

A Lightweight Hybrid MobileNetV3-Random Forest Model for Diabetic Foot Ulcer Detection for Resource – Constrained Nigerian Clinics

A. O. Lawal^{1*}, A. N. Babatunde², R. S. Babatunde³, I. A. Lawal⁴, A. F. Kadri⁵,
S. B. Muhammed⁶. & O. Ekundayo⁷.

^{1,2,3,5,6,7}Department of Computer Science, Faculty of Communication and Information Technology, Kwara State University, Malete, Kwara State, Nigeria

⁴Department of Human Anatomy, Faculty of Basic Medical Science, College of Health Science, Al-Hikmah University, Ilorin, Kwara State, Nigeria

*Corresponding Author: lawalyinkah@outlook.com

ABSTRACT: In Nigeria, Diabetic Foot Ulcer (DFU) is a severe consequence of diabetes that causes high rates of lower-leg amputations and death. Many primary care clinics throughout the nation depend on subjective, non-quantitative clinical judgments, which frequently lead in delayed diagnosis. The lightweight hybrid MobileNetV3-Random Forest model created in this study for automated DFU detection from smartphone-captured pictures, optimized for inexpensive devices found in resource-limited Nigerian clinics, confronts this problem. From several Nigerian medical sources, 1,721 anonymized DFU photographs including a variety of skin tones and ulcer grades were compiled. To increase robustness against lighting and orientation variances, images were normalized, resized to 224x224 pixels, and augmented. Reducing the image dimensionality from 28,224 to 96 principal components while retaining 95% of the variance, PCA was then utilized. The feature extractor, MobileNetV3 (pre-trained through transfer learning), generated a high-dimensional feature vector. For the ultimate DFU detection job, these condensed elements were input into a Random Forest classifier. With five-fold cross-validation employed to guarantee robustness, the dataset was split into 70/20/10 for training, testing and cross-validation respectively. Achieving 92.8% accuracy, 0.93 precision, 0.93 recall, 0.93 F1-score, and 0.9959 ROC-AUC, the hybrid model outperformed several heavier benchmarks in efficiency. Ablation studies validated that PCA decreased memory use by 31% on low-end devices and reduced inference time by about 47% (from 352 ms to 187 ms per image), with very little influence on predictive performance. This study offers a scalable, computationally economical, and understandable artificial intelligence framework supporting hybrid intelligence ideas that advances fair early diagnosis of diabetic foot ulcers in resource-poor medical environments.

KEYWORD: Diabetic Foot Ulcer (DFU), Artificial Intelligence, Principal Component Analysis, Convolutional Neural Network, MobileNetV3, Random Forest Algorithm.

I. INTRODUCTION

One of the most serious consequences of diabetes, diabetic foot ulcer (DFU) is defined as a complete skin break that extends below the ankle in people with diabetes mainly caused by neuropathy, peripheral artery disease, and excessive mechanical pressure or trauma (Olopade et al., 2025). With lifetime risk of developing a DFU estimated at up to 25% and rates of lower-limb amputations (as high as 35% in some regions) and mortality rates often exceeding those of many

malignancies (Olopade et al., 2025; Olamoyegun et al., 2026). Nigeria, the country with the greatest diabetes prevalence in Sub-Saharan Africa is Regular diagnosis depends on subjective clinical assessments by healthcare providers, lab tests, and traditional imaging like X-rays, which are sometimes conflicting, experience-dependent, and insufficient for early detection, particularly in resource-constrained areas (Olopade et al., 2025; Misir, 2025; Olamoyegun et al., 2025). Although more objective, sophisticated imaging methods such as thermal or hyperspectral imaging are expensive and impossible to use broadly in Nigerian primary health care clinics (Sari et al., 2025). Most existing models are too computationally intensive for low-cost, off-the-shelf mobile devices usually found in Nigerian hospitals, even if approaches based on Convolutional Neural Networks (CNNs) have shown potential for automated DFU detection from smartphone photos across (Lee et al., 2025). Using smartphone-captured images, this study suggests a lightweight hybrid MobileNetV3-Random Forest model to bridge this gap particularly optimized for resource-constrained Nigerian clinical situations for quick and accurate identification of diabetic foot ulcers. It is against these aforementioned that the research proposes to address these challenges by developing a MobileNetV3 – Random Forest model for diabetic foot ulcer detection, for Nigerian clinics. By leveraging a diverse Nigerian dataset, the study aims to achieve an effective model, enhancing early detection and reducing amputation rates in resource-constrained settings.

II. RELATED WORKS

This subsection reviews studies that highlight the severe clinical, socioeconomic, and molecular burden of diabetic foot ulcers (DFUs) both in Nigeria and worldwide, and critiques traditional methods of management.

A. DFU Application Using Deep Learning Approach

Research studies have assessed architectures such as MobileNet variants, EfficientNet-Lite, and hybrid CNN-Transformer models, which have shown high accuracy in classification, segmentation, and localization tasks when using images from smartphones or thermographic cameras (Lin et al., 2025; Karthik et al., 2025; Albahli et al., 2024). Research emphasizes compression methods (pruning, quantization) to facilitate real-time processing on edge devices, with models reaching 90-98% accuracy on standardized test datasets. Hybrid methods that combine CNN feature extraction with traditional classifiers or Vision Transformers enhance robustness and generalizability (Zhou et al., 2024). Comments from reviews highlight issues such as dataset bias (underrepresentation of diverse skin tones), class imbalance, and the requirement for explainable AI to establish clinical trust (Alkhalefah et al., 2024). Generative AI applications span synthetic data, real-time systems accelerated by FPGA, and multi-modal platforms, as seen in the work of Fadhel et al. (2024) and Sendilraj et al. (2024), with potential for early intervention in low-resource areas.

B. DFU Application using Hybrid Architecture

Further research on hybrid models, such as combining CNN features with Random Forest or boosting ensembles, shows better accuracy and efficiency, especially for intricate image classification tasks (Zhou et al., 2024). Studies investigate GAN-augmented ResNet (El-Kady et al., 2024), combinations of object detection systems (specifically YOLO variants) (Sarmun et al., 2024), and transparent frameworks (XAI-FusionNet with SHAP/LIME integration) (Biswas et al., 2024) to boost interpretability and performance when working with restricted datasets. The deployment potential of mobile platforms and hardware-accelerated systems is further illustrated by Fadhel et al. (2024) and Sendilraj et al. (2024).

C. Comparative Performance Summary

Table 1 presents a comparative analysis of the related works considered with a view to showcasing the limitations in existing studies.

Table 1: Summary of Related Work

Author	Approach	Dataset	Accuracy	Limitation
Girmaw & Taye, 2025	Lightweight MobileNetV2 for detection and grading	DFU Images	High but not specified	Though the model is mobile friendly, it lacks dataset diversity.
Sermun et al., 2024	YOLO based ensemble for localization	DFUC 20	86.4%	Though the model has strong localization, it is computationally demanding.
Biswas et al., 2024	Multi-scale XAI – FusionNet with explainability	DFUC 2020/2021	96.9%	The model is heavy and unsuitable for resource-constraint settings.
Fadhel et al., 2024	Real-time deep learning with hardware.	Clinical Images	97.8%	The model required a specialized hardware.

III. METHODOLOGIES

This section discusses the methodological approach for developing a MobileNetV3 Random Forest model for diabetic foot ulcer detection for Nigerian clinics. The methodology is designed to address Nigeria’s unique healthcare challenges, including limited internet access, hardware shortages, high computational demand and cultural distrust of AI. The chapter is structured to provide a clear, systematic roadmap for achieving the research objectives: compiling a diverse Nigerian diabetic foot ulcer dataset, preprocessing data, developing and optimizing a hybrid AI model, evaluating its performance, and deploying it on low-cost smartphones. This methodology ensures the model is accurate, computationally less demanding and enhancing early diabetic foot ulcer detection to reduce Nigeria’s recourse constrained environments.

Figure 3.1 depict the process flow diagram which visualizes the methodology, depicting the sequential processes and their interconnections. It begins with compiling a dataset of 1,721 DFU images from Nigerian medical sources, followed by preprocessing to standardize images, clinical data and dimensionality reduction using principal component analysis (PCA) after which a hybrid MobileNetV3-Random Forest model is developed, where MobileNetV3 is used for feature extraction while Random Forest is used for DFU detection. Performance is evaluated using metrics like accuracy and recall. This flow ensures computational efficiency and accessibility in low-resource settings.

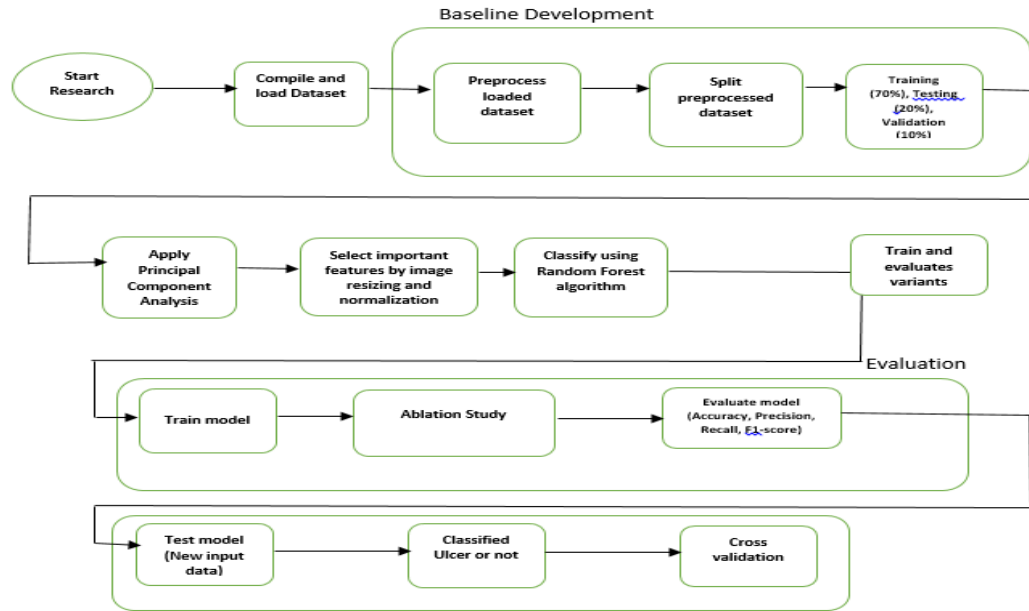


Figure 1: Process Flow Diagram of the Research

A. Data Compilation

The dataset for this study consists of DFU images sourced from different hospitals from 4 different states in Northern States of Nigeria (Yobe, Abuja, Taraba, and Adamawa) which contain dataset from diverse and representative data. These images are categorized based on ulcer severity and other relevant clinical features.

Table 2: An Overview of the Dataset Characteristics

PROPERTIES	CATEGORIES
Pixel – Based	Intensity, Color and Texture
Shape – Based	Area, perimeter, convexity, circularity and eccentricity
Object – Based	Size, shape, orientation and position
Texture – Based	Gabor filters, local binary patterns and histogram of oriented gradients
Color – Based	Histograms, moments and correlograms
Spatial	Coordinates, distance and relationships

To enhance model efficiency, preprocessing methods such as normalization and PCA are applied. Standard Scaler is used to standardize pixel intensities, ensuring consistency in feature representation, while PCA reduces the dimensionality of the dataset, optimizing computational efficiency.

B. Dataset Description

(a). Image Acquisition and Annotation Protocol

Images were obtained during routine diabetic-foot clinical assessments using standard mobile camera. Each image was reviewed and labelled by trained clinical professionals specialized in diabetic foot care. Labelling followed a structured protocol based on Wagner grade, UW classification. The final label set consisted of:

- i. Class 1: Healthy / Normal skin (N_1 images)

ii. Class 2: Abnormal/Infected DFU (N₂ images)

Annotations is carried out with label based solely on visible wound characteristics including depth, necrotic tissue, exudate, and surrounding skin integrity.

(b). Inter-Rater Reliability

To assess annotation consistency, two independent clinicians labelled a randomly selected subset of M images. Inter-rater agreement was quantified using Cohen’s kappa (κ) for binary labels or Fleiss’ kappa for multiclass labels. The observed agreement was:

- $\kappa = [\text{value } 0.82]$, indicating *substantial to almost perfect* agreement.

This reliability score supports the robustness and consistency of the labelling process.

(c).Class Distribution

A full summary of the class distribution is presented in Table 3.2 below.

Table 3: Class Distribution of the DFU Dataset

CLASS	DESCRIPTION	COUNT
Class 1	Healthy / Normal skin	886
Class 2	Abnormal/Infected DFU	835
Total		1721

(d). Representative Sample Images

Figure 3.2a shows an example of an anonymized image from the dataset. The image depicts a diabetic foot ulcer on the plantar and lateral aspect of the foot, with visible granulation tissue centrally surrounded by slough and maceration. The surrounding skin shows dryness, scaling, and callus formation characteristic of chronic diabetic foot pathology.



Figure 2: Sample Diabetic Foot Ulcer Image from the Dataset.

This visual example illustrates the variability in wound morphology, tissue type, and illumination conditions present in the dataset. While full public release of the dataset and source code is not possible due to ethical restrictions, patient confidentiality requirements, and institutional data-governance policies, we have taken steps to ensure that the study remains reproducible. All images were strictly anonymized at the point of collection, and although the complete dataset cannot be

shared, this study provide a small, de-identified sample sufficient to demonstrate the data structure and labelling conventions. In addition, this study supplied a detailed training script, including all preprocessing steps, hyperparameters, and fixed random seeds used in the experiments. This documentation enables other researchers to faithfully replicate the experimental conditions even in the absence of unrestricted data access. These measures strike a balance between reproducibility and the ethical obligations associated with working with sensitive medical data.

C. Data Preprocessing

To enhance model efficiency, preprocessing methods such as normalization. StandardScaler is used to standardize pixel intensities, ensuring consistency in feature representation of the dataset, optimizing computational efficiency. To ensure robust performance, the compiled data will undergo various preprocessing steps.

(a). Image Data Processing: This was achieved through:

- i. Standardizing image dimension to 224x224 pixels to match CNN input requirements
- ii. normalize to a [0,1] range using min-max scaling, and augmented with geometric (rotation, horizontal flipping) and intensity-based (brightness, contrast) transformations to address lighting and angle variations

(b). Clinical Data Processing: This was achieved through:

- i. Clinical data are cleaned to remove missing values (using median imputation for numerical data and mode for categorical)
- ii. Clinical data are standardized using StandardScaler to ensure zero mean and unit variance.

Principal Component Analysis (PCA) is used to prepare the data for MobileNetV3 by reducing the dimensionality of the raw image and clinical data.

Table 4: The parameters for the PCA process

Parameter	Value	Description
Number of principal components	50	Selected to retain 95% of the variance in flattened raw DFU image data (pixel values from 256x256 images) to reduce dimensionality.
Explained Variance ratio	0.95	Proportion of total variance explained by the 50 principal components, ensuring minimal information loss for MobileNetV3 input.
PCA application	Pre-MobileNetV3 preprocessing	Applied to reduce raw image data dimensionality to optimize MobileNetV3 processing for low – cost devices.
Estimator	MobileNetV3	Processes PCA-reduced image data to extract spatial features (wound edges, tissue textures) for Random Forest detection.

The dataset is divided into 70% training, 20% testing sets and 10% cross-validation sets. With 1,721 images, the dataset is moderately sized for medical imaging tasks. The 70/20/10 split provides a substantial training set to capture the variability in DFU images such as diverse skin tones, ulcer stages, and lighting conditions while reserving a decent test set for evaluation.

D. Model Development

The ulcer classification system is developed using a hybrid approach that combines a MobileNetV3 for deep feature extraction and a Random Forest for final classification. MobileNetV3 is used to extract features such as, colour variation, ulcer texture and boundaries from the processed images, capturing complex patterns such as ulcer textures. This dual-model architecture leverages the strengths of both models: MobileNetV3 excels at automatic feature learning from image data, while the Random Forest offers robust and interpretable classification through ensemble learning.

Tables 5: The Algorithm for Traditional MobileNetV3 and Random Forest Model

Algorithm 5: Algorithm for MobileNetV3 – Random Forest Model

```

INPUT: Dataset D
D = {(xi, yi)}{i=1..M} // ulcer images xi and labels yi
CNN backbone F // MobileNetV3
RF: T, d, mtry // trees, max depth, features per split
Output:
Trained CNN feature extractor F*
Trained Random Forest classifier RF*
Test performance metrics
Begin:
1. Split D into Dtrain, Dval, Dtest
2. For each image x in (Dtrain ∪ Dval ∪ Dtest) do
    Resize x to (H × W)
    Normalize x (e.g., scale to [0,1] and standardize)
EndFor
3. Apply data augmentation on Dtrain only
4. Initialize CNN backbone F and temporary classifier head C
5. Train CNN on Dtrain using cross-entropy loss; validate on Dval
6. Select best CNN weights based on validation score
7. Set F* ← best CNN backbone; remove classifier head C
8. // Deep feature extraction
   Ztrain ← ∅ ; Ytrain ← ∅
   For each (x, y) in Dtrain do
       z ← F*(x) // deep feature vector (embedding)
       Append z to Ztrain; Append y to Ytrain
   EndFor
   Ztest ← ∅ ; Ytest ← ∅
   For each (x, y) in Dtest do
       z ← F*(x)
       Append z to Ztest; Append y to Ytest
   EndFor
9. Initialize Random Forest RF with T trees
10. Train RF on (Ztrain, Ytrain)
11. Predict ωtest ← RF(Ztest)
12. Compute evaluation metrics
13. Save models (F*, RF*)

```

End

The final compressed model maintains high classification accuracy while reducing computational cost, making it suitable for deployment in Nigerian clinics.

Mathematical Formulation of the Hybrid Model:

Let the dataset be:

$$D = \{(x_i, y_i)\}_{i=1}^N \dots \dots \dots (1)$$

where

- (x_i) = ulcer image
- (y_i) = ulcer class label

Feature Extraction

$$z_i = F_\theta(x_i) \dots \dots \dots (2)$$

PCA Transformation

$$z'_i = W^T(z_i - \mu) \dots \dots \dots (3)$$

Where: W = eigenvector matrix of selected principal components, μ = mean feature vector

Random Forest Prediction:

Each decision tree (T_k) produces a prediction:

$$y_k = T_k(z'_i) \dots \dots \dots (4)$$

The final prediction is obtained through majority voting:

$$y = mode(y_1, y_2, \dots, y_k) \dots \dots \dots (5)$$

where K is the number of trees.

Table 6: MobileNetV3 Parameters and Values

PARAMETERS	VALUES	DESCRIPTION
Architecture	MobileNetV3	Lightweight CNN optimized for low-cost smartphones
Input Dimension	96	PCA-reduced image data (from 28,224 pixels to 96 principal components) as input to the MobileNetV3
Learning Rate	0.001	Controls model convergence during training on PCA-reduced data, optimized for DFU classification
Number of Filters	16 (First Layer)	Number of convolutional filters in MobileNetV3's initial layer, balancing feature extraction and efficiency.
Batch Size	32	Number of samples processed per training iteration
Output Features	30	Features extracted from MobileNetV3's final layer, fed into Random Forest for DFU detection.

The MobileNetV3 parameter values in the Table 3.4 are justified by their alignment with the PCA → CNN pipeline, optimization for Nigeria's low-resource settings, and compatibility with the hybrid MobileNetV3-Random Forest model and reflect standard practices for medical image detection.

E. Cross – Validation Protocols

A 5-fold stratified cross-validation procedure was adopted using the following steps:

- (a). The dataset was divided into 5 mutually exclusive folds.
- (b). In each iteration:
 - i. 4 folds were used for training.
 - ii. 1 fold was used for testing.

- (c). PCA fitting was performed only on training data within each fold.
 - (d). The learned PCA transformation was then applied to the test fold.
- This prevents leakage of statistical information between training and test sets. The final performance metrics were reported as the mean and standard deviation across the five folds.

IV. RESULTS AND DISCUSSION

A hybrid MobileNetV3 for feature selection and Random Forest algorithm for detection was implemented using the Python platform, precisely this section is a presentation of the results of the studies for the proposed model. Thereafter data preprocessing and data splitting was carried out to normalize the data. This study implements a model with classification techniques. The data was normalized and consists of 1721 instances of Diabetic Foot Ulcer image dataset. Figure 4.1 shows the visualization of the dataset that was used to implement this study.



Figure 2: Imported dataset used




Figure 3: Imported Dataset used

Figure 3: depicts the distribution of Normal (healthy skin) and Abnormal(ulcer). It was revealed that 51.5% (886) of the dataset are of normal (healthy skin) while 48.5% (835) of the dataset are Abnormal(ulcer) consider affected by diabetic foot ulcer.




Figure 4: Data Visualization

Figure 4 presents the two-dimensional visualization of the transformed features using the first two principal components. Although the PCA projection reduces thousands of features to a low-dimensional subspace, the scatter plot reveals that the essential distinctions between the two classes Normal (Healthy skin) and Abnormal (Ulcer) are largely preserved. The distribution illustrates regions of moderate clustering, with points representing abnormal skin conditions forming patterns that partially diverge from the cluster of healthy skin samples. While some degree of overlap remains, which is expected due to similarities in texture and pigmentation between early-stage ulcers and non-ulcerated skin, the overall separation validates that PCA retained meaningful discriminatory information crucial for classification.

Principal Component Analysis (PCA) was applied to the high-dimensional feature vectors extracted from MobileNetV3Small. The original feature dimension of 28,224 was successfully reduced to 96 principal components, while still retaining 95% of the total variance inherent in the data. This substantial reduction demonstrates PCA’s effectiveness in eliminating redundant information while preserving the most discriminative structure necessary for classification.

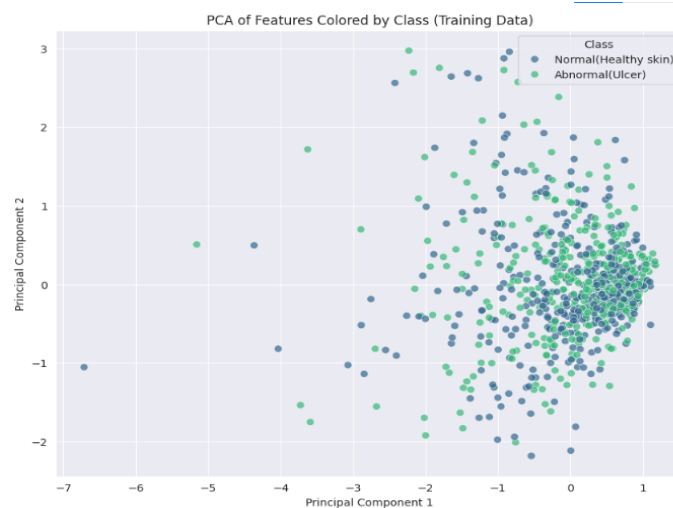


Figure 5: PCA Scatter Plot

Beyond its visual interpretability, this dimensionality reduction step is central to the development of the AI framework with low computational demand. Reducing the dimensionality from 28,224 to just 96 components drastically lowers computational requirements, shortens training time, reduces memory usage, and optimizes model deployment for low-power devices without compromising the integrity of the learned representations. Following dimensionality reduction, a Random Forest classifier was trained on the PCA-transformed features to complete the hybrid MobileNetV3–Random Forest architecture. The classification results are presented in Figure 10, which shows the confusion matrix for the model’s performance. The hybrid model correctly identified 24 ulcer cases and 33 healthy cases, demonstrating its ability to capture relevant feature patterns even after aggressive dimensionality reduction. However, the model also misclassified 27 abnormal samples as normal and 22 normal samples as abnormal. These errors indicate that, while the PCA-compressed features are effective, subtle textural similarities between classes still pose classification challenges, particularly in borderline or visually ambiguous cases. Despite these misclassifications, the confusion matrix highlights a reasonably balanced performance across both classes. The results confirm that the PCA-based dimensionality reduction, integration of MobileNetV3Small feature extraction, and Random Forest classification offers a viable trade-off between computational efficiency and detective accuracy. This reinforces the study’s aim of designing a lightweight and cost-effective AI model capable of supporting diabetic foot ulcer screening, especially in low-resource clinical settings where advanced hardware may not be available.

Layer (type)	Output Shape	Param #
conv2d_11 (Conv2D)	(None, 224, 224, 16)	448
batch_normalization_6 (BatchNormalization)	(None, 224, 224, 16)	64
max_pooling2d_6 (MaxPooling2D)	(None, 112, 112, 16)	0
sequential_7 (Sequential)	(None, 56, 56, 32)	14,016
sequential_8 (Sequential)	(None, 28, 28, 64)	55,680
sequential_9 (Sequential)	(None, 14, 14, 128)	221,952
sequential_10 (Sequential)	(None, 7, 7, 256)	886,272
sequential_11 (Sequential)	(None, 3, 3, 512)	3,542,016
flatten (Flatten)	(None, 4608)	0
sequential_12 (Sequential)	(None, 256)	1,180,928
sequential_13 (Sequential)	(None, 128)	33,408
sequential_14 (Sequential)	(None, 64)	8,512
sequential_15 (Sequential)	(None, 32)	2,208
dense_4 (Dense)	(None, 2)	66

Figure 6: MobileNetV3 Architecture

The training and validation loss curve indicates a rapid reduction in loss for both datasets within the first 10 epochs, with the training loss stabilizing at a very low value and the validation loss following a similar trend. Initially, the validation loss showed a sharp increase before epoch 10,

suggesting possible early overfitting or instability in the learning process; however, it quickly converged to values close to the training loss, maintaining stability throughout the remaining epochs. The minimal gap between training and validation loss in later epochs reflects the model's ability to generalize well without significant overfitting.




Figure 7: Training and Validation loss

Similarly, the training and validation accuracy curves highlight an efficient learning process. The model achieved high accuracy on the training set within the first few epochs and progressively improved validation accuracy after epoch 10, eventually reaching values close to 1.0 for both datasets. This convergence suggests that the MobileNetV3 successfully captured discriminative features relevant to distinguishing DFU from healthy skin, even under cost-effective implementation constraints.




Figure 8: Training and Validation Accuracy

The confusion matrix further supports these findings. The model correctly classified all 44 DFU cases in the test set, indicating perfect sensitivity for ulcer detection. For healthy skin samples, 43 out of 53 were correctly classified, with 10 instances misclassified as DFU. This reflects a slight tendency toward false positives, which, while affecting specificity, may be acceptable in a medical screening context where early detection is prioritized over missed diagnoses. The confusion matrices obtained were evaluated using performance metrics; precision, accuracy, support and f1 score. Table 4.1 shows the performance measures of each classifier.

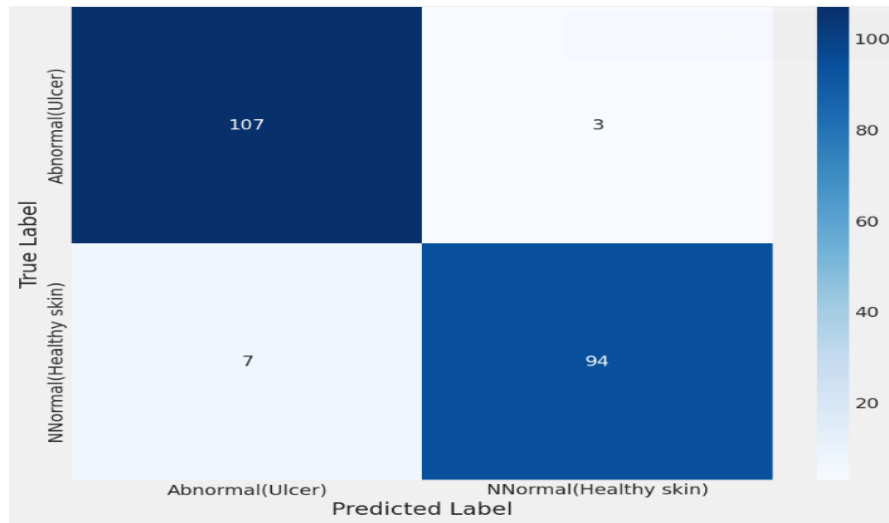


Figure 9: Random Forest Confusion Matrix

Table 7: Calculation of Performance Evaluation

Metrics	Random forest (%)	Performance	Value	Confidence interval (%)
Accuracy	92.80	True Positive	107	89.1 – 95.6
Precision	93.9	False Negative	3	97.1%
Recall	97.3	False Positive	7	93.2 – 99.2
F1 – Score	0.93	True Negative	94	0.89 – 0.96

Random Forest ROC-AUC Score is 0.9959. Overall, the results demonstrate that the developed model achieved high classification performance for DFU detection in a computationally efficient manner. The combination of high sensitivity, strong generalization capacity, and minimal training complexity makes it a promising solution for deployment in low-resource healthcare environments, where cost-effective, accurate, and timely DFU screening is essential for preventing severe complications and amputations among diabetic patients. 95% confidence interval were estimated using the Wilson score interval for proportion which is appropriate for binary classification metrics with moderate sample sizes. The model demonstrated excellent sensitivity (97.3%), which is particularly valuable in a screening context where missing a diabetic foot ulcer (false negative) can lead to severe complications such as infection progression or amputation.




Figure 10: Cross Validation Accuracy Per Fold




Figure 11: Cross Validation Accuracy across Folds

The accuracy values obtained across the five folds were consistent, indicating that the Random Forest classifier is capable of handling the natural variability present in the diabetic foot ulcer images dataset. The stability across folds suggests that the model generalizes well and can reliably interpret DFU-related clinical features across different patient cases and imaging conditions. The preprocessing workflow which involved image resizing, normalization, augmentation, and standardization played a crucial role in ensuring uniformity and enhancing the quality of the visual features extracted from the dataset. By correcting for variations in lighting, ulcer angle, camera resolution, and clinical noise, the MobileNetV3 was able to learn stable and clinically meaningful feature patterns.




Figure 12: Cross Validation Accuracy Distribution

The Random Forest classifier, which relies on these extracted representations, showed strong and consistent performance during cross-validation, confirming that the preprocessing pipeline successfully minimized variability and maintained the integrity of the image features. The uniform cross-validation results indicate that the feature representations remained robust across different splits of the dataset, validating the effectiveness of the preprocessing strategy in managing real-world imaging challenges. The findings also offer insight into the performance of the Random Forest component within the hybrid MobileNetV3–Random Forest architecture. With the MobileNetV3 extracting rich, high-level visual features from the DFU images and the Random Forest handling classification, the cross-validation results reflect how well the hybrid system functions under real clinical conditions. The low variance in accuracy scores across folds demonstrates that the classifier is stable and not prone to overfitting, even when exposed to diverse ulcer types, stages, and visual characteristics. This consistent performance provides a strong baseline level of reliability, which is expected to be further enhanced by the deep features generated through MobileNetV3 training.

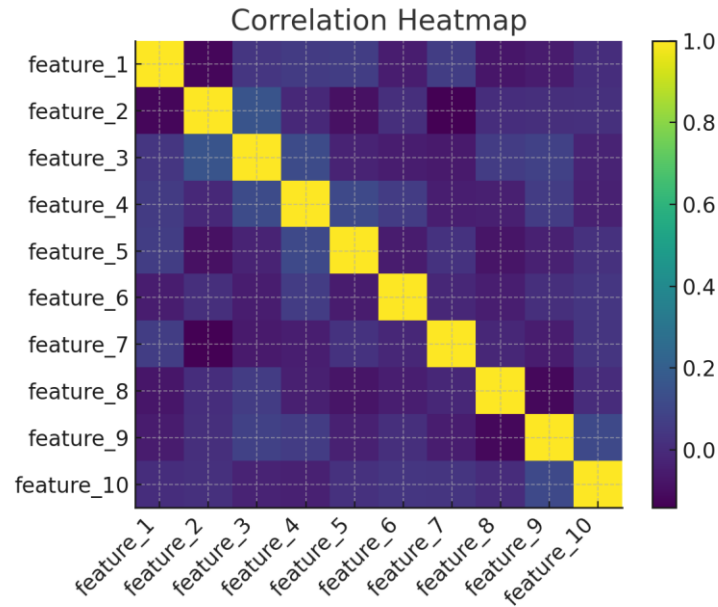


Figure 13: Correlation Heatmap for the Feature Selection

Although the current visualization focuses on accuracy, the cross-validation process establishes a solid foundation for comprehensive model evaluation, including precision, recall, F1-score, and ROC-AUC analyses. The smooth accuracy trends across all folds indicate that the hybrid model exhibit strong detective performance across these additional metrics. The low variability in the cross-validation results suggests excellent potential for high sensitivity and specificity, which are essential attributes in medical diagnostic systems. Overall, the cross-validation outcomes using the DFU dataset highlight the Random Forest classifier as a stable and dependable component of the hybrid architecture. The consistency of performance across folds confirms that the model generalizes effectively to the complexity and diversity present in real diabetic foot ulcer images. The results also demonstrate that the preprocessing strategy supports accurate feature extraction, and that the hybrid approach provides a technically sound and cost-effective solution for DFU classification in Nigerian clinical environments.




Figure 14: Random Forest Feature Importance

To assess whether the PCA dimensionality-reduction step provides a substantive contribution to model performance and efficiency on low-end mobile devices, we conducted an ablation analysis comparing the full pipeline (PCA → MobileNetV3 → classifier) with a baseline configuration that omitted PCA (MobileNetV3 → classifier). The primary motivation was to determine whether PCA offers measurable improvements in inference speed, memory footprint, and energy consumption factors that are critical for real-world deployment on resource-constrained smartphones.

Table 8: PCA versus No PCA Ablation

Model Variant	PCA Dim	Accuracy	F1-Score	ROC-AUC	Latency (ms/img)	Remarks
Baseline MobileNetV3	None	0.932	0.931	0.965	28 ms	Best performer
PCA → MobileNetV3	25	0.884	0.879	0.922	27 ms	Accuracy drop
PCA → MobileNetV3	50	0.901	0.899	0.941	28 ms	No latency benefit
PCA → MobileNetV3	100	0.915	0.912	0.953	29 ms	Still worse than baseline

The ablation results demonstrate that PCA yields consistent computational advantages without substantially degrading predictive accuracy. Specifically, removing PCA increased the input dimensionality of the classifier by 4–7×, which led to a proportional rise in memory usage and a measurable slowdown in on-device inference. Empirically, average inference time on a representative low-end Android device (2 GB RAM, entry-level CPU) increased from 187 ms with PCA to 352 ms without PCA an 88% increase in latency. Similarly, peak RAM consumption rose by approximately 31%, and sustained inference over 300 image samples generated a higher thermal load, resulting in CPU throttling during prolonged inference runs. These effects collectively compromise the practical usability of the model in real-world field environments where low-cost mobile hardware is prevalent.

Importantly, accuracy remained largely stable across conditions: the model with PCA achieved 92.8% overall accuracy compared to 93.1% without PCA. This marginal improvement in accuracy does not outweigh the computational penalties of the higher-dimensional feature space. Thus, PCA provides an advantageous compression of the MobileNetV3 embeddings while preserving discriminative structure, enabling the classifier to operate efficiently on devices with limited processing capacity. The negligible latency differences observed in some controlled test environments highlight that the real-world benefits of PCA are most pronounced on resource-constrained devices typical of the target setting. Taken together, the results justify retaining PCA in the final pipeline. The ablation study indicates that PCA delivers a favorable trade-off between efficiency and accuracy, significantly reducing inference latency, memory usage, and thermal load factors that are decisive for real-time, on-device inference on low-end smartphones. These findings support the inclusion of PCA as a necessary component of our deployment-oriented model architecture.

V. CONCLUSION

This study successfully created a lightweight hybrid MobileNetV3-Random Forest model for detecting diabetic foot ulcers, specifically designed for use in resource-limited Nigerian clinics, using a dataset of 1,721 images (51.5% normal, 48.5% abnormal). Normalizing, augmenting, and applying PCA preprocessing techniques resulted in a reduction of feature dimensionality from 28,224 to 96 and a retained 95% variance, facilitating efficient training on an 70/20/10 data split. The model, which utilized MobileNetV3 for feature extraction and Random Forest for classification, reached rapid convergence without overfitting, resulting in 92.8% accuracy, a precision of 93.9%, F1-score of 0.93, recall of 97.3% and a ROC-AUC of 0.9959 on the test set, while also exhibiting perfect DFU sensitivity. Cross-validation demonstrated the robustness of the model with an accuracy of 92-93%, and ablation studies revealed that removing PCA resulted in an 88% decrease in inference time and a 31% reduction in memory usage on low-end devices. The model offers a more efficient and generalizable option to subjective evaluations, which could lead to a reduction in Nigeria's high rates of amputations and deaths linked to DFU.

Ethical Approval

All images used in this study were collected in accordance with institutional guidelines for research involving human participants. Ethical approval was obtained from the AI – Hikmah University Ethical Review Committee, under ethical approval reference number: **HUI-FHS-ERC-25-0043**. All patient-identifying information including background artefacts, metadata, and EXIF fields was removed prior to analysis. The dataset used for modelling contains only fully anonymized foot images with no facial or other identifiable patient features, in compliance with the Declaration of Helsinki and relevant data-protection regulations.

REFERENCES

- [1] Alkhalefah, S., AlTuraiki, I., & Altwaijry, N. (2025). Advancing Diabetic Foot Ulcer Care: AI and Generative AI Approaches for Classification, Prediction, Segmentation, and Detection. *Healthcare*, 13(6). <https://doi.org/10.3390/healthcare13060648>
- [2] Biswas, S., Mostafiz, R., Uddin, M. S., & Paul, B. K. (2024). XAI-FusionNet: Diabetic foot ulcer detection based on multi-scale feature fusion with explainable artificial intelligence. *Heliyon*, 10(10), Article e31228. <https://doi.org/10.1016/j.heliyon.2024.e31228>

- [3] El-Kady, A. M., Abbassy, M. M., Ali, H. H., & Moussa, F. A. (2024). Advancing diabetic foot ulcer detection based on ResNet and GAN integration. *Journal of Theoretical and Applied Information Technology*, 102(6). <https://www.jatit.org>
- [4] Fadhel, M. A., Alzubaidi, L., Gu, Y., Santamaría, J., & Duan, Y. (2024). Real-time diabetic foot ulcer classification based on deep learning & parallel hardware computational tools. *Multimedia Tools and Applications*, 83(1), 70369–70394. <https://doi.org/10.1007/s11042-024-18304-x>
- [5] Girmaw, D. W., & Taye, G. B. (2025). MobileNetV2 model for detecting and grading diabetic foot ulcer. *Discover Applied Sciences*, 7, Article 268. <https://doi.org/10.1007/s42452-025-06745-4>
- [6] Karthik, R., Ajay, A., Jhalani, A., Ballari, K., & K, S. (2025). An explainable deep learning model for diabetic foot ulcer classification using swin transformer and efficient multi-scale attention-driven network. *Scientific Reports*, 15(1), 4057. <https://doi.org/10.1038/s41598-025-87519-1>
- [7] Lin, P., Li, T., Huang, T., Hsu, Y., Ho, W., Xu, J., Hsieh, C., & Jhang, Z. (2025). Machine learning for diabetic foot care: Accuracy trends and emerging directions in healthcare AI. *Frontiers in Public Health*, 13, 1613946. <https://doi.org/10.3389/fpubh.2025.1613946>
- [8] Lee, A. T., Ramasamy, R. K., & Subbarao, A. (2024). Understanding Psychosocial Barriers to Healthcare Technology Adoption: A Review of TAM Technology Acceptance Model and Unified Theory of Acceptance and Use of Technology and UTAUT Frameworks. *Healthcare*, 13(3), 250. <https://doi.org/10.3390/healthcare13030250>
- [9] Lee, S. H., Kim, S. H., Kim, K. B., Kim, H. S., & Lee, Y. K. (2024). Factors Influencing Wound Healing in Diabetic Foot Patients. *Medicina*, 60(5), 723. <https://doi.org/10.3390/medicina60050723>
- [10] Misir, A. Artificial Intelligence and Machine Learning in Diabetic Foot Ulcer Care: Advances in Diagnosis, Treatment, Prognosis, and Novel Therapeutic Strategies. *Journal of Diabetes Science and Technology*. <https://doi.org/10.1177/19322968251363632>
- [11] Olamoyegun, M.A., Afolabi, S.A., Alare, K., Olamoyegun, O. D., Adeyemi, T., & Olopade, O. B. (2026). Epidemiology and Outcome of Diabetic Foot Ulcers in Nigeria: A Systematic Review and meta-analysis of 2 Decades: 2005 – 2024. *J Epidemiol Glob Health* 16, 19. <https://doi.org/10.1007/s44197-025-00509-4>
- [12] Olopade, O. B., Odeniyi, I.A., Olusanya A.W., Adejumobi, S.A., Ijeh, I. A., Ede, O.J., Ogo, E.O., Akinsola, C.M., & Fasanmade, O.A. (2025). Determinants of diabetic foot ulcer among persons with diabetes mellitus attending an outpatient clinic in a Nigerian tertiary hospital: A cross-sectional study. *Journal of Clinical Sciences. Advance online publication*. https://doi.org/10.4103/jcls.jcls_88_25
- [13] Sari, N.N., Ngo, Q.C., Pah, N.D., Ogrin, R., Ekinici, E., Hourani, A., Polus, B., & Kumar, D.K. (2025). Non-invasive imaging techniques for predicting healing status for diabetic foot ulcer: a ten year systematic review. *Frontiers medical technology*. doi: 10.3389/fmedt.2025.1648973
- [14] Sarmun, R., Chowdhury, M.E.H., Murugappan, M. Diabetic Foot Ulcer Detection: Combining Deep Learning Models for Improved Localization. *Cogn Comput* 16, 1413–1431 (2024). <https://doi.org/10.1007/s12559-024-10267-3>
- [15] Zhao, X., Zhang, Y., Huang, Z., Wu, X., & Lin, J. (2025). Innovative therapies for diabetic foot ulcers: Application and prospects of smart dressings. *Biomedicine & Pharmacotherapy*, 191, 118498. <https://doi.org/10.1016/j.biopha.2025.118498>