Journal of Engineering Issues and Solutions

Skin Cancer Classification Using Transfer Learning With MobileOne: A Deep Learning Approach

Dilli Hang Rai[1], Sabin Kafley[2*], Om Prakash Dhakal[3]

[1] *Department of Computer Science and Information Technology, Tribhuvan University, Institute of Science and Technology (IoST), Godawari College (Affiliated to TU), Itahari, Nepal. email: dillihangrai.078@godawari.edu.np*

[2] *Department of Electronics and Computer Engineering, Tribhuvan University, Institute of Engineering, Purwanchal Campus, Dharan, Nepal. email: sabinkafley@ioepc.edu.np*

[3] *Department of Electronics and Computer Engineering, Tribhuvan University, Institute of Engineering, Purwanchal Campus, Dharan, Nepal. email: omprakash@ioepc.edu.np*

**Corresponding email: sabinkafley@ioepc.edu.np*

Submitted: April 13, 2025; Revised: June 15, 2025; Accepted: June 21, 2025

https://doi.org/10.3126/joeis.v4i1.81609

Abstract

Skin cancer is common and a rising cause of death globally. Early detection using deep learning helps save lives by enabling early and effective treatment. In this study, we utilized a pre-trained lightweight model, MobileOne, to classify malignant or benign skin cancer images. We combined two publicly available Kaggle datasets (ISIC 2019–2020 Malignant or Benign and Skin Cancer: Malignant vs. Benign). The proposed pre-trained model MobileOne outperformed existing models like MobileNetV3, Xception Net, Inception V3, VGG16, Densenet-121, ResNet50, VGG19, ViT b16, and ViT b32, achieving accuracy of 92.61%, precision of 0.9276, recall of 0.9261, F1-score of 0.9261, and ROC of 0.98. The results suggest that combined datasets with the lightweight MobileOne model improve accuracy and offer potential for fast, low-latency mobile skin cancer diagnosis, especially in resource-limited and remote settings.

Keywords —benign, fine-tuning, lightweight model, malignant, mobileone, pre-trained model, skin cancer datasets, transfer learning

1. Introduction

Skin has three layers: epidermis, dermis, and subcutaneous tissue. Skin cancer is caused mainly by UV radiation, genetics, and harmful chemicals. It starts in the epidermis with uncontrolled cell growth. Skin cancers are primarily categorized as benign (non-cancerous) or malignant (cancerous). Common benign skin lesions include seborrheic keratosis, dermatofibroma, and benign nevi (moles). Malignant skin cancers primarily consist of basal cell carcinoma, squamous cell carcinoma, and melanoma, which are more aggressive and require prompt treatment. The GLOBOCAN 2022 report states that melanoma is the 17th most common

cancer globally, with about 331,722 cases and 58,667 deaths, while non-melanoma skin cancer ranks 5th, with 1,234,533 cases and 69,416 deaths; incidence is highest in Oceania, North America, and Europe, especially Australia, while Asia and Africa lowest, but mortality rates vary little by region, and non-melanoma deaths currently exceed melanoma deaths, with Human Development Index (HDI) positively correlating with melanoma incidence and mortality, but negatively with non-melanoma mortality (Wang, M., Gao, X., & Zhang, L., 2025).

Clinical skin cancer detection is time-consuming, often invasive due to biopsies, and relies on specialist availability. It's less accessible in remote areas and hard to scale for mass screening due to manual inspection needs. Artificial intelligence, particularly deep learning, significantly enhances skin cancer detection when used prospectively in primary care patients, which could add significant clinical value for primary care physicians assessing skin lesions for melanoma (Papachristou, P., et al., 2024).

Traditional statistical and machine learning methods using handcrafted features frequently show limited generalization and struggle to capture complex patterns in skin lesions. MobileOne's highly efficient design features millisecond inference times on mobile devices through reparameterization Vasu, P. (K. A., Gabriel, J., Zhu, J., Tuzel, O., & Ranjan, A. ,2023). Therefore, the implementation of an AI-based mHealth app for the detection of skin cancer in the hands of patients or as a diagnostic tool used by GPs in primary care appears feasible (Smak Gregoor, A. M., Sangers, T. E., Eekhof, J. A. H., Howe, S., Revelman, J., Litjens, R. J. M., et al.,2023). In this study, we combined two Kaggle datasets—ISIC 2019-2020 Malignant or Benign and Skin Cancer: Malignant vs. Benign—to enhance model performance and fine-tuned the pre-trained MobileOne (S4 variant) for skin cancer classification. Input images were processed using two approaches: SGD without median filtering with a StepLR scheduler, and Adam with median filtering using a ReduceLROnPlateau scheduler. Both approaches applied data augmentation, froze backbone layers, and replaced the top classification layer for two-class classification. After training, classification and evaluation were performed, and the results were compared with existing methods. We achieved an impressive test accuracy of 92.61%, precision of 92.76%, recall, and F1-Score of 92.61% and ROC of 0.98.

Our paper is arranged as follows. Section 2 provides the background and literature review needed for the study. The training data sources, preprocessing methods, and fine-tuning the pre-trained model for skin cancer classification are discussed in Section 3. The experiments are outlined, and the results are analyzed and compared in Sections 4, 5, and 6. Finally, we summarize this study in Section 7.

2. Background Study and Literature Review

Recent advances in skin cancer classification have increasingly leveraged lightweight convolutional neural networks (CNNs) and fine-tuned transfer learning approaches to achieve efficient and accurate classifications of dermoscopic images.

(F. Mahmud, M. M. Mahfiz, M. Z. I. Kabir, and Y. Abdullah, 2023) used pre-trained models XceptionNet, EfficientNetV2S, InceptionResNetV2, EfficientNetV2M, and modern explainable artificial intelligence (XAI) to address the problem of skin cancer detection (Mahmud, F., Mahfiz, M. M., Kabir, M. Z. I., & Abdullah, Y.,2023, December). They used HAM1000 datasets and resized the image to 224x224, split the dataset as 80:10:10 training, test, and validation sets, and augmented the dataset. Their models detected 7 types of skin cancer, consisting of both malignant and benign classes: AKIEC, BCC, BKL, DF,

MEL, NV, and VASC. They got 88.72% using XceptionNet and visualized the classified and misclassified cancer images using XAI techniques, SmoothGrad, and Faster Score-CAM.

(A. M. Ibrahim, M. Elbasheir, S. Badawi, A. Mohammed, and A. F. M. Alalmin, 2023) used the Visual Geometry Group 16 (VGG16) architecture on a Kaggle dataset, applying data cleaning processes like noise and hair removal, contrast enhancement, and resizing the input images. They also applied segmentation (edge-based, region-based, morphological, active-contours, histogram-based) and feature extraction techniques (asymmetry, diameters, compactness, borders, blue & white veil). They were able to achieve 84.242% accuracy by applying several preprocessing techniques.

(Alrabai, A., Echtioui, A., and Kallel, F., 2024) used InceptionV3 and Xception Kaggle: Skin Cancer: Malignant vs. Benign. They applied resizing, normalization, data augmentation, and fine-tuned a model using maxpooling, batch normalization, dropout 0.5, and final layers with sigmoid classification. InceptionV3 and Xception achieved 89.11% and 87.03% accuracy.

(Bello, A., Ng, S.-C., and Leung, M.-F. ,2024) used the ISIC 2020 dataset with 1497 malignant images and 1800 benign images to classify malignant or benign. They used data augmentation, an image generator, and normalization to multiple methods like CNN, SVM, Random Forest Classifier, fine-tuned deep learning models such as EfficientNetB0, ResNet34, VGG16, InceptionV3, and DenseNet121. Among all the used methods, they were only able to achieve the highest accuracy of 87.00% on fine-tuned DenseNet-121.

(Alrabai, A., Echtioui, A., and Kallel, F., 2025) used ResNet50, VGG19, ViT b16, and ViT b32 with the kaggle dataset, training 1497 malignant and 1800 benign images, and testing 360 benign and 300 malignant images. Among all the proposed models, they achieved an accuracy of 89.09% on ResNet50 using XAI techniques LIME and Integrated Gradients.

Mehboob, S., Bukhari, M., Shah, Y. A., Khan, S., Sharif, M., and Ijist Jr. (2025) used MobileNetV3 Small and MobileNetV3 Large for classifying the malignant or benign image using the ISIC Archive. They used 3297 images to train their models, of which 1497 images belong to the malignant class, and 1800 images belong to the benign class. Despite implementing regularization techniques such as dropout, the model may still overfit to specific patterns in the training data, hindering its ability to generalize effectively to unseen examples (Mehboob, S., Bukhari, M., Shah, Y. A., Khan, S., Sharif, M., & Ijist Jr. ,2025). . But they got MobileNetV3 Small and MobileNetV3 Large with an accuracy of 87.00% and 89.00%.

MobileOne is a lightweight CNN backbone by Apple, designed for mobile devices with ultra-fast inference (under 1 ms on iPhone 12) and high accuracy (up to 79.4% Top-1

on ImageNet) Vasu, P. (K. A., Gabriel, J., Zhu, J., Tuzel, O., & Ranjan, A. ,2023).. It features a two-stage design: a complex training phase with multi-branch convolutions (3×3 depth-wise, 1×1 pointwise), Batch Normalization, optional Squeeze-and-Excite blocks, and trivial branches; followed by a re-parameterized, single-branch inference phase. This simplified structure without branches or skip connections and using only ReLU minimizes memory use and ensures efficient deployment.

Table 1: Existing Analysis of Related Work

Study	Model	Dataset	Accuracy
F. Mahmud, M. M. Mahfiz, M. Z. I. Kabir and Y. Abdullah (2023)	XceptionNet	HAM1000	88.72%
A. M. Ibrahim, M. Elbasheir, S. Badawi, A. Mohammed, and A. F. M. Alalmin (2023)	VGG-16	Kaggle	84.242%
Bello, A., Ng, S.-C., & Leung, M.-F. (2024)	Densenet-121	ISIC-2020	89.09%
A. Alrabai, A. Echtioui, and F. Kallel (2024)	InceptionV3	Kaggle	87.00%
A. Alrabai, A. Echtioui, and F. Kallel (2024)	MobileNetV3 Small, MobileNetV3 Large	ISIC Archive	87.00% 89.00%
Mehboob, S., Bukhari, M., Shah, Y. A., Khan, S., Sharif, M. and Ijist Jr. (2025)	ResNet50, VGG19, ViT b16, ViT b32	Kaggle	89.09% 86.21% 86.97% 83.33%

Collectively, these studies reflect a growing trend toward balancing accuracy and efficiency in skin cancer classification models. We proposed and utilized a lightweight pre-trained CNN MobileOne model and achieved accuracy, precision, recall, and F1-score, ROC of 92.61%, 92.76%, 92.61%, and 0.98. To the best of our knowledge, our transfer learning

approach, MobileOne S4 has not been explored in the related fields.

3. Methodology

This section describes the proposed methodology for classifying the malignant and benign skin cancer images. Data collection, preprocessing, augmentation, fine-tuning, classification process, and evaluation metrics are all part of the methodology. The overall process for skin cancer classification is illustrated in Fig. 1.

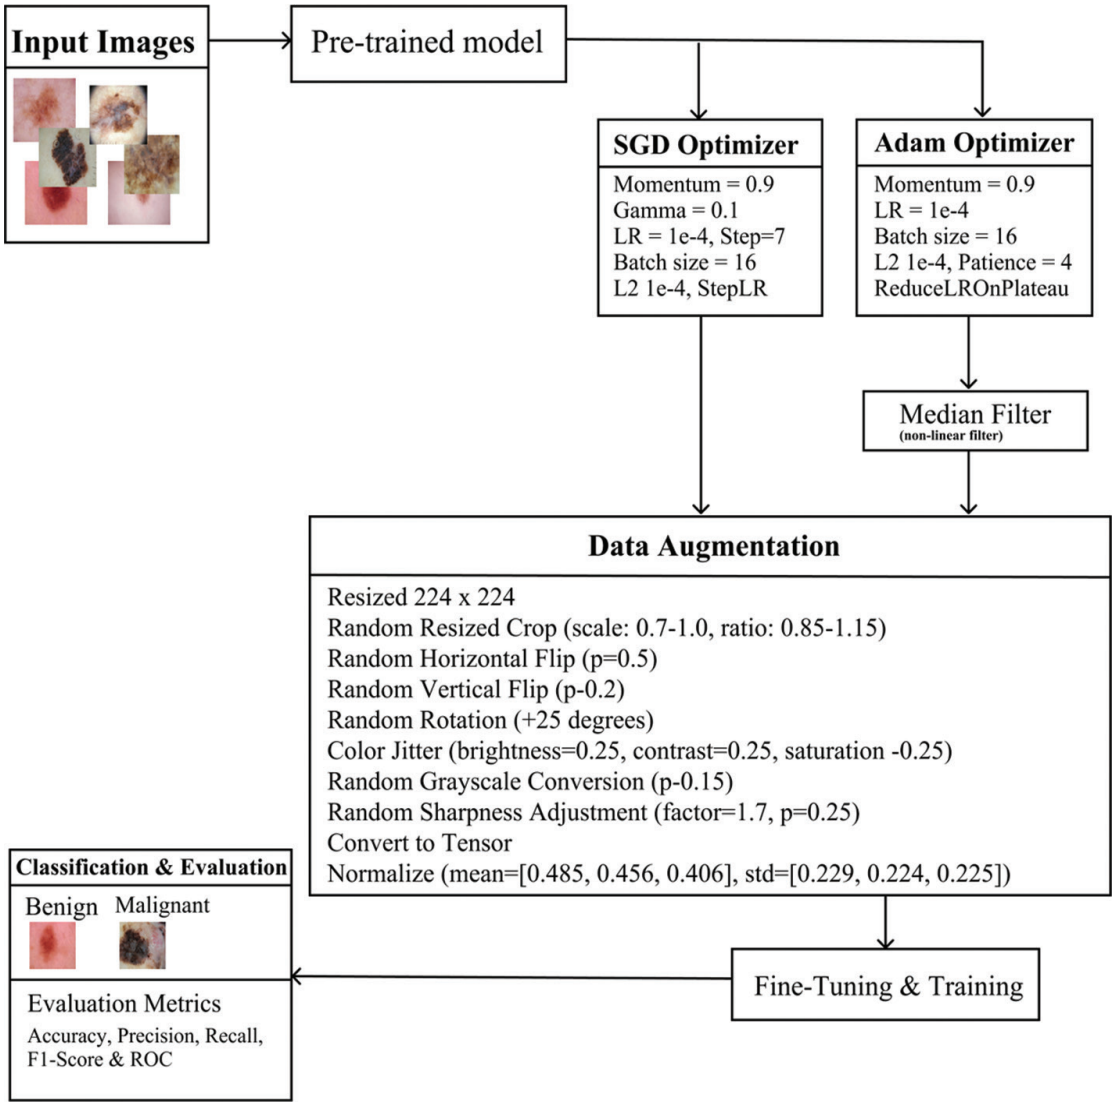

Figure 1: Fine-tuning pipeline using the MobileOne pre-trained model. A median filter is applied before data augmentation only when using the Adam optimizer. Data augmentation is applied for both Adam and SGD optimizers.

3.1 Dataset

In this study, we combined the publicly available Skin Cancer: Malignant vs. Benign and ISIC 2019-2020 Malignant or Benign datasets from Kaggle, which contain labeled images of skin lesions categorized as either malignant (cancerous) or benign (non-cancerous). A total of 13,186 images were split in a 70:15:15 ratio for training, validation, and testing. To achieve equal class distribution across training, validation, and test splits, the dataset was manually balanced using undersampling techniques. The training set contains 4617 malignant and 4617 benign images. Both the test and validation sets contain 988 malignant and 988 benign images each.

3.2 Data preprocessing

All Input images are resized to 224×224 and normalized. The input image is normalized channel-wise using the following formula: The normalization of each pixel x_c in channel c ∈ {R, G, B}is performed as

$$x'_c = \frac{x_c - \mu_c}{\sigma_c} \quad (1)$$

where μ = [0.485,0.456,0.406] and σ = [0.229,0.224,0.225] are the mean and standard deviation values.

3.3 Median Filter (Non-Linear Filtering)

Median filters are quite popular because, for certain types of random noise, they provide

excellent noise-reduction capabilities, with considerably less blurring than linear smoothing filters of similar size (Gonzalez, R. C., & Woods, R. E.,2018). Given an input image I and a window size k×k, for each pixel (i,j), define the window Wi,j as the set of pixel values:

$$w_{i,j} = \{ I(m,n) \mid m \in [i-r, i+r], n \in [j-r, j+r] \} \quad (2)$$

where $r = \lfloor \frac{k}{2} \rfloor$ is the half window size.

The set $w_{i,j}$ contains all pixel values in the neighborhood centered at (i,j). Sorting $w_{i,j}$ and

taking the median gives the filtered output pixel value:

$$I_{filtered}(i,j) = \text{median}(w_{i,j}) \quad (3)$$

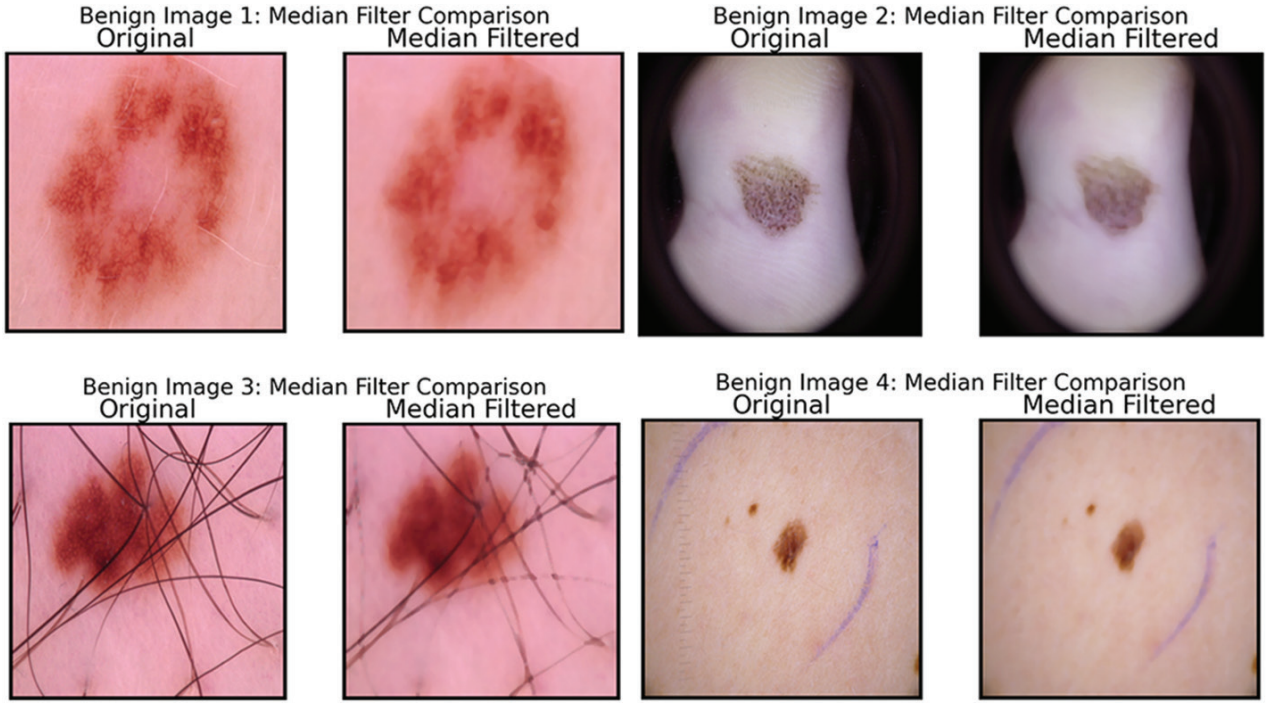

Figure 2: Visual comparison between the original benign and median-filtered benign image

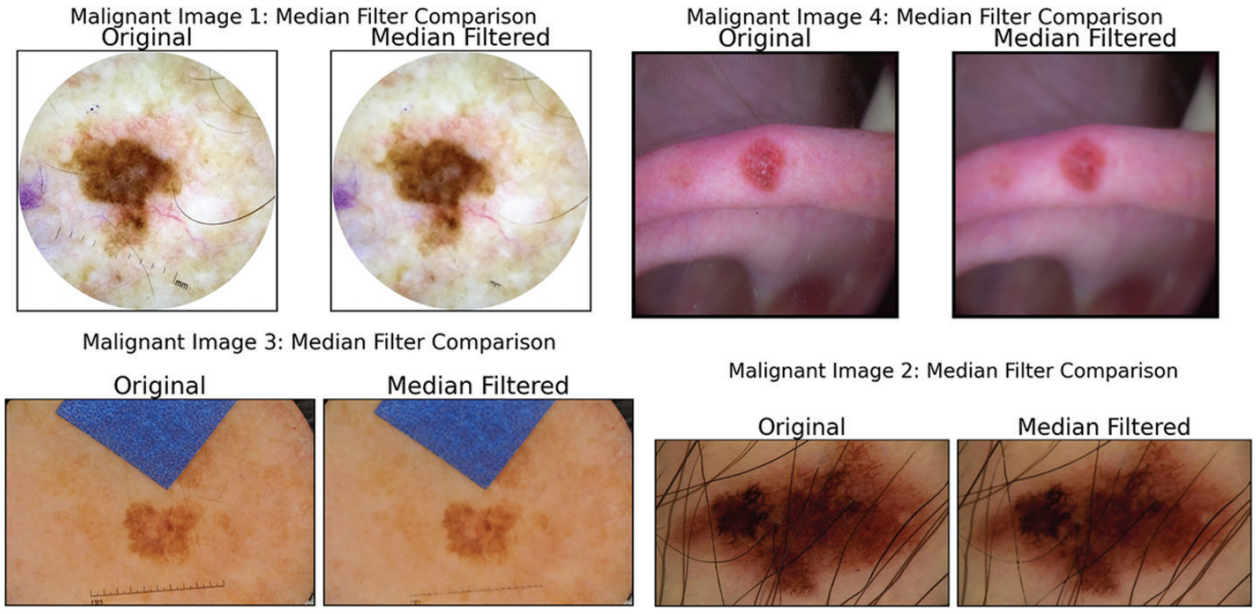

Figure 3: Visual comparison between the original malignant and median-filtered malignant image

3.4 Data augmentation

Data augmentation is applied to enhance model generalization and reduce overfitting. These included resizing, random cropping, horizontal and vertical flipping, random rotation, random color jittering, random grayscale conversion, random sharpness adjustment, and normalization. Validation and test images are only resized and normalized without augmentation.

3.5 Classification and Fine-Tuning

The Original Pre-trained model is trained on ImageNet with 1000 classes. We freeze all the except the top layer for the classification. The top classification layer is removed and replaced with a new classification layer for classes 2. The model is fine-tuned and able to make classifications accurately using the loss function CrossEntropy.

3.5.1 Original Pre-trained Model:

Let $\theta=\{\theta_1,\theta_2,\ldots,\theta_L\}$ denote the parameters of the network, where L is the total number of layers. Given an input image x of size height × width × channels, the model outputs logits:

$$z = f(x;\theta) = [z_1, z_2, \ldots, z_{1000}] \quad (4)$$

The predicted class probabilities are obtained by applying the softmax function:

$$\hat{y}_i = \frac{e^{z_j i}}{\sum_{j=1}^{1000} e^{z_j}},\ i = 1,2,\ldots.1000 \quad (5)$$

where $\hat{\mathbf{y}} = [\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \ldots\ldots \hat{\mathbf{y}}_{1000}]$ is a probability distribution satisfying $\sum_i \quad \hat{y}_i = 1.$

The true label vector $\mathbf{y} = [\mathbf{y}_{1,\ y2,} \ldots\ldots \mathbf{y}_{1000}]$ is one-hot encoded such that $\mathbf{y}_i \in \{0,1\}$ and .

The cross-entropy loss function is defined as

$$L(\theta\ ) = -\sum_{i=1}^{1000} y_i \log(\hat{y}) \quad (6)$$

3.5.2) Modified Classification Head:

For the skin cancer classification task, the original 1000-class output layer is replaced with a two-class output layer:

$$z' = f'(x;\theta') = [z_1', z_2'] \ (7)$$

The SoftMax probabilities for the two classes are:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^{2} e^{z_j'}},\ i = 1,2 \ (8)$$

The loss function for fine-tuning is:

$$L(\theta') = -\sum_{i=1}^{2} y_i \log(\hat{y}) \quad (9)$$

The objective is to optimize

$$\theta'^{*} = \arg\min_{\theta'} \quad E(x,y) \sim D[L(\theta')] (10)$$

During fine-tuning, input samples are propagated through frozen backbone layers, producing fixed feature embeddings with non-trainable parameters. These embeddings feed into trainable classification layers, where task-specific learning occurs. Backpropagation updates only the trainable layers' weights, preserving backbone representations while optimizing final predictions.

3.6 Evaluation Metrics

To assess the proposed fine-tuned model's performance, standard metrics were employed. Accuracy measures the overall correctness of predictions. Precision indicates the proportion of correctly identified positive cases among all predicted positives. Recall reflects the model's ability to detect actual positive cases. The F1-score balances precision and recall into a single metric. The ROC curve visualizes the trade-off between true positive rate and false positive rate across different classification thresholds. Here, TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \ (11)$$

$$Recall = \frac{TP}{TP+FP} \ (12)$$

$$TPR = Recall = \frac{TP}{TP + FN}, FPR = \frac{FP}{FP + TN} (13)$$

$$\Pr e\,cision = \frac{TP}{TP + FP} (14)$$

$$F1 = 2 \times \frac{\Pr e\,cision \times Recall}{\Pr e\,cision + Recall} (15)$$

Table 2: Classification Matrix for Binary Classification

Actual/Predicted	Positive	Negative
Positive	True Positive (TP)	False Negative (FN)
Negative	False Positive (FP)	True Negative (TN)

4. Experiment

In our experiment, the pre-trained MobileOne model was fine-tuned using both SGD and Adam optimizers. Median filtering was applied only in the Adam configuration, which

led to notable improvements in accuracy, as reported in Table IV. The hyperparameters used in the fine-tuning process are summarized in Table III.

To maintain computational efficiency, training was conducted using a batch size of 16 over 15 epochs. Stochastic Gradient Descent (SGD) with a momentum factor of 0.9 was employed for its stable convergence characteristics and reduced oscillatory behavior near local minima. Adam optimizer—with default momentum coefficients—was employed, demonstrating faster initial convergence. But it showed variance in later training stages as shown in Fig. 7, due to its adaptive learning rate behavior and the dynamic adjustments from the ReduceLROnPlateau scheduler.

Table 3: Fine-tuned Hyperparameters for Pre-trained Model

Hyperparameters	**MobileOne (SGD)**	**MobileOne (Adam)**
Optimizer	SGD	Adam
Momentum	0.9	0.9
Gamma	0.1	-
Epochs	15	15
Filter	-	Median
Learning Rate	0.0001	0.0001
Batch Size	16	16
L2 Regularization	0.0001	0.0001
Scheduler	StepLR (Step Size = 7, Gamma = 0.1)	ReduceOnPlateau (mode = min, patience=4, min_lr = 0.0000001)
Loss Function	CrossEntropy	CrossEntropy

5. Results and Analysis

This section describes the results and shows comparative analysis. The pre-trained models MobileOne with SGD and MobileOne with Adam accurately predicted two types of skin cancer images. Table 4 shows the results from the tests of the pre-trained models.

Table 4: Classification Report in Weighted Average

Model	Optimizer	Accuracy	Precision	Recall	F1-Score	ROC
MobileOne	SGD	90.23%	0.9025	0.9023	0.9032	0.96
MobileOne	Adam	92.61%	0.9276	0.9261	0.9261	0.98

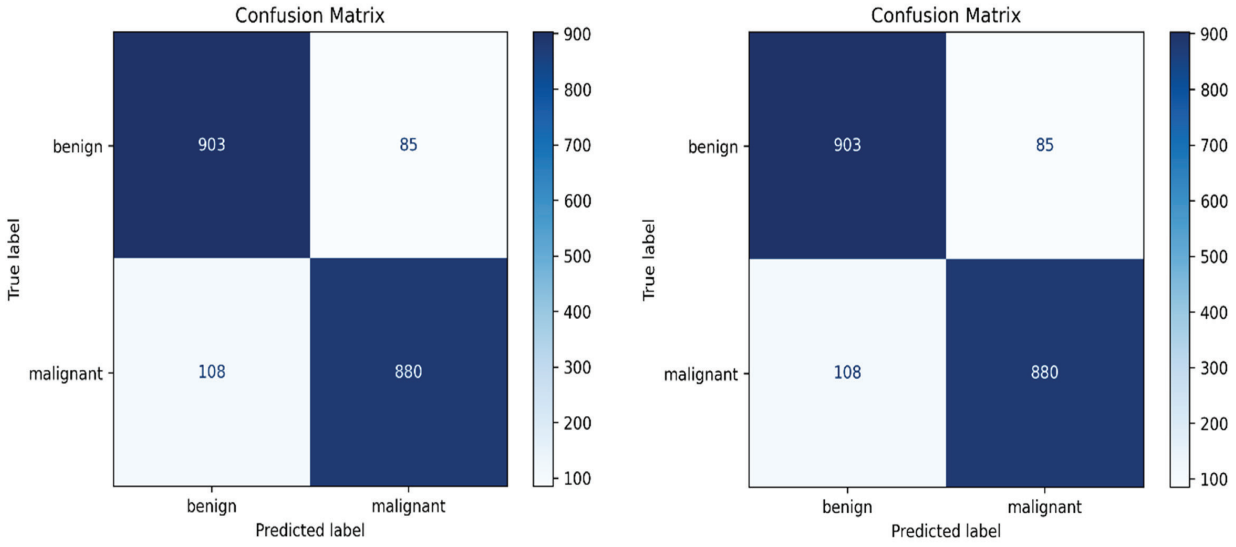

Figure 4: Confusion matrices for the MobileOne model fine-tuned with (a) the SGD optimizer and (b) Adam optimizer.

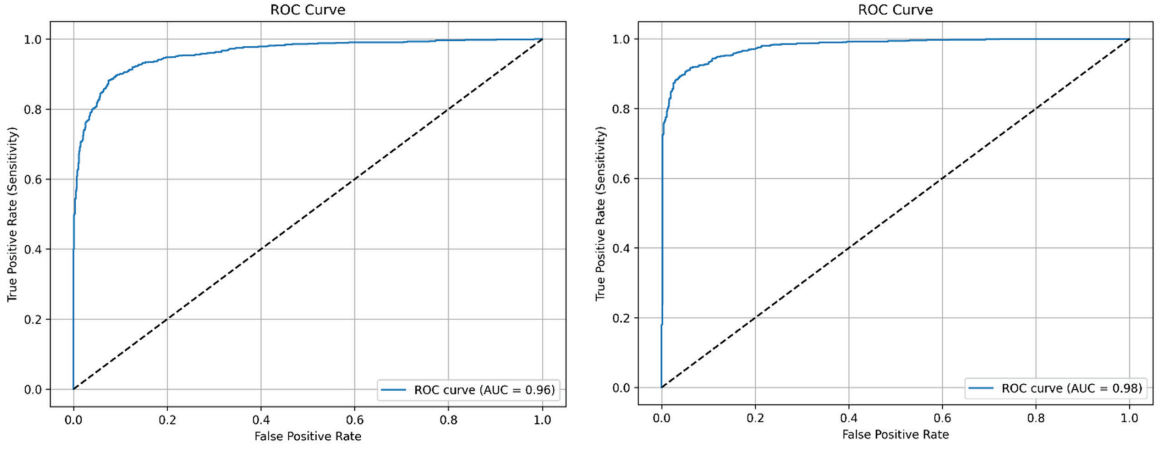

Figure 5: ROC Curve matrices for MobileOne model fine-tuned with (a) SGD optimizer and (b) Adam optimizer.

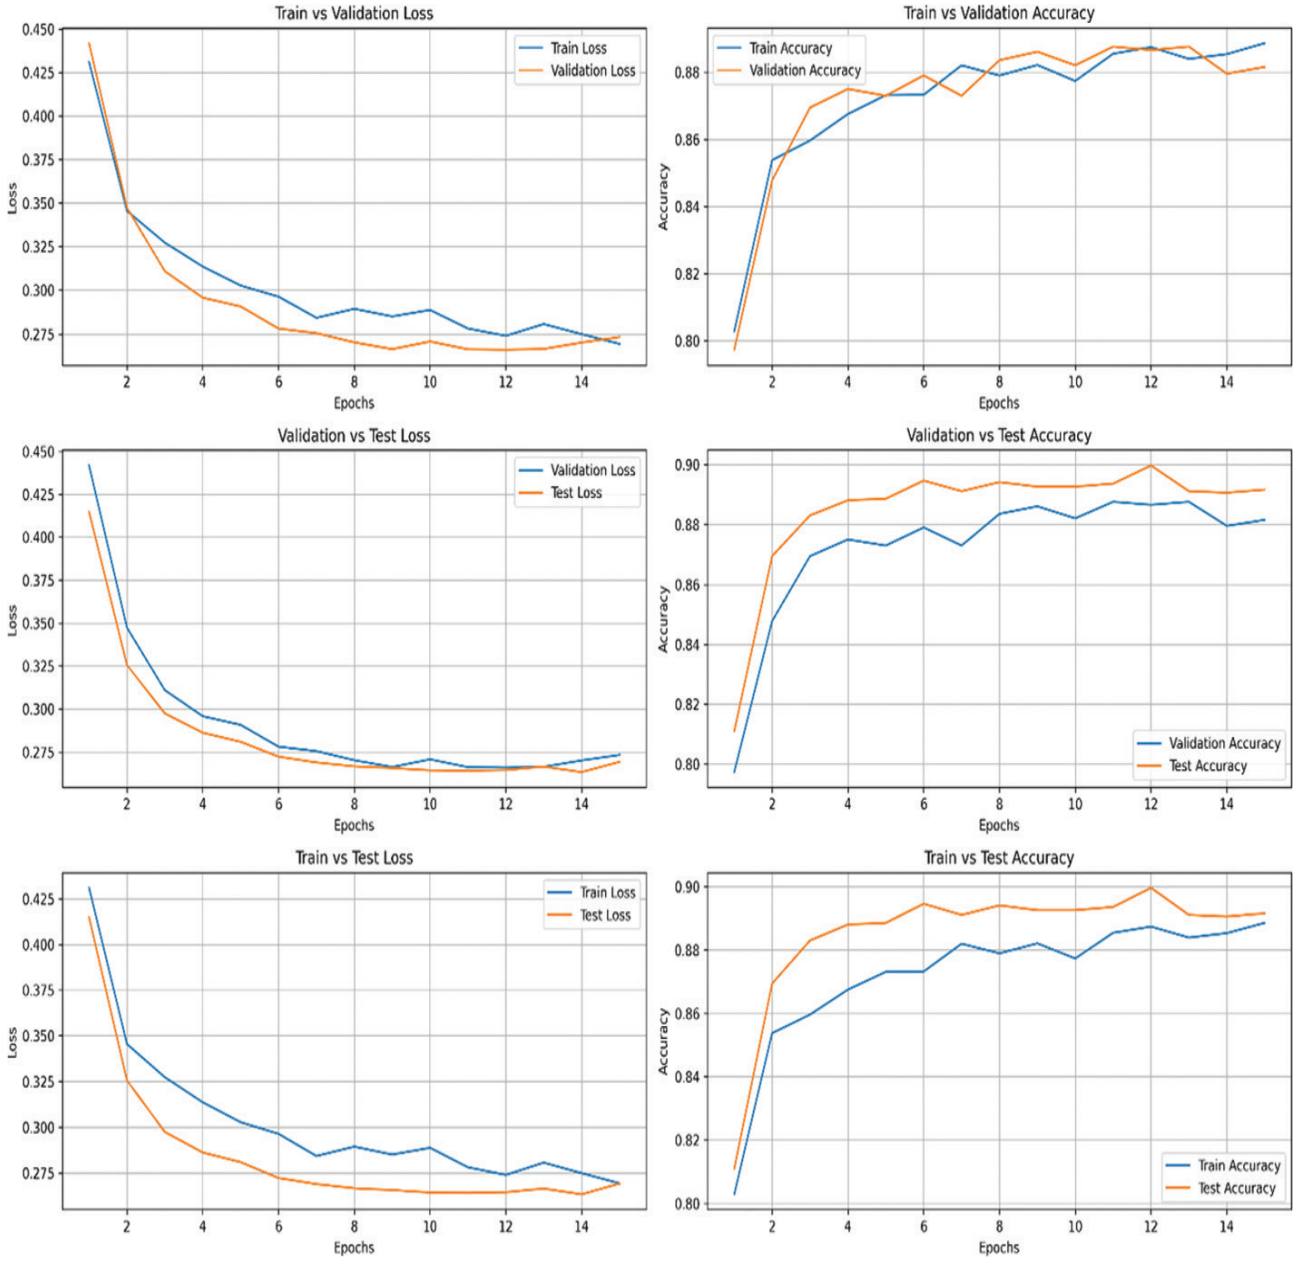

Figure 6: Training History of MobileOne with SGD Optimizer

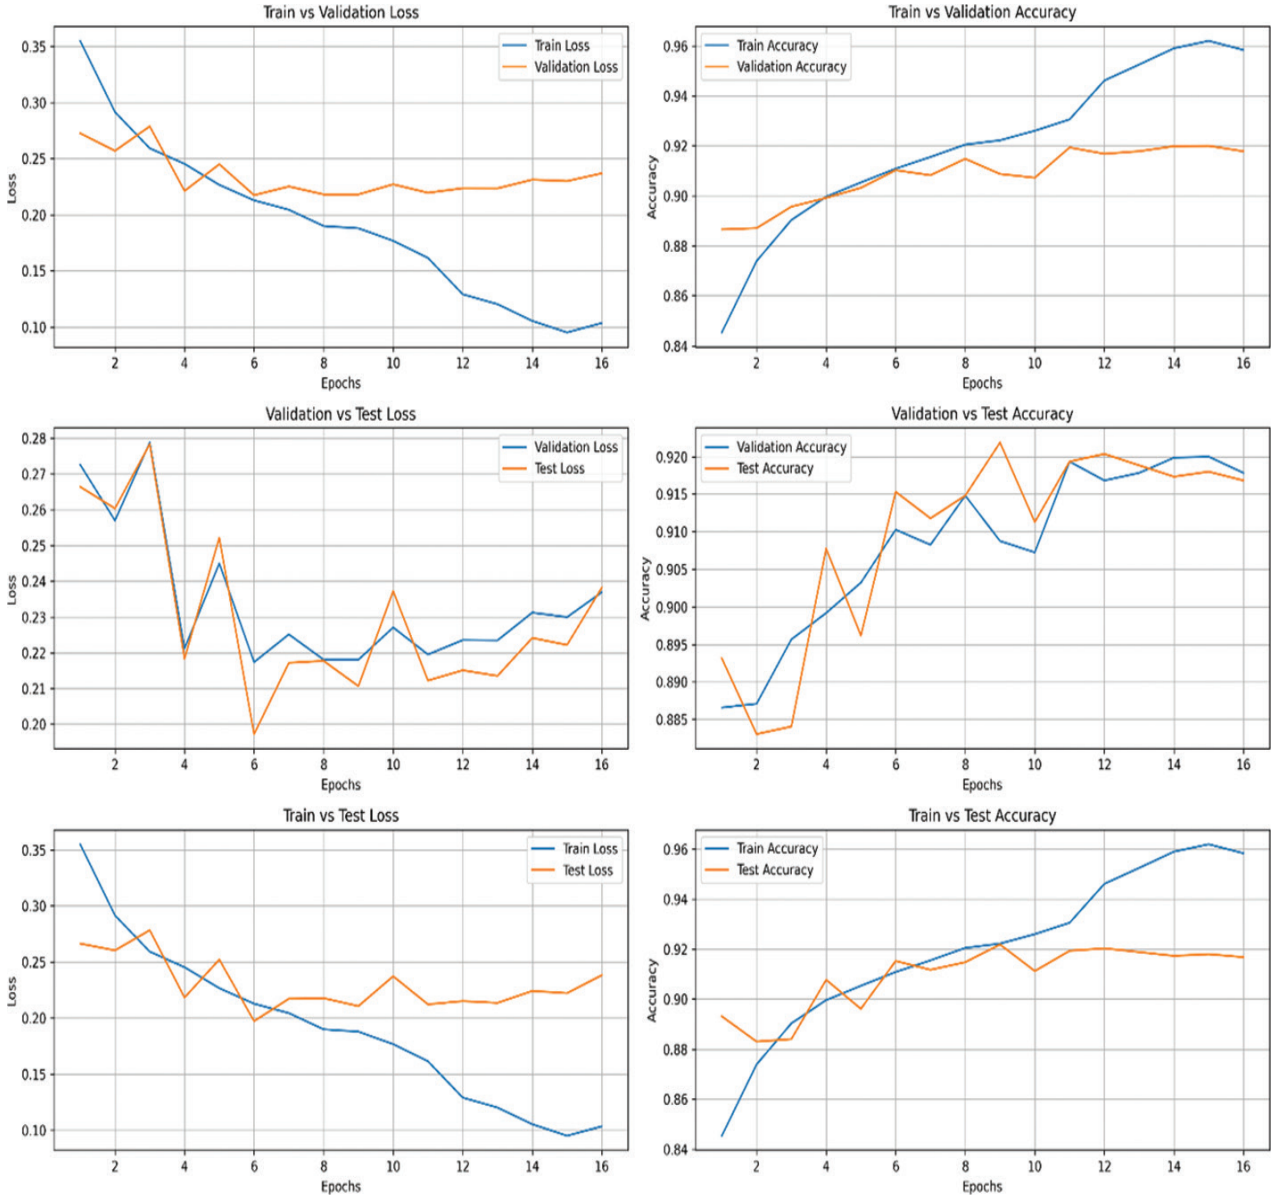

Figure 7: Training History of MobileOne with Adam Optimizer

6. Results Comparison

This section presents a comparative analysis of the proposed MobileOne model's performance against state-of-the-art models. We fine-tuned the model without relying on specialized preprocessing techniques, initializing a network with transferred features from almost any number of layers can produce a boost to generalization performance after fine-tuning to a new dataset (Yosinski, J., Clune, J., Bengio, Y., & Lipson, H.,2014). As summarized in Table V, MobileOne demonstrates superior accuracy with both SGD and Adam optimizers, outperforming prominent models MobileNetV3, DenseNet121, Resnet50, VGG19, XceptionNet, Inception V3, VGG16, and Vision Transformers.

Table 5: Comparison of Proposed Methods' Accuracy with Previous Methods

Paper	Models	Accuracy
Mehboob, S., Bukhari, M., Shah, Y. A., Khan, S., Sharif, M., & Ijist Jr. (2025).	MobileNetV3 Small MobileNetV3 Large	87.00% 89.00%
Bello, A., Ng, S.-C., & Leung, M.-F. (2024).	DenseNet-121	87.00%
Ibrahim, A. M., Elbasheir, M., Badawi, S., Mohammed, A., & Alalmin, A. F. M. (2023).	Inception V3	89.09%
Wang, M., Gao, X., & Zhang, L. (2025).	ResNet50 VGG19 ViT b16 ViT b32	89.09% 86.21% 86.97% 83.33%
Mahmud, F., Mahfiz, M. M., Kabir, M. Z. I., & Abdullah, Y. (2023, December)	XceptionNet	88.72%
Djaroudib, K., Lorenz, P., Bouzida, R. B., & Merzougui, H. (2024).	VGG-16	84.5%
Proposed Method	**MobileOne(SGD) MobileOne (Adam)**	**90.23% 92.61%**

7. Conclusions

Skin cancer remains one of the most rapidly increasing cancers globally, with significant mortality rates. To address this healthcare issue, our study introduces a pre-trained model, MobileOne—a lightweight and efficient convolutional model designed for mobile devices, offering high accuracy and low latency in milliseconds. The proposed method applies manual undersampling to balance the dataset, median filtering for noise removal, and data augmentation to counter the over-fitting and enhance generalization. The model was trained and evaluated on a merged dataset from public sources and assessed using standard evaluation metrics. The MobileOne model demonstrated promising results with 92.61% accuracy, surpassing state-of-the-art lightweight CNN models in accuracy, indicating its potential as a reliable classification tool.

Our study demonstrates the potential for early detection of skin cancer, mobile and remote screening, reduction of the burden on healthcare systems, and encourages further research. This study also faced several challenges, including dataset imbalance, limited diversity in skin tones, variability in test accuracy with dataset size, and adaptability of pre-trained models, data noise, and visual artifacts commonly found in clinical skin images.

However, different problems require different approaches when applying deep learning models (Akter, M.S., Shahriar, H., Sneha, S., & Cuzzocrea, A., 2022). In future work, we plan to integrate Explainable AI (XAI) techniques and enhance the current model by developing hybrid architectures and incorporating additional lightweight models. Ultimately, our goal is to improve both accuracy and generalization across datasets with diverse skin tones.

Data Availability

This study utilizes publicly available datasets from Kaggle:

1. https://www.kaggle.com/datasets/fanconic/skin-cancer-malignant-vs-benign/data
2. https://www.kaggle.com/datasets/sallyibrahim/skin-cancer-isic-2019-2020-malignant-or-benign

References

Vasu, P. K. A., Gabriel, J., Zhu, J., Tuzel, O., & Ranjan, A. (2023). *MobileOne: An improved one millisecond mobile backbone.* In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).* https://doi.org/10.48550/arXiv.2206.04040

Mehboob, S., Bukhari, M., Shah, Y. A., Khan, S., Sharif, M., & Ijist Jr. (2025). Enhanced skin cancer classification with MobileNetV3 and morphological preprocessing: A deep learning-based extension. *International Journal of Innovative Science and Technology, 7,* 1–12.

Bello, A., Ng, S.-C., & Leung, M.-F. (2024). Skin cancer classification using fine-tuned transfer learning of DENSENET-121. *Applied Sciences, 14*(17), 7707. https://doi.org/10.3390/app14177707

Ibrahim, A. M., Elbasheir, M., Badawi, S., Mohammed, A., & Alalmin, A. F. M. (2023). Skin cancer classification using transfer learning by VGG16 architecture (case study on Kaggle dataset). *Journal of Intelligent Learning Systems and Applications, 15*(3), 61–75. https://doi.org/10.4236/jilsa.2023.153005

Wang, M., Gao, X., & Zhang, L. (2025). Recent global patterns in skin cancer incidence, mortality, and prevalence. *Chinese Medical Journal, 138*(2), 185–192. https://doi.org/10.1097/CM9.0000000000003416

Alrabai, A., Echtioui, A., & Kallel, F. (2025). Exploring pre-trained models for skin cancer classification. *Applied System Innovation, 8*(2), 35. https://doi.org/10.3390/asi8oz20035

Mahmud, F., Mahfiz, M. M., Kabir, M. Z. I., & Abdullah, Y. (2023, December). An interpretable deep learning approach for skin cancer categorization. In *2023 26th International Conference on Computer and Information Technology (ICCIT).* https://doi.org/10.1109/ICCIT60459.2023.10508527

Akter, M. S., Shahriar, H., Sneha, S., & Cuzzocrea, A. (2022). Multi-class skin cancer classification architecture based on deep convolutional neural network. In *2022 IEEE International Conference on Big Data (Big Data)*(pp. 5404–5413). https://doi.org/10.1109/BigData55660.2022.10020302

Smak Gregoor, A. M., Sangers, T. E., Eekhof, J. A. H., Howe, S., Revelman, J., Litjens, R. J. M., et al. (2023). Artificial intelligence in mobile health for skin cancer diagnostics at home (AIM HIGH): A pilot feasibility study. *eClinicalMedicine, 60,* 101020. https://doi.org/10.1016/j.eclinm.2023.101020

Papachristou, P., et al. (2024). Evaluation of an artificial intelligence-based decision support for the detection of cutaneous melanoma in primary care: A prospective real-life clinical trial. *British Journal of Dermatology, 191*(1), 125–133. https://doi.org/10.1093/bjd/ljae021

Djaroudib, K., Lorenz, P., Bouzida, R. B., & Merzougui, H. (2024). Skin cancer diagnosis using VGG16 and transfer learning: Analyzing the effects of data quality over quantity on model efficiency. *Applied Sciences, 14*(17), 7447. https://doi.org/10.3390/app14177447

Gonzalez, R. C., & Woods, R. E. (2018). *Digital image processing* (4th ed.). Pearson.

Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? In *Proceedings of the 28th International Conference on Neural Information Processing Systems (NeurIPS)* (Vol. 2, pp. 3320–3328).