

Advances in dermatological imaging: enhancing skin melanoma classification for improved patient outcomes

Debadutta Sahoo, Soumya Mishra

Department of Electronics and Communication Engineering, C.V. Raman Global University, Bhubaneswar, India

Article Info

Article history:

Received Sep 18, 2025
Revised Dec 5, 2025
Accepted Jan 1, 2026

Keywords:

Alex net architecture
Deep learning
Diagnostic accuracy
Medical image analysis
Melanoma detection
Skin cancer classification

ABSTRACT

The study presents an enhanced AlexNet-based deep learning system for binary classification of melanoma skin cancer as either benign or malignant using two paired dermoscopic and clinical image datasets. The study evaluates the resilience of the models across different image sets with common preprocessing and specific data augmentation, using a melanoma dataset containing 10,000 images and a benign versus malignant dataset with 3,600 images. The AlexNet refinement exceeded several standard machine learning (ML) classifiers and other deep architectures on the two datasets with practical training times, gaining 97.12% and 96.21% in balanced accuracy. The training proceeded with SGD as optimiser and cross-entropy as loss on 256×256 images. Benchmarking against support vector machine (SVM), k-nearest neighbour (KNN), and other convolutional neural networks (CNNs) designs shows that the selected architecture and hyperparameters achieved the highest performance on cost-effective computation for the routine melanoma triage. The report highlights the need for external validation, incorporation into dermatological workflows, and explainability to improve trust, diminish dataset bias, and support the safe clinical deployment in practice.

This is an open access article under the [CC BY-SA](#) license.



Corresponding Author:

Debadutta Sahoo
Department of Electronics and Communication Engineering, C.V. Raman Global University
Bhubaneswar 752054, India
Email: debaduttasahoo2255@gmail.com

1. INTRODUCTION

Skin cancer is one of the most prevalent tumours worldwide, and melanoma typically makes news due to its rapid spread and high death rate. Millions of new cases of skin cancer are reported by the WHO each year, underscoring the importance of early detection and effective treatment. If the initial symptoms are disregarded, melanoma, which begins in pigment-producing melanocytes, can turn fatal [1], [2]. Doctors now have sharper optics and additional eyes to detect skin issues thanks to exciting advancements in medical imaging and machine learning (ML) tools [3]. On the other hand, routine inspections that mostly rely on a specialist's glance may occasionally overlook or misread small signs. In this case, deep learning techniques—particularly convolutional neural networks, or CNNs—fill that need as trustworthy helpers for skin photo tagging [2], [4], [5]. The popular AlexNet architecture has already been modified for this specific task; it was once revolutionary in routine picture searches [1]. The current effort focuses on two primary categories of skin cancer: the ostensibly benign but obviously cancerous moles. The team aims to increase diagnosis rates while maintaining a manageable level of false alarms by optimising a pretrained version of AlexNet. In order to improve dermatologists' clinical practice skills, the current study intends to add to this expanding body of knowledge by showcasing how well the AlexNet model classifies skin cancer [1], [2], [6]–[10].

2. RELATED WORK

The advancements in melanoma detection have resulted in a surge in cancer diagnostic computing research. The first systems utilized classical image analysis and ML procedures in which features like colour and texture, as well as lesion shape, were created by domain experts and classified by support vector machines (SVMs), k-nearest neighbours (KNNs), decision trees, random forests, naive Bayes, and functional link artificial neural network (FLANN) [11], [12]. These techniques, while baselines were useful, suffered from poor performance due to considerable intra-class variability and image noise. The dead images were often present in low-quality dermatoscopic and clinical photographs. The precise and timely diagnosis of melanoma and skin lesion categorization, which deep learning excels in analyzing, came from deployable CNNs. Over the years, studies have documented almost every architecture to be competing with a dermatologist in performance, if not outperforming them [13]–[18]. Typical CNNs, LeNet, VGG16, ResNet, GoogLeNet, Inception, GANs, and the recent transformer-inspired networks have all been studied. These approaches, however, require not only a wealth of annotated datasets and high-end computers but also often draw upon considerable learning to finalize a solution and capture images with complex visual patterns that would be almost impossible to describe manually [2], [19].

Several works continue to be slowed down because of the same practical issues that make translation challenging: limited access to diverse and well-annotated dermatoscopic images, privacy and regulatory constraints of medical data sharing, and difficulties in building costly and scalable data ingestion, annotation, training, and validation pipelines. Consequently, more recent works have developed new models, but also attempted to build the requisite infrastructure and legal frameworks that would enable the use of AI in clinical practice while safeguarding anonymity [20]–[22]. However, most of the available and developed research continues to examine a very narrow range of ML baselines, or a singular deep learning architecture, and typically, using a single dataset with different data preparation/evaluation methodologies [23], [24]. Consequently, there is little evidence to support the optimal models and the trade-off between structure, flexibility, and cost, and there are little evidence and guidance available to physicians to help them select the most appropriate technologies for clinical melanoma screening [25].

The current work addresses this gap by performing a methodical, side-by-side comparison under the same experimental framework among several conventional and deep learning approaches for skin cancer classification. This article assesses a fine-tuned AlexNet against multiple classic ML classifiers, such as SVM, KNN, random forest, and fast library for approximate nearest neighbours, and other deep networks, including a standard CNN, LeNet, VGG16, ResNet, GoogLeNet, and a generative adversarial network-based model, adopting a coherent preprocessing, training, and testing procedure on two complementary publicly available datasets: the melanoma skin cancer dataset, consisting of 10,000 images, and the skin cancer malignant vs. benign dataset. The article allows for a more complete benchmark by jointly considering accuracy, AUC, precision, recall, F1-score, and training time on both datasets and also demonstrates that an optimized AlexNet can achieve superior diagnostic performance while still being computationally viable for integration into practical dermatology workflows [26].

3. OBJECTIVES

Evaluate deep learning techniques: to assess the effectiveness of deep learning architectures, particularly the AlexNet model, in classifying skin cancer from medical images. Compare with ML models: to compare the performance of the AlexNet model against traditional ML classifiers, such as SVM, KNN, and random forests, in terms of accuracy, precision, recall, and F1 score. Utilize diverse datasets: to utilize two distinct datasets, namely the melanoma skin cancer dataset and the skin cancer: malignant vs. benign dataset, to ensure a comprehensive evaluation of the model's performance across varying image characteristics. Optimize model performance: to fine-tune the AlexNet model by optimizing hyperparameters, including learning rate, batch size, and image preprocessing techniques, to achieve the highest possible classification accuracy and efficiency. Visualize results: to present the performance metrics and model architecture through flowcharts and graphs, facilitating a clear understanding of the model's structure and its classification capabilities. Contribute to clinical applications: to contribute to the field of dermatology by providing insights into the potential of AI-driven diagnostic tools for early detection and accurate classification of melanoma skin cancer, ultimately aiming to enhance patient outcomes [27].

4. METHODOLOGY

In this study, the AlexNet architecture—a pretrained deep learning model—was used and refined for the binary classification of melanoma skin cancer into benign and malignant groups [6], [7]. The size that the

network requires, 256 by 256 pixels, was applied to each picture that was pushed into the pipeline. A fast two-unit linear layer that spits either benign or malignant material was used to complete the task, replacing the last fully connected layer [1], [6], [7]. To ensure smooth and quick learning, pixels were then normalised using the traditional ImageNet mean and standard deviation values (0.485, 0.456, 0.406) / (0.229, 0.224, 0.225). Stochastic gradient descent (SGD) at 0.001 with 0.9 momentum was used for training [1], [28]; this pair accelerates progress while controlling noisy jumps. During training, cross-entropy loss was used as the objective function to measure classification mistakes efficiently [28]. Together, these hyperparameters and architectural modifications allowed the AlexNet model to attain high classification accuracy with reliable training results and realistic execution times, establishing it as a reliable method for diagnosing melanoma skin cancer [6], [7], [29].

5. DATASET DETAILS

The first collection of dermatoscopic images is taken from Kaggle's melanoma skin cancer dataset of 10,000 images, which has a size of 300×300 pixels [6]. The model has a lot of visual detail to work with because the photos are clearly divided into groups that are benign and malignant, as in Figure 1. The remaining 1,000 images are held back for final testing, while the remaining 9,600 are designated for training [6], [8]. The second collection, which can be found in the skin cancer: malignant vs. benign Kaggle project, contains 3,600 clinical images of size 224×224 pixels that are evenly divided between 1,800 benign and 1,800 malignant moles [7]. The combination of the two datasets pushes the system to address issues that it will encounter in actual clinics by displaying cancer on a range of skin tones, sizes, colours, and textures. These image collections are being used because doctors urgently need a way to spot melanoma early and get patients the right care fast [1], [6], [7].

Melanoma is dangerous—it spreads quickly and can be deadly if the diagnosis comes too late. When clinicians search for early melanoma indications, the pixels in the first collection—which was crowdsourced by Kaggle users—show minute details from dermatoscopic photographs taken with expensive skin-scanning devices. Every image has a thorough label that is either human or malware-formed following clinical examinations and microscopy slides [6], [7], [16]. This allows students to rely on such labels when they develop new supervised code. As the provider observes, the second collection, which was borrowed from the ISIC archive, collects regular clinic photos of skin patches, combining benign and malignant cancer in about equal proportions [1], [7]. The teams removed camera artefacts and size mismatches by trimming, resizing, and normalising the pixels to make each batch easier to input into neural nets or vintage machines. Combining the two sources enables researchers to create models that learn more comprehensive patterns than any one clinic can identify.

Data splitting and evaluation strategy: in this implementation, all images in the generated dataset were placed into a single folder structure and loaded through an image-folder interface. During training, those images were then shuffled randomly and presented repeatedly to the network. Given metrics reflect performance on the training distribution rather than a strictly independent validation set, because the same set of images was used to calculate accuracy, precision, recall, and F1-score at the end of training. Although this setup gives an initial idea of how well the fine-tuned AlexNet can fit the melanoma classification task, it can also inflate generalization since the model is tested on photos that it has already seen during tuning. The next steps in this development will focus on separating training, validation and test training sets that are not overlapping—for example 70%, training 15%, validation 15% testing efforts could be employed to avoid bias and have an unbiased representation of your algorithm's true applicability within an actual production environment. Additionally, implementing cross-validation across multiple different dataset sets can provide you with an unbiased assessment of your algorithm's performance and add to the reliability of the estimates you provide.




Figure 1. Sample image from the combined dataset illustrating benign and malignant skin cancer

6. DATA PRE-PROCESSING

The pre-processing stage of the data was essential for getting the melanoma skin cancer and malignant versus benign mole image datasets ready for efficient training and the best possible performance from the AlexNet model, as in Figure 2. The following are important pre-processing actions and specifics of hyperparameter optimisation:

- i) Image standardisation: all input images were resized uniformly to 256×256 pixels to maintain consistent dimensions suitable for the AlexNet architecture. The images were normalised using channel-wise mean values of (0.485, 0.456, 0.406) and standard deviations of (0.229, 0.224, 0.225), aligning with the pretrained ImageNet model's expected input distribution [30].
- ii) Data augmentation: to improve dataset diversity and decrease overfitting, training methods like random rotation, horizontal and vertical flipping, scaling, shearing, and translation were used [2], [17], [23], [31].
- iii) Training configuration: the final fully-connected layer of the pretrained AlexNet model was adjusted to provide two classes (malignant and benign). Using SGD as the optimiser, the momentum was set at 0.9, and the learning rate was set at 0.001. Cross-entropy loss was the classification loss function, which was suitable for jobs involving binary classes.
- iv) Model optimisation: outperforming frameworks for deep learning for the two datasets, the 10,000-images melanoma skin cancer dataset had a model accuracy of 97.12% after 6453 seconds of training. The skin cancer: malignant vs. benign dataset, in contrast, demonstrated the efficacy and efficiency of the AlexNet model in skin cancer classification with an accuracy of 96.21% and a training time of 1700 seconds.




Figure 2. Overview of the methodology: data processing, augmentation, training

7. EXPERIMENTS

This research presents a comprehensive analysis of how various skin cancer spots can be sorted using deep learning and conventional ML techniques described in Figure 3 [1]–[4], [6], [7], [23]. Two image pools were employed in the test runs: the malignant-versus-benign set and the 10,000-image melanoma skin cancer dataset [6], [7]. To identify the method that best classifies cancer, models from both families were scored side by side [25], [32].

7.1. Deep learning networks

Standard CNNs, ResNet, LeNet, GoogLeNet, VGG16, InceptionV3, GANs, and AlexNet were among the powerful deep learning models that the researchers investigated [1]–[4], [11], [23]. The skin cancer datasets were used to refine each model, which began with pretrained weights from AlexNet. Accuracy and training pace varied significantly across systems, despite the fact that nearly all of them generated positive results. For example, LeNet and VGG16 were faster but trailed behind, whereas ResNet and GoogLeNet achieved respectable accuracies. When combined, the findings unequivocally demonstrate that overall design, depth, and complexity influence—and occasionally impede—skin cancer categorisation performance [33], [34].

7.2. Machine learning networks

SVMs, KNNs, decision trees, random forests, naive Bayes, and flexible local approximate nearest neighbour are some of the well-known classifiers that the team also examined and visualised [1], [2], [4], [5], [11], [13], [14], [20], [21], [28], [31]. Before making a final approximation, these models relied on manually created features taken from a typical image-processing pipeline. Deeper neural networks won the

gold medal despite SVM and KNN slipping into the mid-80% accuracy range because of their superior ability to handle the noise and fine details present in skin cancer images [1]. Nevertheless, the older algorithms established reliable baselines and served as a reminder of the advantages and disadvantages of the traditional toolkit.




Figure 3. Visual representation of image preprocessing steps: from raw input to normalization

7.3. Proposed model

AlexNet comprises eight layers in total, with three fully connected layers coming after five convolutional layers, as shown in Figure 4 [1]. Together, they enable the model to pick up different details in the depth of unprocessed photos. The layer is followed by a 3×3 max pooling layer with a stride of 2. As it enhances generalisation and decreases the spatial dimensions of the feature maps, it captures the most prominent features. The model learns the most relevant information or features of the picture, which reduces the computing load in addition to learning to generalise factually [1], [2], [6], [7].

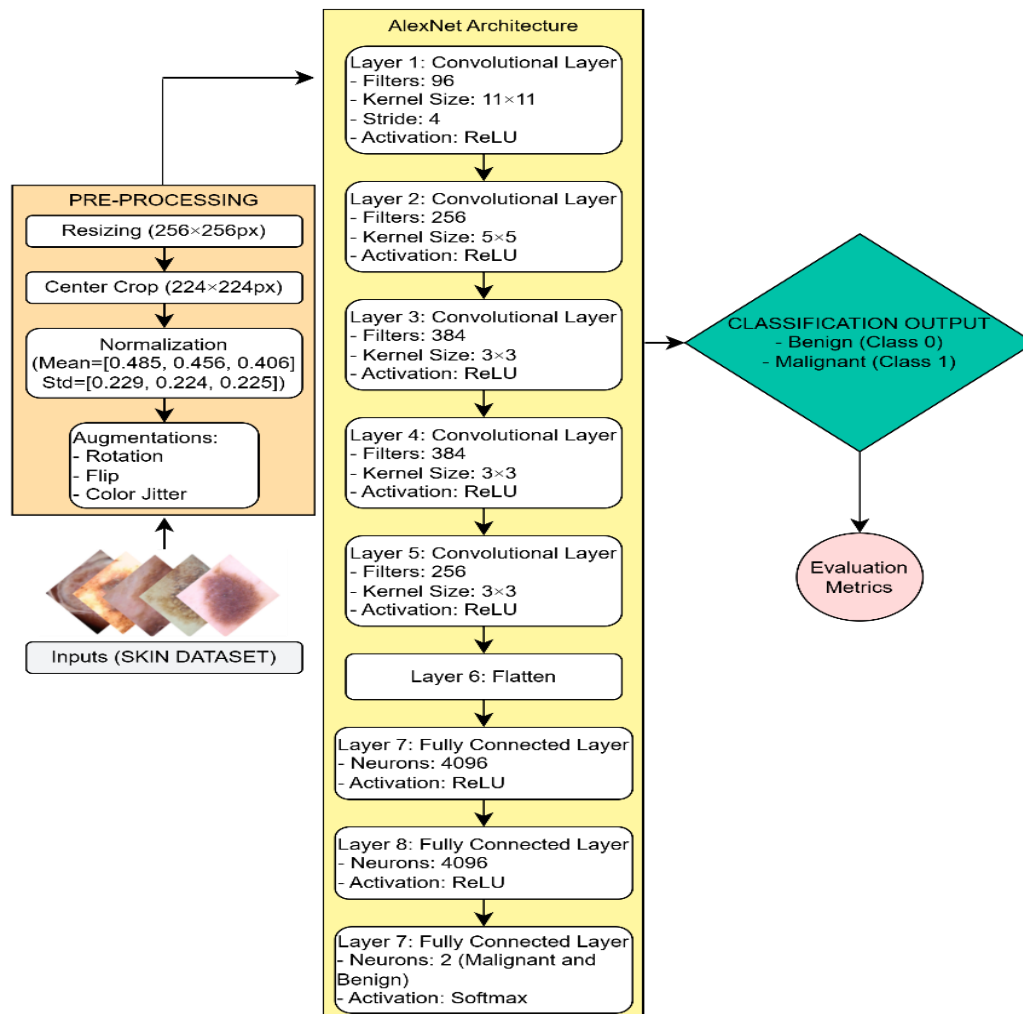


Figure 4. Overview of the proposed model architecture and evaluation process

For instance, the second layer contains 256 5×5 filters, while the third, fourth, and fifth layers have 384 and 256 3×3 filters, respectively. Together with the first layer, all of these layers carefully describe the rough to exact patterns or shapes, which is crucial for determining if a skin disease is benign or cancerous [6], [7]. Rectified linear unit (ReLU) activation functions are also incorporated into all of these levels; these activation functions are consistent throughout the conflicting layers. Since the AlexNet model achieves a good balance of accuracy, precision, recall, F1 score, and manageable training time, we decided to employ it. The network produced remarkable results after adjusting a few hyperparameters for each dataset: a 0.001 learning rate with 0.9 momentum using SGD, a 256×256 -pixel input size, and cross entropy loss.

On the melanoma set, the accuracy was 97.12% in 6453 seconds, and on the malignant-versus-benign set, it was 96.21% in 1700.85 seconds. These impressive results, combined with reasonable training times, imply that AlexNet may be a dependable tool for routine clinical diagnoses. The model's effectiveness is further demonstrated by the training times of 6453 seconds for the melanoma dataset and 1700.85 seconds for the malignant vs. benign dataset, which make it a viable option for practical dermatology applications. The findings of this study not only highlight the AlexNet model's potential to improve diagnostic precision, but they also add to the expanding corpus of research on the use of AI in medical image processing.

8. RESULTS AND DISCUSSION

The new AlexNet model consistently and accurately detects melanoma spots as either benign or malignant, according to tests conducted on two distinct collections of skin cancer images. On the bigger melanoma skin cancer dataset, it achieved 97.12% accuracy and 96% AUC after approximately 6453 seconds. On the balanced Malignant vs Benign set, it recorded 96.21% accuracy and 96% AUC in approximately 1700 seconds, visualised in Figure 5.




Figure 5. Average time complexity comparison of the deep learning models in seconds

The model demonstrated its strength when presented with a variety of clinical images by producing balanced precision, recall, and F1 scores across both runs, as displayed in Table 1. In terms of raw accuracy, area-under-curve values, and training speed, AlexNet outperformed several more recent deep learning and conventional ML approaches in head-to-head comparisons, demonstrating the importance of the selected architecture and adjusted hyperparameters. When combined, the findings imply that the system may be a reliable tool for early, precise melanoma identification, which would enhance physician judgment and, eventually, patient outcomes [6], [35]. Figure 6 illustrates the average performance analysis of ML models evaluated using two distinct skin cancer datasets.

Table 1. Average performance comparison of the two datasets across evaluation metrics

DL model	Accuracy (%)	Precision (%)	Recall (%)	F1 score (%)
AlexNet	97.12	96	95	97
CNN	91	95	85	90
LeNet	88.45	46.79	47	46.95
GoogLeNet	87	47	51	49
SNN	87	86	87	87
FLANN	86	89	80	84
ResNet	66.86	47	16	24
GAN	53	52	67	58
VGG16	51	50	44	47




Figure 6. Average performance analysis of ML models using two distinct skin cancer datasets

9. LIMITATIONS AND FUTURE DIRECTIONS

Noteworthy limitations must be taken into consideration regarding the said AlexNet-based methodology, despite the accuracy, AUC, and balanced precision–recall achieved on two publicly accessible skin cancer datasets being relatively good. To begin with, the melanoma skin cancer dataset and skin cancer malignant vs. benign dataset are culled from particular online repositories and, therefore, may not accurately represent the variation of the real-world clinical population, imaging devices, and processes of acquisition. Under-representation of certain skin tones, the early-stage melanomas, or lesion type, may lead to dataset bias and limit the applicability of the reported performance to other hospitals or regions [1]. In addition, such studies often use the labels created by dataset developers. While these labels may be of poor quality, there is no way to control the differences that may exist within the dataset itself during the course of the study. Even with the popularity of these datasets, we suspect that there is no true ground for the metrics used, and thus no true basis for the metrics, due to the existence of what we have called ‘labelling noise’, differences related to the histopathology approval at hand, and the control of the diagnosis. Also, the study lacks tests on a more significant domain shift, such as images taken from a variety of dermatoscopes or smartphone cameras, as the focus of the assessment is on two similar datasets. This may indicate a gap in the study concerning such a domain shift.

Lastly, none of the trials have been done in the setting of a future study on readers (dermatologists or other healthcare providers). Therefore, many of the important issues, including trust in the model by the users, how interpretable and how well workflows can be incorporated into daily practice, and how much AI assistance affects the speed of diagnoses and how quickly decisions can be made based on diagnosis, were not tested. The current model only does an image-level classification and does not yet do lesion segmentation or provide structured explanations that would have been useful in proving the clinical validity of the models.

The limitations of our findings provide many avenues for future research. Future research should corroborate our model's results by validating it against independent, multi-centre datasets with more extensive variations in skin types, lesion types, and image acquisition methods, as well as by utilizing cross-dataset evaluation methods to quantify the impact of domain shift. Additionally, uncertainty estimation and calibration should be incorporated into future studies to allow better quantification of the certainty associated with individual predictions, particularly for borderline or infrequently observed instances. Lastly, future studies should provide for the incorporation of explanations of the AI's decisions through the use of

techniques like saliency maps and class-activation visualizations, and the conduct of prospective studies employing human-AI collaboration in actual clinical settings [4], [36] before AI-based decision-support systems can be implemented and integrated into clinical practice.

10. CONCLUSION

This study confirms that fine-tuned AlexNet excels at classifying skin cancer images, achieving 97.12% accuracy on the melanoma skin cancer dataset and 96.21% on the skin cancer malignant vs. benign dataset, while outperforming traditional ML methods and several competing deep networks. The model should be incorporated into the dermatology triage process with saliency maps for clinical use so that the model is transparent, validated through multicentre trials across various populations, and used as a collaborative tool to facilitate early detection of melanoma and improve patient care. In dermatological practice, where our model will serve as a triaging tool, we propose that the clinical implementation of our model will occur in the future, with an emphasis on using saliency mapping techniques to highlight the relevant visual features that aid in the diagnostic process. To validate the use of our model in the clinical setting, we plan to conduct multi-site studies that include a large number of patients accessed by multiple photographic capture devices. Our model acts as a supplementary tool to aid dermatologists; thus, both community and practice experiences will support our model's use to expedite melanoma detection and improve the standard of care practices for all patients.

FUNDING INFORMATION

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Debadutta Sahoo	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	
Soumya Mishra	✓	✓		✓	✓	✓	✓		✓	✓	✓	✓		

C : **C**onceptualization

M : **M**ethodology

So : **S**oftware

Va : **V**alidation

Fo : **F**ormal analysis

I : **I**nterpretation

R : **R**esources

D : **D**ata Curation

O : Writing - **O**riginal Draft

E : Writing - Review & **E**ditng

Vi : **V**isualization

Su : **S**upervision

P : **P**roject administration

Fu : **F**unding acquisition

CONFLICT OF INTEREST STATEMENT

The authors state no conflict of interest.

DATA AVAILABILITY

Data availability does not apply to this paper as no new data were created or analyzed in this study.

REFERENCES

- [1] K. Sudharson, R. Santhiya, and S. Varsha, "A deeply connected model for high-precision classification of benign and malignant skin lesions," in *2025 International Conference on Advanced Computing Technologies (ICoACT)*, Mar. 2025, pp. 1–6, doi: 10.1109/ICoACT63339.2025.11004814.
- [2] N. Darapaneni *et al.*, "Early detection of skin cancer - solution for identifying and defining skin cancers using AI," in *2022 Interdisciplinary Research in Technology and Management (IRTM)*, Feb. 2022, pp. 1–6, doi: 10.1109/IRTM54583.2022.9791676.
- [3] International Agency for Research on Cancer World Health Organization, "Global cancer observatory," International Agency for Research on Cancer World Health Organization, 2022.
- [4] P. Kaushik, Y. Chopra, A. Kajla, M. Poonia, A. Khan, and D. Yadav, "AI-powered dermatology: achieving dermatologist-grade skin cancer classification," in *2024 IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI)*, 2024, doi: 10.1109/IATMSI60426.2024.10502664.
- [5] A. Esteva *et al.*, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639,

- pp. 115–118, Feb. 2017, doi: 10.1038/nature21056.
- [6] M. H. Javid, “Melanoma skin cancer dataset of 10000 images,” Kaggle. [Online]. Available: <https://www.kaggle.com/datasets/hasnainjaved/melanoma-skin-cancer-dataset-of-10000-images>
 - [7] C. Fanconi, “Skin cancer: malignant vs. benign,” Kaggle. [Online]. Available: <https://www.kaggle.com/datasets/fanconic/skin-cancer-malignant-vs-benign>
 - [8] O. Akinrinade and C. Du, “Skin cancer detection using deep machine learning techniques,” *Intelligence-Based Medicine*, vol. 11, 2025, doi: 10.1016/j.ibmed.2024.100191.
 - [9] G. S. Virgens, J. A. Teodoro, E. Iarussi, T. Rodrigues, and D. T. Amaral, “Enhancing and advancements in deep learning for melanoma detection: a comprehensive review,” *Computers in Biology and Medicine*, vol. 194, 2025, doi: 10.1016/j.combiomed.2025.110533.
 - [10] A. A. Shirsath, T. Shah, and V. A. Shirsath, “Deep learning based system for early detection and classification of melanoma skin cancer—a systematic review,” *Journal of Carcinogenesis*, vol. 24, no. 2, pp. 26–34, 2025, doi: 10.64149/J.Carcinog.24.2.26-34.
 - [11] S. Jain, V. Jagtap, and N. Pise, “Computer aided melanoma skin cancer detection using image processing,” *Procedia Computer Science*, vol. 48, pp. 735–740, 2015, doi: 10.1016/j.procs.2015.04.209.
 - [12] E. Vocaturo, D. Perna, and E. Zumpano, “Machine learning techniques for automated melanoma detection,” *Proceedings-2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 2310–2317, 2019, doi: 10.1109/BIBM47256.2019.8983165.
 - [13] A. Mahbod, G. Schaefer, I. Ellinger, R. Ecker, A. Pitiot, and C. Wang, “Fusing fine-tuned deep features for skin lesion classification,” *Computerized Medical Imaging and Graphics*, vol. 71, pp. 19–29, 2019, doi: 10.1016/j.compmedimag.2018.10.007.
 - [14] H. A. Haenssle et al., “Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists,” *Annals of Oncology*, vol. 29, no. 8, pp. 1836–1842, Aug. 2018, doi: 10.1093/annonc/mdy166.
 - [15] H. Bhatt, V. Shah, K. Shah, R. Shah, and M. Shah, “State-of-the-art machine learning techniques for melanoma skin cancer detection and classification: a comprehensive review,” *Intelligent Medicine*, vol. 3, no. 3, pp. 180–190, 2023, doi: 10.1016/j.imed.2022.08.004.
 - [16] F. Alam, A. Ullah, D. Shah, S. Ali, and M. Tahir, “Artificial intelligence in melanoma detection: a review of current technologies and future directions,” *International Journal of Intelligent Systems*, no. 1, Jan. 2025, doi: 10.1155/int/3164952.
 - [17] M. Kreouzi et al., “Deep learning for melanoma detection: a deep learning approach to differentiating malignant melanoma from benign melanocytic nevi,” *Cancers*, vol. 17, no. 1, Dec. 2024, doi: 10.3390/cancers17010028.
 - [18] H. L. Gururaj, N. Manju, A. Nagarjun, V. N. M. Aradhy, and F. Flammini, “DeepSkin: a deep learning approach for skin cancer classification,” *IEEE Access*, vol. 11, pp. 50205–50214, 2023, doi: 10.1109/ACCESS.2023.3274848.
 - [19] U. B. Ansari and T. Sarode, “Skin cancer detection using image processing,” *International Research Journal of Engineering and Technology (IRJET)*, vol. 4, no. 4, pp. 2875–2881, 2017.
 - [20] V. Gulshan et al., “Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs,” *JAMA - Journal of the American Medical Association*, vol. 316, no. 22, pp. 2402–2410, 2016, doi: 10.1001/jama.2016.17216.
 - [21] J. De Fauw et al., “Clinically applicable deep learning for diagnosis and referral in retinal disease,” *Nature Medicine*, vol. 24, no. 9, pp. 1342–1350, Sep. 2018, doi: 10.1038/s41591-018-0107-6.
 - [22] T.-M. Chiu, Y.-C. Li, I.-C. Chi, and M.-H. Tseng, “AI-driven enhancement of skin cancer diagnosis: a two-stage voting ensemble approach using dermoscopic data,” *Cancers*, vol. 17, no. 1, 2025, doi: 10.3390/cancers17010137.
 - [23] S. Aburass, O. Dorgham, J. Al Shaqsi, M. A. Rumman, and O. Al-Kadi, “Vision transformers in medical imaging: a comprehensive review of advancements and applications across multiple diseases,” *Journal of Imaging Informatics in Medicine*, 2025, doi: 10.1007/s10278-025-01481-y.
 - [24] M. Stefańczyk and T. Bocheński, “Mixing deep learning with classical vision for object recognition,” *Journal of WSCG*, vol. 28, no. 1–2, pp. 147–154, 2020, doi: 10.24132/JWSCG.2020.28.18.
 - [25] R. Sabir and T. Mehmood, “Classification of melanoma skin cancer based on image data set using different neural networks,” *Scientific Reports*, vol. 14, no. 1, Nov. 2024, doi: 10.1038/s41598-024-75143-4.
 - [26] P. N. S. C. Teja, T. B. Krishna, A. K. R. Poreddy, and P. Kokil, “Lightweight deep learning model for melanoma classification in dermoscopy images for smart healthcare,” in *2024 International Conference on Wireless Communications Signal Processing and Networking (WiSPNET)*, Mar. 2024, pp. 1–6, doi: 10.1109/WiSPNET61464.2024.10532923.
 - [27] M. F. Almuftareh, N. Tariq, M. Humayun, and F. A. Khan, “Melanoma identification and classification model based on fine-tuned convolutional neural network,” *Digital Health*, vol. 10, Sep. 2024, doi: 10.1177/20552076241253757.
 - [28] Y. Qin et al., “Autofocus layer for semantic segmentation,” *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2018, pp. 603–611, doi: 10.1007/978-3-030-00931-1_69.
 - [29] Y. Mousa, R. Taha, R. Kaur, and S. Afifi, “Melanoma classification using deep learning,” in *Proceedings, Image and Video Technology: 11th Pacific-Rim Symposium, PSIVT 2023, Auckland, New Zealand, November 22–24, 2023*, 2024, pp. 259–272, doi: 10.1007/978-981-97-0376-0_20.
 - [30] G. Veronesi, N. Curti, A. Gardini, G. Querzoli, G. Castellani, and E. Dika, “Machine learning to detect melanoma exploiting nuclei morphology and spatial organization,” *Scientific Reports*, vol. 15, no. 1, Jul. 2025, doi: 10.1038/s41598-025-02913-z.
 - [31] N. Darapaneni et al., “COVID 19 severity of pneumonia analysis using chest x rays,” *2020 IEEE 15th International Conference on Industrial and Information Systems (ICIIS) Proceedings*, pp. 381–386, 2020, doi: 10.1109/ICIIS51140.2020.9342702.
 - [32] M. Q. Khan et al., “Classification of melanoma and nevus in digital images for diagnosis of skin cancer,” *IEEE Access*, vol. 7, pp. 90132–90144, 2019, doi: 10.1109/ACCESS.2019.2926837.
 - [33] H. C. Reis and V. Turk, “Fusion of transformer attention and CNN features for skin cancer detection,” *Applied Soft Computing*, vol. 164, 2024, doi: 10.1016/j.asoc.2024.112013.
 - [34] V. Venugopal, N. I. Raj, M. K. Nath, and N. Stephen, “A deep neural network using modified EfficientNet for skin cancer detection in dermoscopic images,” *Decision Analytics Journal*, vol. 8, Sep. 2023, doi: 10.1016/j.dajour.2023.100278.
 - [35] N. Melarkode, K. Srinivasan, S. M. Qaisar, and P. Plawiak, “AI-powered diagnosis of skin cancer: a contemporary review, open challenges and future research directions,” *Cancers*, vol. 15, no. 4, 2023, doi: 10.3390/cancers15041183.
 - [36] A. A. Louison and B. B. Sujitha, “MOON-NET: deep learning based melanoma skin lesion classification using SE-DNet and DBN,” in *2025 IEEE International Conference on Compute, Control, Network & Photonics (ICCCNP), Bangalore, India*, 2025, pp. 1–6, doi: 10.1109/icccnp63914.2025.11233427.

BIOGRAPHIES OF AUTHORS

Debadutta Sahoo     is an enthusiastic electronics and communication engineer skilled in PCB design, VLSI, and signal processing. As a licensed HAM operator, I thrive in R&D and troubleshooting within embedded systems. Proficient in C and Python, His passionate about leveraging AI to innovate and enhance technology for effective problem-solving in the engineering domain. He can be contacted at email: debaduttasahoo2255@gmail.com.



Soumya Mishra     is a seasoned professional with over 16 years of diverse experience encompassing teaching, research and development (R&D), and administration. Currently serving as the IQAC institutional coordinator, custodian of space club-CGU, and International Programs & Outreach Coordinator, she brings extensive expertise in enhancing academic quality assurance and fostering global partnerships. Her research specialization lies in wireless communication, signal processing, and the application of artificial intelligence/machine learning (AI/ML) in these domains. She can be contacted at email: soumya@cgu-odisha.ac.in.