iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: VII Month of publication: July 2026

DOI: https://doi.org/10.22214/ijraset.2026.82253

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

A Hybrid IncepX-Ensemble Model for Multiclass Skin Lesion Detection and Classification

Krishna Priya MP[1], Althaf Yazin PP[2], Nihidha Parvin M[3], Prof. Richu Shibu[4], Prof. Shruthi Susan Moncy[5]
[1]Department of Computer Science and Engineering, Mar Athanasius College of Engineering (Autonomous), Kothamangalam

Abstract: Skin cancer is one of the most common cancers worldwide, and melanoma is especially dangerous if it is not de-tected early. Early diagnosis can improve survival rates. Dermo-scopic images are widely used by dermatologists to examine skin lesions. But manual analysis can be difficult and depends heavily on clinical experience. So automated methods for skin lesion classification have gained much attention. A hybrid deep learning model is proposed for multiclass skin lesion classification, aiming to improve diagnostic accuracy and robustness across diverse lesion types. The model combines two pretrained architectures, InceptionV3 and Xception, to extract features from dermoscopic images. The features obtained from both networks are fused to improve the overall representation. Transfer learning is applied to make use of pretrained weights, and data augmentation techniques are used to reduce overfitting and improve model stability. The proposed method is evaluated on a seven-class dermoscopic image dataset. The results show that the ensemble model performs better than the individual models. In addition, explainable AI techniques are used to highlight the important regions in the images that influence the model's predictions, which improves interpretability and supports its practical use.

I. INTRODUCTION

The number of reported skin cancer cases has been grad-ually increasing around the world in recent years. Early diagnosis is very important to reduce the death rate and improve survival rates. It has created a strong need for reliable diagnostic support systems. Among the different types of skin cancers, melanoma is especially dangerous because it is highly aggressive and can spread quickly if it is not detected at an early stage. Dermoscopy is a commonly used imaging method that helps dermatologists observe skin lesions more clearly without the need for invasive procedures. This technique provides clearer view of the lesions. However, the final diagnosis still depends on the doctor's experience. Differences in doctors' experience can lead to inconsistent diagnoses. This creates a need for automated support systems. Artificial intelligence, especially deep learning, is widely used in medical image analysis today. Convolutional Neural Networks (CNNs) are commonly used for image classification tasks, including disease detection. Because they can learn repeated patterns directly from images. In skin lesion analysis, CNNs can identify important features such as color, texture, and shape. This minimises the need for manual feature extraction.

However, many existing studies fully depend on a single deep learning model. A single architecture may not always capture the complex and multi-scale characteristics of dermo-scopic images. The main challenge in classification is that different skin lesion classes often appear visually similar. Ensemble learning is a practical solution to these limitations. The performance of the model improves when multiple deep learning systems are combined. This is because features from different models are merged. Also, Prediction stability im-proves when results are combined. As a result, individual weaknesses are reduced naturally when multiple models work together. In our work, we use a combined approach based on Incep-tionV3 and Xception to classify different types of skin lesions. The main goal is to achieve high accuracy and consistent results for real-world medical applications.

II. RELATED WORKS

Initially, skin lesion classification mainly relied on manually extracted features. These features were based on character-istics such as color, shape, and border. The ABCD rule (Asymmetry, Border irregularity, Color variation, Diameter) is commonly used to derive these features. However, its success rate depends on human expertise. Traditional machine learning models like Support Vector Machines, Random Forest, and k-Nearest Neighbors were used for classification. These methods were also sensitive to image quality, and their performance decreased when images contained noise, poor lighting, or artifacts such as hair. Deep learning has enhanced the analysis of dermoscopic skin images. CNN models can automatically learn important features from images without the need for handcrafted fea-ture extraction. Several models, such as AlexNet, VGG16, ResNet, InceptionV3, and Xception, have been widely used for skin lesion classification. They perform better than traditional methods because they can capture complex patterns in the images. Public datasets like ISIC, PH2, and HAM10000 have supported the training and evaluation of these models.

As a result, classification performance has improved compared to earlier methods. Moreover, the availability of large annotated datasets has further improved their performance. Several challenges still remain in skin lesion classification. One major issue is class imbalance. This occurs because most datasets contain more images of non-cancerous lesions than cancerous ones, making it difficult for models to correctly identify rare but serious cases. To address this issue, data augmentation techniques are used to increase the diversity of training data. Methods such as rotation, flipping, scaling, zooming, and slight changes in brightness and colour help im-prove model performance. To further handle these challenges, advanced techniques such as transfer learning and ensemble learning are used. Transfer learning is widely used in medical imaging tasks. Pretrained models, such as those trained on ImageNet, are fine-tuned for specific tasks like skin lesion classification. This approach is useful when only a limited amount of medical data is available. This helps to enhance model performance. In addition, ensemble learning combines multiple models instead of using only a single model. These methods help overcome the limitations of a single model. As a result, accuracy is improved, and predictions become more stable. In addition to improving performance, it is also important to understand how these models make decisions. Recent studies focus on both model accuracy and in-terpretability. Explainable AI methods such as Grad-CAM, LIME, and saliency maps are used to highlight regions in the input image that influence predictions. This makes the decision process more transparent to medical professionals. This helps increase trust in the models. It also address the black-box nature of AI models. However, some challenges still remain, including class imbalance and reduced performance across different datasets. Many models perform well in controlled environments but face difficulties in real-world applications due to overfitting and underfitting. Combining multiple deep learning models into a single system can help handle these issues. This helps improve performance and reduce the limi-tations of individual models. Previous studies show that skin lesion classification systems perform better when transfer learning and ensemble deep learning methods are used together. Many works also use data augmentation techniques to improve model generalisation and reduce overfitting caused by limited dermoscopic datasets. In addition, explainable AI methods are becoming important for understanding model predictions and improving trust in medical applications. These observations motivated the devel-opment of the proposed Hybrid IncepX-Ensemble framework with integrated explainability analysis.

III. METHODS

This section describes the proposed approach for skin lesion classification. The method combines a deep ensemble model with explainable AI techniques. It helps to improve both accuracy and interpretability. The overall pipeline includes four main stages: data preparation, feature extraction using a dual-model ensemble, classification, and explanation-based evaluation. In addition, a modified LIME method is introduced to generate more stable and lesion-focused explanations.

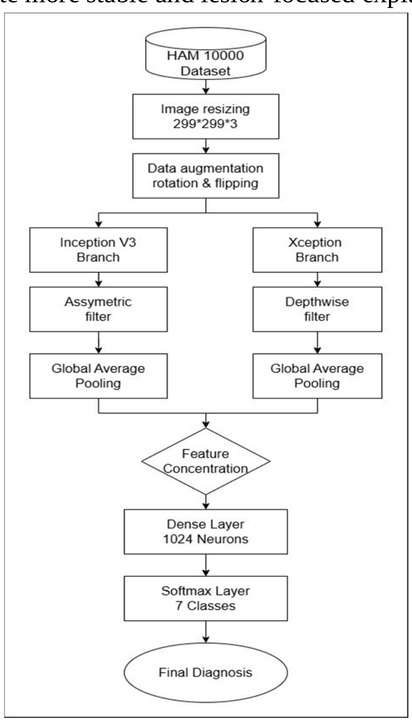

Fig. 1. Dual-branch ensemble architecture using InceptionV3 and Xception for feature extraction.

A. Data Preparation

The HAM10000 dataset was used to train and test the IncepX model. It has dermoscopic images of seven types of skin lesions. All images were resized to 299 × 299 × 3 so they fit the input size of the model.The pixel values of the images were normalized to the range of 0 to 1.

This helps the model train faster and improves stability during learning. Data augmentation was applied to help the model learn better and avoid overfitting.

The images were randomly rotated and flipped. This added variations in how the lesions looked. Thus the model can recognise features better and work well on unseen images during testing.

B. Proposed Deep Ensemble Architecture

We design a two-branch CNN architecture to extract com-plementary features from skin lesion images. Features are extracted using both InceptionV3 and Xception networks.

Both networks are initialized with pre-trained ImageNet weights. This helps the designed Incepx model use previously learned features and reduces training time. The model takes the pre-processed images as input, as described in the data preparation section. This ensures consistency during training.

The first branch uses InceptionV3 to extract features from the input images. It captures information at different scales using filters of different sizes.

It helps in identifying textures and edges in skin lesions. The second branch uses Xception, which applies depthwise separable convolutions.

This allows the model to learn fine details and small variations in channel-wise patterns. Each branch applies Global Average Pooling to reduce the feature map size while keeping important information. The pooled features from both branches are then combined using concatenation to form a single feature vector. This step combines the strengths of both models and improves the overall representation. The combined features are passed through a fully connected layer with 1024 neurons and ReLU activation.

A dropout layer is added to reduce overfitting. Finally, a softmax layer predicts the probability for each of the seven lesion classes. The ensemble architecture helps the model capture both local and global features. This makes it more robust to variations in lesion appearance.

C. Training and Validation Details

The model was trained using the Adam optimizer, and the training was carried out for multiple epochs. Categorical crossentropy was used as the loss function for multi-class classification.

During training, the model was evaluated on a separate validation set to monitor its performance on unseen data. The training dataset consisted of 38,569 images, while the validation dataset contained 938 images across seven classes.

Validation was performed at each epoch to track accuracy and loss. Callbacks such as EarlyStopping, ReduceLROnPlateau, and ModelCheckpoint were used to improve performance and avoid overfitting. EarlyStopping helps stop training when there is no improvement and ReduceLROnPlateau adjusts the learning rate when validation performance slows down. Also, ModelCheckpoint saves the best-performing model.

D. Explainability Framework

The explainability framework used in our work combines multiple post-hoc methods to better understand the model's predictions, as shown in Fig. 2. Instead of depending on a single method, three XAI techniques, such as Grad-CAM, LIME, and SHAP are used together.

Each method provides a different type of explanation, which helps in understanding the model more clearly.

Grad-CAM is used to generate heatmaps that highlight important regions in the input image. These regions show

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

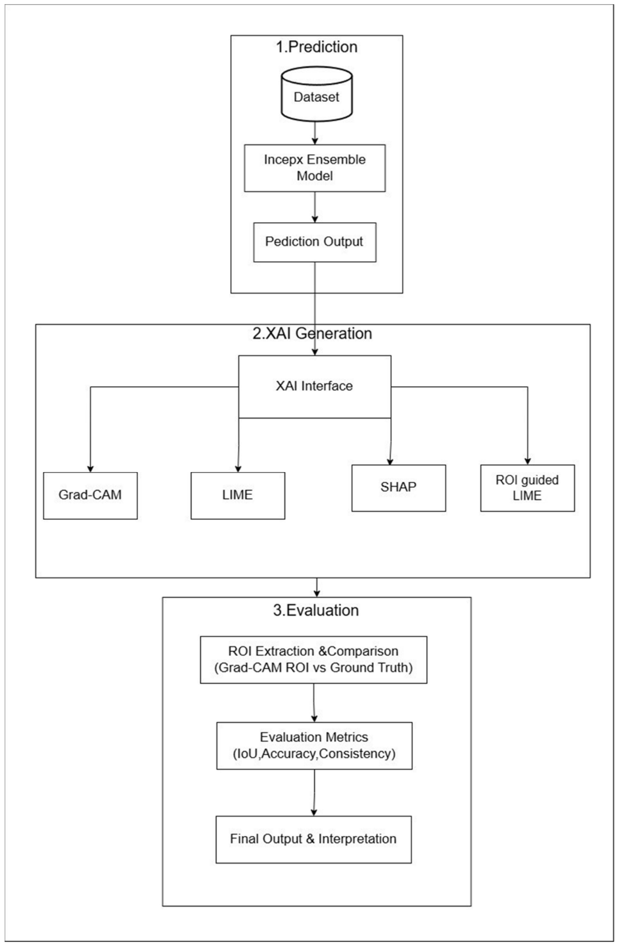

Fig. 2. Explainability generation and evaluation pipeline using Grad-CAM, LIME, and SHAP.

where the model is focusing while making a prediction. This is useful for identifying key areas of skin lesions that influence the decision.

LIME is used to explain individual predictions. It works by making small changes to the input image and observing how the prediction changes. This helps in understanding which parts of the image are important for a specific prediction.

SHAP is used to measure the contribution of different features to the final output. It assigns importance values to different regions of the image, which helps in understanding how each part influences the prediction.

The explanations generated by these methods are compared with known lesion regions to check how well the model is focusing on relevant areas. This helps in evaluating both the accuracy and reliability of the model. By combining multiple explanation methods, the framework provides a better understanding of the model's decision-making process and increases trust in clinical applications.

E. Modified LIME for Lesion-Focused and Stable Explana-tions

LIME is commonly used to explain model predictions, but its basic version is not suitable for medical images. It treats all regions of an image equally and makes random changes to superpixels. Because of this, it may highlight areas such as normal skin, hair, or lighting variations instead of the actual lesion. This reduces the reliability of the explanation in medical use.

To address this problem, a modified version of LIME is used in this work. First, a Grad-CAM heatmap is generated from the trained model to identify important regions. A threshold is applied to keep only the main activated area. This step helps in defining a clear region of interest (ROI).

LIME is then applied only within this ROI instead of the full image. This reduces the influence of background regions and focuses the explanation on the lesion area. The segmentation process is limited to this region, which results in more meaningful and relevant superpixels.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

To improve stability, random perturbations used in standard LIME are reduced. Instead, structured and consistent changes are applied to the segmented regions. Each region is modified in a controlled way, and the effect on model prediction is observed. This makes the explanations more consistent across multiple runs. A surrogate model based on Ridge regression is used to assign importance to each superpixel. This helps in identifying which parts of the lesion contribute most to the prediction. The importance values are normalized to make the explanation maps clearer and easier to compare across different images. The modified LIME method gives more focused and con-sistent explanations. It highlights the lesion areas more clearly and avoids showing unnecessary regions. This makes the explanations more reliable and easier to use in clinical settings.

IV. RESULTS

This section presents the performance of the proposed Hy-brid IncepX-Ensemble model. The evaluation includes training behavior, classification performance, and class-wise analysis. The results show how well the model learns and how ac-curately it classifies different skin lesion types. The results also highlight the reliability of the integrated explainable AI techniques.

A. Training and Validation Performance

Fig. 3 shows the training and validation accuracy and loss over epochs. The training accuracy increases gradually as the model learns from the data. At the same time, the validation accuracy also improves and remains close to the training curve. The results suggest that the model generalises well to unseen data. The loss curves show a decreasing trend for both training and validation. The decreasing loss confirms effective learning during training. There is no large gap between training and validation performance. It suggests that overfitting is limited. Overall, the curves show stable and consistent learning be-haviour.

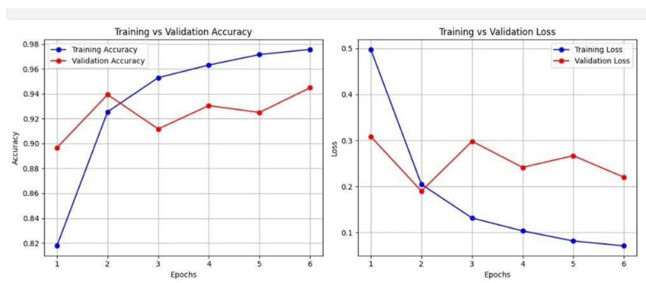

Fig. 3. Training vs. validation accuracy and loss curves of the proposed Hybrid IncepX-Ensemble model.

B. ROC Curve Analysis

The ROC curves for all seven classes are shown in Fig. 4. The curves move close to the top-left corner, which indicates good classification performance. High AUC values are ob-served for most classes. It shows that the model can clearly distinguish between different lesion types.

Even though the classes have visual differences, the model performs well across them. The learned decision boundaries remain stable and effective for multi-class classification.

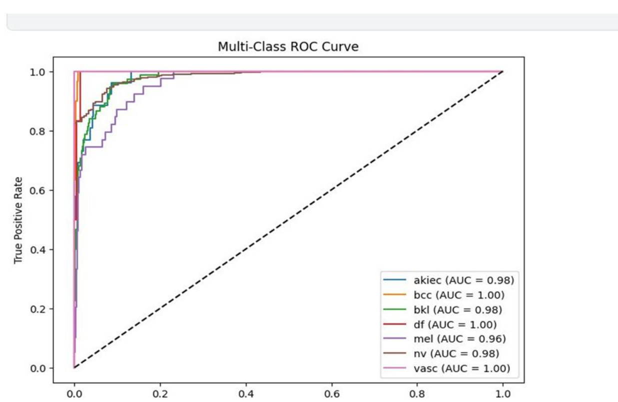

Fig. 4. Multi-class ROC curves for the seven skin lesion categories.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

C. Confusion Matrix

Fig. 5 presents the confusion matrix of the model. Most of the predictions are concentrated along the diagonal, which indicates correct classification. Most samples are classified correctly by the model.

Some misclassifications can be observed between visually similar classes, such as melanoma and melanocytic nevi. These errors are expected because of the similarity in appearance. However, the number of such errors is relatively small com-pared to correct predictions.

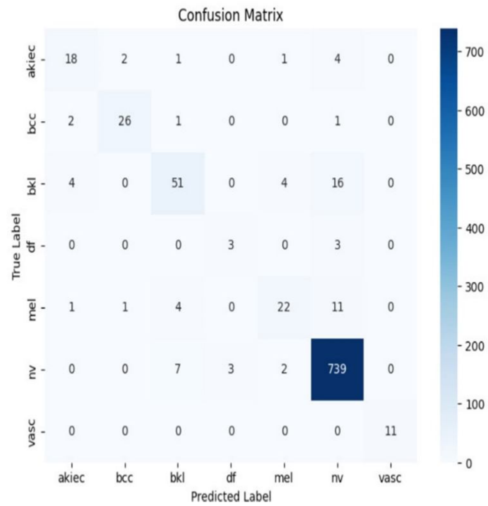

Fig. 5. Confusion matrix of the proposed model on the test dataset.

D. Classification Metrics

The classification report is shown in Fig. 6. It includes precision, recall, and F1-score for each class, along with overall accuracy.

Most classes show high precision and recall values. The results confirm good performance in identifying different lesion types. The F1-score values are also consistent across classes. The model shows a good balance between precision and recall. The overall accuracy of the model further confirms its effectiveness in multi-class skin lesion classification.

	precision	recall	f1-score	support
akiec	0.7200	0.6923	0.7059	26
bcc	0.8966	0.8667	0.8814	30
bkl	0.7969	0.6800	0.7338	75
df	0.5000	0.5000	0.5000	6
mel	0.7586	0.5641	0.6471	39
nv	0.9548	0.9840	0.9692	751
vasc	1.0000	1.0000	1.0000	11
accuracy			0.9275	938
macro avg	0.8038	0.7553	0.7768	938
weighted avg	0.9232	0.9275	0.9242	938

Fig. 6. Classification report showing precision, recall, F1-score, and overall accuracy for all lesion classes.

To further evaluate the effectiveness of the proposed ensem-ble approach, the performance of the Hybrid IncepX-Ensemble model was compared with the individual backbone networks, namely InceptionV3 and Xception. Fig. 7 shows the overall accuracy obtained by each model.

Among the compared models, the ensemble model achieved the highest accuracy. The improvement indicates that combin-ing the strengths of both architectures helps the model learn more discriminative features from dermoscopic images. While both InceptionV3 and Xception individually produced strong results, the ensemble approach provided better generalisation and more stable classification performance on the test dataset. The results demonstrate that feature fusion from multiple deep learning architectures can improve prediction reliability in multi-class skin lesion classification tasks. Fig. 8 presents a

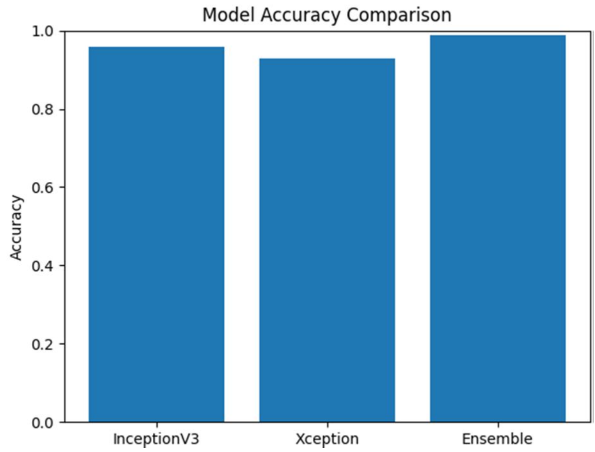

Fig. 7. Comparison of overall classification accuracy between InceptionV3, Xception, and the proposed Hybrid IncepX-Ensemble model.

comparison of precision, recall, and F1-score for the evaluated models. The ensemble model achieved consistently higher values across all three metrics compared to the individual architectures.

The precision score indicates that the ensemble model re-duces incorrect positive predictions more effectively. Similarly, the recall value shows improved capability in identifying lesion samples correctly, which is important in medical image analysis where missed detections can affect diagnosis. The higher F1-score further confirms that the proposed model maintains a balanced performance between precision and recall.

The comparative results suggest that the ensemble strategy improves both robustness and classification consistency across different lesion categories.

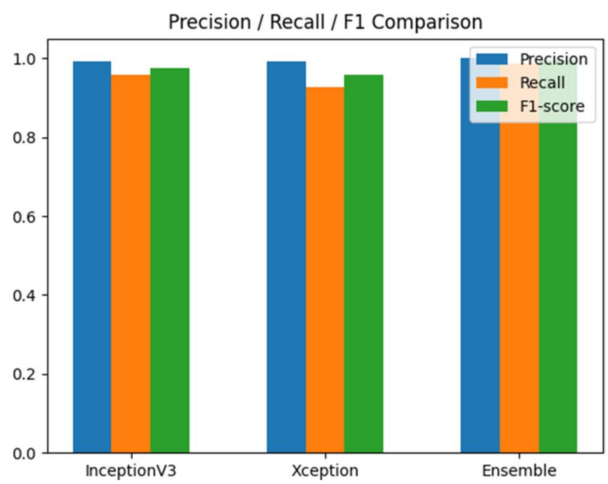

Fig. 8. Comparison of precision, recall, and F1-score between InceptionV3, Xception, and the proposed Hybrid IncepX-Ensemble model.

E. Explainable AI (XAI) Analysis

1) Comparison of XAI Methods

Different explainability methods were used to understand how the model makes decisions. Fig. 9 shows the outputs of Grad-CAM, LIME, and SHAP for a sample image. Each method highlights important regions in a different way.

Grad-CAM shows the main regions that influence the prediction. These areas usually cover the lesion and nearby important structures. LIME focuses on smaller parts of the image and shows which segments support the prediction. SHAP gives importance values to different regions, showing how each part affects the output.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

When these methods are viewed together, they give a clearer idea of how the model works. Most of the highlighted regions fall inside the lesion area, which shows that the model is focusing on meaningful features rather than background noise.

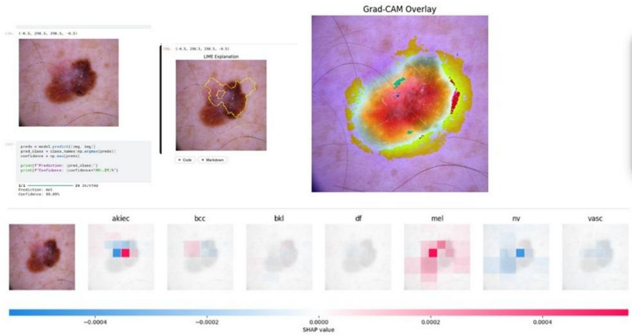

Fig. 9. Combined explainability results for a single dermoscopic image using Grad-CAM, LIME, and SHAP, highlighting clinically relevant regions influencing the model's prediction.

2) Evaluation of Modified LIME

To further study ex-plainability, the modified LIME method was compared with the standard LIME approach. This comparison was done for difficult cases such as false negatives and false positives. Fig. 10 shows the quantitative results. The modified LIME gives higher values for ROI-to-background ratio, which means more focus is placed on the lesion area. It also shows better overlap with Grad-CAM regions, as seen from improved IoU values. Another important observation is the confidence drop. When important regions identified by modified LIME are removed, the model confidence decreases more. This shows that the highlighted areas are truly important for the prediction.

3) Visual Analysis of Error Cases

Fig. 11 shows a visual comparison for a false negative case. The Grad-CAM map highlights the main lesion area clearly. However, standard LIME spreads its attention to surrounding regions and does not align well with the lesion. On the other hand, the modified LIME focuses more closely on the lesion. The highlighted regions are more compact and better aligned with the actual area of interest. The difference map also shows that most of the additional focus from the modified method lies within the lesion.

Fig. 10. Quantitative comparison of standard and modified LIME explanations for false negative (FN) and false positive (FP) cases using ROI-to-background ratio, intersection over union (IoU), and confidence drop metrics.

This improvement is important because false negatives are critical in medical diagnosis. A better explanation in such cases increases trust in the model.

The results show that the proposed Hybrid IncepX-Ensemble model performs well for multi-class skin lesion classification. The model achieved high accuracy along with good precision, recall, and F1-score values for different lesion categories. When compared with the individual InceptionV3 and Xception models, the ensemble model produced better overall performance.

The training and validation graphs showed stable learning suitable for computer-aided skin lesion analysis and other medical image analysis applications. Further improvements can be made by testing the model on larger and more diverse datasets.

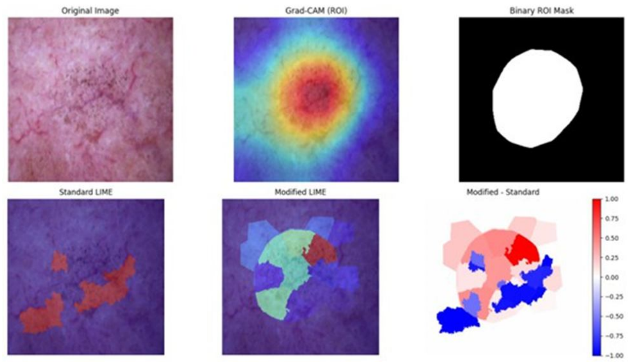

Fig. 11. Visual comparison of standard and modified LIME explanations for a representative false negative dermoscopic image, shown alongside the original image, Grad-CAM heatmap, binary lesion mask, and the difference map, highlighting improved lesion-focused interpretability.

during training with less overfitting. The ROC curves and confusion matrix also showed that the model can classify most lesion types correctly, including some visually similar classes. The explainability methods such as Grad-CAM, LIME, SHAP, and modified LIME helped in understanding how the model makes predictions. Most of the highlighted regions were focused on the lesion area instead of the background. Compared to standard LIME, the modified LIME method pro-vided more focused and meaningful explanations, especially
in difficult cases.

Overall, the results suggest that the proposed framework provides both good classification performance and better in-terpretability, which is important for medical image analysis applications.

V. CONCLUSION

Our work presented a deep learning-based framework for multi-class skin lesion classification using dermoscopic im-ages. The proposed Hybrid IncepX-Ensemble model combines InceptionV3 and Xception to learn different visual features from skin lesion images. By integrating both architectures, the model achieved improved classification performance compared to individual models.

The experimental results showed that the proposed model performs effectively across the seven lesion categories of the HAM10000 dataset. The training and validation behaviour indicated stable learning with limited overfitting. The ROC analysis, confusion matrix, and classification metrics further demonstrated the capability of the model to distinguish be-tween different lesion types with good overall performance.

In addition to classification performance, explainability was also incorporated into the framework using Grad-CAM, LIME, SHAP, and a modified LIME approach. The explainability results showed that the model mainly focuses on clinically rel-evant lesion regions during prediction. Compared to standard LIME, the modified LIME method generated more focused and meaningful explanations, especially in challenging cases. Overall, the proposed framework combines strong classi-fication capability with improved interpretability, making it

VI. FUTURE WORKS

While the proposed Hybrid IncepX-Ensemble model shows strong performance, there are several areas for further im-provement. One important direction is to evaluate the model on larger and more diverse datasets. The current study mainly uses the HAM10000 dataset, while real clinical images may vary due to differences in imaging devices, lighting conditions, and patient characteristics. Testing the model on more varied data can improve its robustness and generalisation.

Another area for future work is the stability of explainability methods. While techniques such as Grad-CAM and LIME provide useful insights, their outputs may vary slightly across different runs. Studying the consistency of these explanations can help improve their reliability, which is important for medical applications.

The model can also be extended for real-world clinical use. It can be integrated into medical systems where doctors can upload images and receive predictions along with visual explanations. Such a system can assist in early diagnosis and support decision-making, especially in remote healthcare settings.

In addition, the proposed framework can be explored for other types of medical images such as MRI, CT, and X-ray. The ensemble design may help in extracting useful features across different imaging modalities. Future work can also include combining image data with other patient information to improve diagnostic accuracy.

REFERENCES

[1] N. Parvin *et al.*, "Early detection of multiclass skin lesions using transfer learning-based incepx-ensemble model," *Under Review*, 2025.
[2] K. M. Hosny, M. A. Kassem, and M. M. Fouad, "Explainable deep inherent learning for multi-classes skin lesion classification," Applied Soft Computing, vol. 159, p. 111624, 2024.
[3] G. Akilandasowmya, G. Nirmaladevi, S. Suganthi, and A. Aishwariya, "Skin cancer diagnosis: Leveraging deep hidden features and ensemble classifiers for early detection and classification," Biomedical Signal Processing and Control, vol. 88, p. 105306, 2024.
[4] H. Ding, Q. Huang, and A. Alkhayyat, "A computer aided system for skin cancer detection based on developed version of the archimedes optimization algorithm," Biomedical Signal Processing and Control, vol. 90, p. 105870, 2024.
[5] M. A. Khan, K. Muhammad, M. Sharif, T. Akram, and S. Kadry, "Intelligent fusion-assisted skin lesion localization and classification for smart healthcare," Neural Computing and Applications, vol. 36, no. 1,
[6] pp. 37–52, 2024.
[7] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, 2017.
[8] G. Cai, Y. Zhu, Y. Wu, X. Jiang, J. Ye, and D. Yang, "A multimodal transformer to fuse images and metadata for skin disease classification," The Visual Computer, vol. 39, no. 7, pp. 2781–2793, 2023.
[9] Y. S. Alsahafi, M. A. Kassem, and K. M. Hosny, "Skin-net: A novel deep residual network for skin lesions classification," Journal of Big Data, vol. 10, no. 1, p. 105, 2023.
[10] N. U. Gilal, S. A. M. Ahmed, J. Schneider, M. Househ, and M. Agus, "Mobile dermatoscopy: Class imbalance management based on blurring augmentation, iterative refining and cost-weighted recall loss," Journal of Imaging and Graphics, vol. 11, no. 2, pp. 161–169, 2023.
[11] A. Alhudhaif, B. Almaslukh, O. O. Aseeri, O. Guler, and K. Polat, "A nonlinear automated multi-class skin lesion detection system using soft attention cnns," Chaos, Solitons and Fractals, vol. 170, p. 113409, 2023.
[12] V.-D. Hoang, X.-T. Vo, and K.-H. Jo, "Categorical weighting domination for imbalanced classification with skin cancer," IEEE Access, vol. 11,
[13] pp. 105 170–105 181, 2023.
[14] K. M. Hosny and M. A. Kassem, "Refined residual deep convolutional network for skin lesion classification," Journal of Digital Imaging, vol. 35, no. 2, pp. 258–280, 2022.
[15] P. Tang, X. Yan, Y. Nan, S. Xiang, S. Krammer, and T. Lasser, "Fusionm4net: A multi-stage multi-modal learning algorithm for skin lesion classification," Medical Image Analysis, vol. 76, p. 102307, 2022.
[16] Y. Wang, Y. Feng, L. Zhang, J. T. Zhou, Y. Liu, and R. S. M. Goh, "Adversarial multimodal fusion with attention mechanism for skin lesion classification," Medical Image Analysis, vol. 81, p. 102535, 2022.
[17] Y. Liu, Z. Wang, Z. Li, J. Li, and T. Li, "Multiscale ensemble of convolutional neural networks for skin lesion classification," IET Image Processing, vol. 15, no. 10, pp. 2309–2318, 2021.
[18] S. S. Chaturvedi, J. V. Tembhurne, and T. Diwan, "A multi-class skin cancer classification using deep convolutional neural networks," Multimedia Tools and Applications, vol. 79, pp. 28 477–28 498, 2020.
[19] P. Tschandl, C. Rosendahl, and H. Kittler, "The ham10000 dataset: A large collection of dermatoscopic images," Scientific Data, vol. 5, 2018.
[20] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," Proceedings of CVPR, 2016.
[21] F. Chollet, "Xception: Deep learning with depthwise separable convo-lutions," Proceedings of CVPR, 2017.
[22] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," Proceedings of CVPR, 2009.