to correctly categorise pictures taken of people they have never seen before, patient-level separation is more appropriate.

The chosen CNN architectures required conventional preprocessing procedures on all clinical pictures before to training. To make the model more resistant to changes in face orientation and picture arrangement, data augmentation methods including horizontal flipping and random cropping were used during training. In order to decrease overfitting and preserve visually important clinical information, these preprocessing approaches boost sample variety without changing the underlying illness features.

Table I summarises the distribution of face photos across the six illness categories. Model training used suitable weighting procedures to decrease classification bias, despite the dataset exhibiting class imbalance, especially for Lupus Erythematosus and Actinic Keratosis. Keeping the dataset's original composition allows for a fair comparison of all examined CNN architectures and guarantees consistency with the experimental design.

Table I. Distribution of Facial Skin Disease Images

Disease Type	Total Images	Facial Images	Training Images	Testing Images
Basal Cell Carcinoma (BCC)	689	578	623	100
Actinic Keratosis (AK)	277	108	219	58
Lupus Erythematosus (LE)	1273	781	1188	85
Rosacea (ROS)	318	183	263	55
Seborrheic Keratosis (SK)	1127	593	1075	52
Squamous Cell Carcinoma (SCC)	710	413	638	72
Total	4394	2656	4006	388

V. CNN-BASED CLASSIFICATION FRAMEWORK

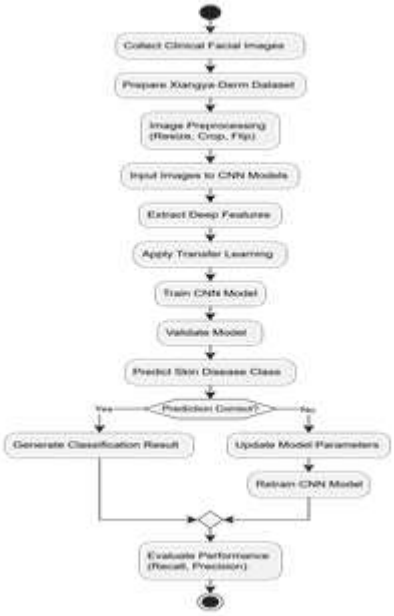

Fig. 1. Activity Diagram for CNN-Based Facial Skin Disease Classification

Automatic facial skin disease identification from clinical photos is achieved by means of the suggested classification framework that makes use of deep learning algorithms. This system uses Convolutional Neural Networks (CNNs) to learn hierarchical visual representations from the input photos instead of depending on manually generated image descriptors. The models are able to discriminate between dermatological disorders that are clinically identical because to their automated feature learning capabilities, which allows them to recognise complicated lesion features including colour change, texture, shape, and border information. The whole system includes steps for cleaning up images, extracting features from pretrained convolutional neural networks (CNNs), learning to transfer across them, and finally, illness categorisation.

In order to make sure the chosen CNN architectures work with each face picture, typical preprocessing is first used. To make the model more resistant to changes in face orientation and picture placement, image augmentation methods like horizontal flipping and random cropping are used during training. The dataset has an uneven distribution of samples across illness categories; to mitigate the impact of this imbalance and enhance learning for under-represented groups, the loss function incorporates class-weighting algorithms. By preserving the lesions' original clinical features, these preprocessing methods improve the deep learning models' generalisability.

The main difference between the CNN designs is the way they extract features. In order to train complicated models effectively, ResNet-50 incorporates residual connections that allow gradient propagation over deep networks.

Inception-v3 uses multi-scale convolutional procedures to concurrently record picture characteristics at several receptive fields. DenseNet121 promotes feature reuse and reduces duplicate parameter learning by establishing dense interconnections among succeeding layers. Xception improves computational efficiency while preserving excellent feature representation capacity by replacing standard convolution operations with depthwise separable convolutions. Lastly, Inception-ResNet-v2 integrates Inception modules with residual learning's benefits, enabling the network to acquire more accurate and detailed representations for face skin disease categorisation.

According to the first experimental setup, all networks use input photos with a resolution of 300 × 300. The fully connected layer that comes before the final convolution stage is substituted by Global Average Pooling (GAP) and then a 1 × 1 convolution in order to decrease the amount of trainable parameters without losing crucial spatial information. Next, the feature representations that were recovered are fed into a Softmax classifier that has a 1024-dimensional fully connected layer. This classifier then produces probability values that match the six categories of face skin diseases. The final prognosis is based on the illness label that is most likely to be correct. To maintain the original experimental technique, it is important that the core architectural components of each CNN model, such as dense blocks, Inception modules, Inception-ResNet blocks, transition layers, and residual blocks, do not change.

To enhance automated face skin disease detection while preserving the original experimental design, the suggested system integrates transfer learning with robust feature extraction. The work builds a consistent foundation for comparing classification performance and identifies the most effective model for computer-aided dermatological diagnosis by analysing different state-of-the-art CNN architectures under similar preprocessing and training settings.

VI. EXPERIMENTAL CONFIGURATION

Automated face skin disease classification using clinical photos was the intended goal of the experimental study, which aimed to examine the efficacy of several CNN architectures. To provide a fair comparison among the chosen deep learning models, all tests were carried out under the same circumstances. The classification capacity of each network is measured during the assessment, which takes into account the original dataset, preprocessing approach, and training methods. The limits and capabilities of each CNN architecture for computer-aided face dermatological diagnosis may be objectively evaluated using this standardised experimental design.

A subset of face images taken from the Xiangya-Derm clinical dataset was used for the research. In order to meet the input specifications of the chosen CNN architectures, all photos were downsized to 300 × 300 pixels prior to model training. In order to train a model that is resilient to changes in face orientation and picture composition, we used standard image augmentation methods like horizontal flipping and random cropping to enhance the variety of the images. Weighted loss functions were used to mitigate the effects of class imbalance and enhance learning for minority classes due to the dataset's differential sample sizes across illness categories. To keep things consistent throughout the experiment, these preprocessing activities were the same for all of the networks that were examined.

We chose five pretrained CNN architectures—ResNet-50, Inception-v3, DenseNet121, Xception, and Inception-ResNet-v2—to compare and contrast. After being trained with pretrained ImageNet weights, each model was transferred to the face skin disease dataset for fine-tuning. By making use of pretrained parameters, networks may adjust to dermatological picture properties while drawing on previously taught visual representations, leading to better convergence and better classification performance. In order to maintain consistency with the original implementation, all of the CNN models' original architectural components, such as residual blocks, Inception modules, dense blocks, transition layers, and depthwise separable convolutions, were kept unchanged throughout the tests.

Each network's classification layer minimises the amount of trainable parameters while keeping critical spatial information retrieved by the convolutional layers. It comprises of Global Average Pooling (GAP) followed by a 1 × 1 convolution. Next, the feature vectors are processed by a fully connected layer with 1024 dimensions. Lastly, they are fed into a Softmax classifier, which determines the probability distribution for each of the six categories of face skin diseases. The category with the greatest confidence score is used to determine the projected class. By standardising the classification process, we can be confident that variations in CNN architectures' feature extraction capabilities, and not in the classifiers themselves, are the root cause of performance disparities.

An independent test set of 388 clinical photos of the face was used to assess the trained models' diagnostic performance. We used Recall and Precision, two useful metrics for overall diagnostic capabilities and illness category specific classification efficacy, to evaluate the model's performance. By consistently using the same criteria for assessment in all tests, we can directly compare the five CNN architectures and find the best model for automated face skin disease categorisation.

The experimental setup makes use of transfer learning as well, by pretraining the convolutional neural network (CNN) models using clinical photos taken of various parts of the body and then refining them using images of faces. By using this approach, networks may learn about dermatology in general before they learn about features unique to individual faces. Next, we compare the performance of pretrained models with that of models trained just on face pictures to determine the efficacy of transfer learning. This analysis sheds light on the

advantages of knowledge transfer for computer-aided dermatological diagnosis.

VIII. CONCLUSION

Convolutional Neural Networks (CNNs) are shown to be successful in this research for automatically classifying face skin disorders using clinical photographs. This study tests five popular convolutional neural network (CNN) architectures—ResNet-50, Inception-v3, DenseNet121, Xception, and Inception-ResNet-v2—on a subset of facial photos taken from the Xiangya-Derm dataset. The goal is to determine which architecture is best at identifying six prevalent skin illnesses that affect the face. The experimental study verifies that deep learning models are capable of learning discriminative visual characteristics from clinical face photos, which paves the way for trustworthy computer-assisted dermatological diagnosis.

Overall, the tested models' classification performance was best by Inception-ResNet-v2, according to a comparison of the examined architectures. In addition, pretrained networks that use transfer learning are able to adapt to face skin disease categorisation by using generalised dermatological traits gained from pictures gathered across many body areas. This greatly increases diagnostic capabilities. The experimental results show that the original findings hold true, with an average recall of 77.0% and an average precision of 70.8%. On the independent facial-image test dataset, there are particularly strong recall values for Lupus Erythematosus (92.9%), Basal Cell Carcinoma (89.2%), and Seborrhoeic Keratosis (84.3%). In order to improve clinical image analysis using CNNs, these results prove that transfer learning is useful.

There are still certain obstacles to overcome, even if the suggested framework does a good job at diagnosing. On rare occasions, misclassification occurs due to an imbalance in the quantity of training samples across various disease categories and the resemblance in appearance between some skin conditions, especially Actinic Keratosis and Seborrhoeic Keratosis. Based on these findings, it seems that bigger and more evenly distributed face dermatology datasets, as well as deeper learning architectures that can learn more discriminative disease-specific characteristics, can lead to even better results.

In sum, the research shows that AI has great promise as a tool to aid clinical dermatology by analysing skin disorders of the face quickly, consistently, and objectively. The accuracy and clinical reliability of computer-aided diagnostic systems are anticipated to rise in tandem with the availability of bigger annotated datasets and the ongoing evolution of CNN architectures. Intelligent systems like these can help dermatologists make better decisions and make preliminary skin assessments more accessible through healthcare apps on smartphones and the cloud. This will lead to earlier diagnoses, better patient management, and better dermatological care overall.

REFERENCES

1. R. J. Hay et al., "The global burden of skin disease in 2010: An analysis of the prevalence and impact of skin conditions," Journal of Investigative Dermatology, vol. 134, no. 6, pp. 1527–1534, Jun. 2014.
2. C. Huang et al., "Association between daily soft drink consumption and moderate-to-severe acne vulgaris in Chinese adolescents," The Journal of Pediatrics, vol. 204, pp. 256–262, Jan. 2019.
3. X. Deng et al., "Development and validation of the Rosacea-specific Quality-of-Life instrument among Chinese patients," PLoS ONE, vol. 13, no. 2, pp. 1–13, Feb. 2018.
4. W. Chen et al., "Leptin-transgenic pig model reveals novel insights into systemic lupus erythematosus," PLoS Biology, vol. 16, no. 8, pp. 1–23, Aug. 2018.
5. Y. Xie et al., "F1F0-ATP synthase regulates keratinocyte differentiation in healthy and diseased skin," Scientific Reports, vol. 7, no. 1, pp. 1–11, Dec. 2017.
6. A. Bewley, "The neglected psychological aspects of skin disease," BMJ, vol. 358, p. j3208, Jul. 2017.
7. W. Chen et al., "SLCO1B1 and SLC22A1 polymorphisms associated with acitretin response in psoriasis patients," Scientific Reports, vol. 8, no. 1, pp. 1–8, 2018.
8. H. A. Haenssle et al., "Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition," Annals of Oncology, vol. 29, no. 8, pp. 1836–1842, Aug. 2018.
9. D. Avdeenko et al., "Deep learning techniques for skin cancer diagnosis," Journal of Healthcare Engineering, vol. 2020, Art. no. 8895279, 2020.
10. J. Zhang, Y. Xie, Y. Xia, and C. Shen, "Attention residual learning for skin lesion classification," IEEE Transactions on Medical Imaging, vol. 38, no. 9, pp. 2092–2103, Sept. 2019.
11. A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, Feb. 2017.
12. N. Codella et al., "Skin lesion analysis toward melanoma detection using deep learning network ensembles," Journal of Biomedical and Health Informatics, vol. 23, no. 2, pp. 501–510, Mar. 2019.
13. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 770–778.
14. C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception architecture for computer vision," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 2016, pp. 2818–2826.

15. G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 2017, pp. 4700–4708.
16. F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 2017, pp. 1800–1807.
17. C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, "Inception-v4, Inception-ResNet and the impact of residual connections on learning," in Proc. AAAI Conf. Artificial Intelligence, Phoenix, AZ, USA, 2017, pp. 4278–4284.
18. J. Deng et al., "ImageNet: A large-scale hierarchical image database," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Miami, FL, USA, 2009, pp. 248–255.
19. O. Russakovsky et al., "ImageNet Large Scale Visual Recognition Challenge," International Journal of Computer Vision, vol. 115, no. 3, pp. 211–252, Dec. 2015.
20. S. Minaee et al., "Image segmentation using deep learning: A survey," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 44, no. 7, pp. 3523–3542, Jul. 2022.
21. G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, Dec. 2017.
22. S. R. Kassani, P. H. Kassasni, M. J. Wesolowski, K. A. Schneider, and R. Deters, "Automatic detection of coronavirus disease (COVID-19) in X-ray and CT images: A machine learning based approach," Biomedical Signal Processing and Control, vol. 68, p. 102559, Jul. 2021.
23. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Proc. Int. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI), Munich, Germany, 2015, pp. 234–241.
24. M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in Proc. 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Savannah, GA, USA, 2016, pp. 265–283.
25. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learning Representations (ICLR), San Diego, CA, USA, 2015.
26. P. Ramesh Babu, “Enhancing Brain Tumor Classification Through Feature Selection with Beetle-Swarm Optimization,” International Journal of Electrical and Electronics Research, vol. 12, no. 4, pp. 1399–1406.
27. P. Ramesh Babu, “Intelligent Disease Prediction and Personalized Health Management System,” Nanotechnology Perceptions, vol. 20, no. S6, pp. 907–920, 2024.