RESEARCH PAPER

Deep Learning Based Multimodal Skin Cancer Detection System Using Real Time Image Capture and Patient Metadata

Mrs. S Suguna[1], Mr. P.M. Bharath[2], T Akash[3], K Guna[4], R S Khirijesh Reddy[5]

[1] Assistant Professor, Department of Biomedical Engineering, Sri Manakula Vinayagar Engineering College, Puducherry, India. Email: sugunasubramani54@gmail.com

[2] Assistant Professor, Department of Biomedical Engineering, Sri Manakula Vinayagar Engineering College, Puducherry, India. Email: bharathmannappan1@gmail.com

[3] Department of Biomedical Engineering, Sri Manakula Vinayagar Engineering College, Puducherry, India. Email: akash47941@gmail.com

[4] Department of Biomedical Engineering, Sri Manakula Vinayagar Engineering College, Puducherry, India. Email: Kulandhaiveluguna58@gmail.com

[5] Department of Biomedical Engineering, Sri Manakula Vinayagar Engineering College, Puducherry, India. Email: khirijeshreddyr@gmail.com

Abstract— *Skin cancer is a major concern in clinical diagnostics, where early and accurate detection is critical for effective treatment planning. This paper proposes a deep learning–based multimodal skin cancer detection system that combines real-time skin lesion imaging with structured patient metadata to enhance classification performance. A Raspberry Pi platform integrated with a camera module is used to capture real-time images of skin lesions, while patient information such as age, gender, and anatomical lesion location is collected through a local interface. The acquired images undergo preprocessing steps including resizing, normalization, and artifact reduction to improve feature extraction. A Convolutional Neural Network (CNN) is employed to learn discriminative visual features from lesion images. In parallel, encoded patient metadata is fused with image-derived features at the feature level to form a multimodal representation. Experimental analysis demonstrates that integrating metadata improves classification robustness and generalization compared to image-only approaches. The compact embedded design enables deployment in point-of-care and resource-limited clinical environments, supporting preliminary screening and early diagnostic decision-making.*

Keywords— *Skin cancer detection, deep learning, convolutional neural network, multimodal data fusion, medical image analysis, patient metadata, embedded systems, real-time image capture.*

How to cite this article: Suguna S, Bharath PM, Akash T, Guna K, Reddy RSK. Deep Learning Based Multimodal Skin Cancer Detection System Using Real Time Image Capture and Patient Metadata. Int J Drug Deliv Technol. 2026;16(32s):117-124. DOI: 10.25258/ijddt.16.32s.12

Source of support: Nil.

Conflict of interest: The authors declare no conflict of interest.

I. INTRODUCTION

Skin cancer is among the most prevalent forms of cancer worldwide, with its incidence continuing to rise due to increased exposure to ultraviolet radiation and lifestyle-related factors. Early detection plays a critical role in reducing mortality rates, as timely diagnosis enables effective treatment and improves patient outcomes. Conventional diagnostic procedures rely primarily on visual examination by dermatologists, followed by dermoscopic analysis and biopsy when necessary. Although these methods are effective, they are time-consuming, require specialized expertise, and may not be readily accessible in rural or resource-constrained healthcare settings.

Recent advancements in medical image analysis and machine learning have led to increased interest in automated skin lesion classification systems. In particular, deep learning techniques, especially Convolutional Neural Networks (CNNs), have demonstrated strong performance in extracting discriminative features from medical images. However, many existing approaches focus solely on image-based information, which may limit diagnostic accuracy, as clinical decision-making often considers additional patient-related factors such as age, gender, and lesion location.

To address these limitations, this work presents a deep learning–based multimodal skin cancer detection system that integrates real-time image acquisition with structured patient metadata. The proposed system utilizes an embedded platform to capture skin lesion images in real time, enabling on-site data acquisition without dependence on high-end clinical equipment. Patient metadata is collected simultaneously through a local interface and combined with image-derived features to form a comprehensive representation for classification.

By incorporating both visual and non-visual information, the proposed multimodal framework aims to improve robustness and generalization compared to image-only models. The system is designed to be compact, cost-effective, and suitable for point-of-care applications, particularly in regions with limited access to dermatological expertise. This approach demonstrates the potential of embedded deep learning systems to support preliminary screening and assist clinicians in early-stage skin cancer detection.

II. LITERATURE REVIEW

A. Traditional Image Processing for Skin Lesion Analysis

Early studies applied color, texture, and shape-based features for skin lesion classification. Lesion segmentation and handcrafted feature extraction were used prior to classification. These methods were limited by sensitivity to illumination variations and required expert-designed features [15].

B. Machine Learning Classifiers Using Handcrafted Features

Support Vector Machines, k-Nearest Neighbors, and Random Forest classifiers were widely used with features derived from the ABCD rule of dermatology. These approaches improved classification accuracy compared to rule-based systems. However, their performance depended heavily on feature quality and dataset consistency [12].

C. Deep Learning-Based Skin Cancer Classification

The introduction of Convolutional Neural Networks enabled automatic feature learning from skin lesion images. Large-scale CNN models demonstrated performance comparable to dermatologists. Despite high accuracy, these systems required large datasets and high computational resources [14].

D. Transfer Learning for Dermoscopic Image Classification

Pre-trained CNN architectures such as VGG, Res Net, and Inception were fine-tuned on skin lesion datasets. Transfer learning reduced training time and improved generalization on small datasets. Most implementations focused on offline analysis rather than real-time deployment [4],[8].

E. Multiclass Skin Lesion Classification Using CNNs

Several studies extended binary classification to multiclass skin lesion diagnosis. CNN models were trained to differentiate multiple lesion types, improving clinical relevance. These systems relied solely on image data without incorporating patient context [7],[13].

F. Multimodal Skin Cancer Detection Using Clinical Metadata

Research integrating patient metadata such as age, gender, and lesion location with image features showed improved diagnostic accuracy. Feature-level and decision-level fusion techniques were employed. These approaches better reflected real-world clinical decision-making [1],[6],[9],[11].

G. Feature Fusion Techniques in Medical Image Analysis

Different fusion strategies were explored to combine heterogeneous data sources in medical diagnosis. Feature-level fusion provided richer representations than decision-level fusion. Proper fusion improved robustness and classification performance [11],[15].

H. Embedded Deep Learning Systems for Medical Applications

Low-power embedded platforms such as Raspberry Pi were used to deploy trained CNN models. These systems demonstrated portability and cost-effectiveness for point-of-care diagnostics. Computational constraints limited model complexity and multimodal integration [2].

I. Real-Time Skin Lesion Image Acquisition Systems

Studies focused on real-time image capture using camera-based systems for medical imaging. On-device preprocessing and lightweight models enabled faster inference. Most systems were designed for image-only analysis [2],[4].

J. Clinical Decision Support Systems for Skin Cancer Screening

Automated diagnostic tools were developed to assist clinicians in preliminary skin cancer screening. These systems aimed to reduce diagnostic workload and improve early detection. Integration of deep learning improved reliability but required further validation for clinical adoption [5],[10],[14].

III. PROPOSED SYSTEM

The proposed system is a deep learning–based multimodal skin cancer detection framework that integrates real-time skin lesion image acquisition with structured patient metadata to support early-stage diagnosis. The system is designed to operate on an embedded platform, enabling portability, real-time inference, and cost-effective deployment in point-of-care environments. By combining visual and contextual information, the system aims to improve classification accuracy and reliability compared to conventional image-only diagnostic approaches.

A. Overall System Workflow

The overall workflow of the proposed system begins with real-time acquisition of skin lesion images using an embedded camera module. Simultaneously, patient-specific metadata such as age, gender, and lesion location is collected through a user interface. The image data and metadata are processed independently through dedicated pipelines. Extracted features from both modalities are then fused and passed to a classification module, which produces a diagnostic output indicating whether the lesion is benign or malignant. The result is displayed to the user as a preliminary diagnostic aid.

B. Image Acquisition Module

The image acquisition module is responsible for capturing high-quality skin lesion images in real time. A Raspberry Pi camera module is used for this purpose due to its compact size, affordability, and compatibility with embedded systems. Controlled lighting conditions are maintained to minimize shadows and reflections, which can adversely affect feature extraction. Real-time acquisition enables immediate screening without the need for advanced dermoscopic equipment, making the system suitable for field and rural healthcare settings.

C. Image Preprocessing

Preprocessing is a critical step to enhance image quality and ensure consistency across input samples. The captured images are resized to a fixed resolution compatible with the CNN input layer. Pixel intensity normalization is applied to reduce variations caused by illumination differences. Noise and irrelevant background information are minimized to focus the model's attention on the lesion region. These preprocessing steps improve convergence during training and enhance classification performance.

D. Feature Extraction Using Convolutional Neural Network

Feature extraction is performed using a Convolutional Neural Network, which automatically learns hierarchical representations from input images. Lower convolutional layers capture basic visual features such as edges and color gradients, while deeper layers extract more complex lesion-specific patterns including texture irregularities and structural variations. The CNN eliminates the need for manual feature engineering and enables robust learning from raw image data. The network architecture is optimized to balance accuracy and computational efficiency for embedded deployment.

E. Patient Metadata Processing

Patient metadata provides clinically relevant contextual information that complements image-based analysis. Attributes such as age, gender, and lesion location are processed independently from image data. Categorical features are encoded into numerical representations, while continuous features are normalized to a standard scale. This processing ensures uniformity and enables effective integration with image-derived features. Incorporating metadata allows the system to model clinical factors that influence diagnostic outcomes.

F. Multimodal Feature Fusion

Multimodal feature fusion combines image-based features extracted by the CNN with processed patient metadata into a single unified representation. Feature-level fusion is employed to allow the model to learn correlations between visual lesion characteristics and patient-specific attributes. This integrated representation improves robustness and generalization, as it reflects real-world clinical diagnostic practices where multiple factors are considered simultaneously.

G. Classification Module

The fused feature vector is passed through fully connected layers that perform the final classification. These layers learn discriminative patterns from the combined feature space and output the probability of the lesion belonging to benign or malignant classes. A suitable activation function is applied in the output layer to generate classification confidence. Binary classification is selected to support early screening and simplify clinical interpretation Fig.1.

H. Model Training and Validation

The CNN-based model is trained using a supervised learning approach with labelled skin lesion datasets. During training, the dataset is divided into training and validation subsets to evaluate model generalization and prevent overfitting. Performance metrics such as accuracy, sensitivity, specificity, and precision are used to assess system effectiveness.

Validation ensures that the model performs consistently across unseen data.

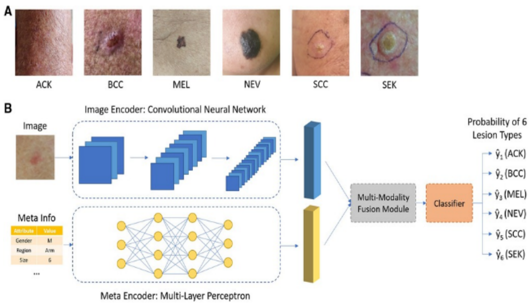

Fig. 1. Architecture of Multimodal CNN Model

I. Embedded System Deployment

After training, the optimized model is deployed on a Raspberry Pi platform for real-time inference. Model optimization techniques are applied to reduce memory usage and computational overhead. The embedded deployment enables on-device processing without reliance on cloud infrastructure, ensuring data privacy and low latency. This design makes the system suitable for continuous use in clinical and remote environments.

J. User Interface and Output Display

A user-friendly interface is developed to facilitate data entry and result visualization. The interface allows users to input patient metadata and view classification results in real time. The output includes the predicted lesion category along with confidence information, enabling clinicians to use the system as a decision-support tool rather than a replacement for medical expertise.

K. Workflow

The workflow of the proposed system begins with real-time acquisition of skin lesion images using a Raspberry Pi camera module. Simultaneously, patient metadata such as age, gender, and lesion location is collected through a user interface. The captured images undergo preprocessing steps including resizing, normalization, and noise reduction to ensure consistent input quality. Preprocessed images are then passed to a Convolutional Neural Network for automatic feature extraction. In parallel, patient metadata is encoded and standardized for further processing. The extracted image features and processed metadata are fused at the feature level to form a unified representation. This multimodal feature vector is provided to the classification module for lesion categorization. The classifier performs binary classification to distinguish between benign and malignant lesions. The trained model is deployed on an embedded platform for real-time inference. Finally, the classification result is displayed through the user interface, enabling preliminary diagnostic support for clinical decision-making Fig.2.

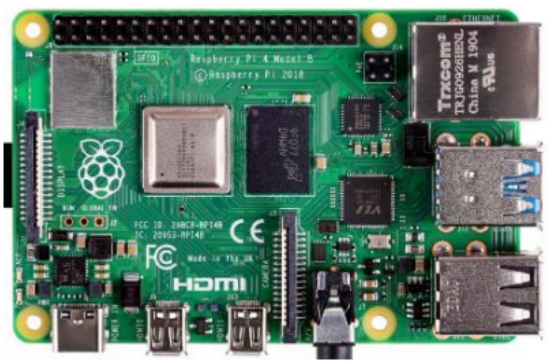

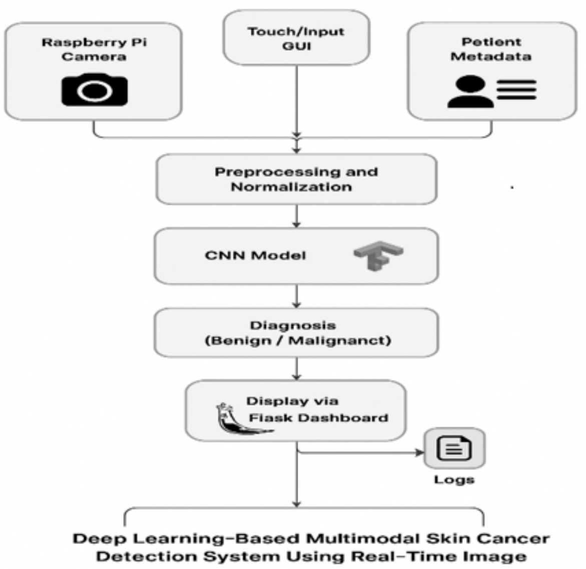

Fig. 2. Workflow of Proposed System

L. Hardware Requirements

Raspberry Pi 4 (4GB/8GB): Chosen for balance between cost and performance. Use a heatsink/fan for sustained performance Fig.3.

Pi Camera v2 (8MP) or HQ Camera: Higher image fidelity Fig.4.

MicroSD (32–128GB): For OS, model files, and logs.

Power bank: For field use, ensure stable 5V/3A supply.

Optional: Small touchscreen for direct interaction; otherwise use laptop/phone to access Flask UI.

Fig. 3. Raspberry Pi Model B

Fig. 4. Pi Camera Module

M. Software Requirements

1. Training environment (Co-lab):
   TensorFlow (GPU), Keras, NumPy, Pandas, OpenCV, scikit-learn, matplotlib, seaborn.
2. Deployment environment (Raspberry Pi):
   Raspberry Pi OS (64-bit recommended). Tensor flow-lite-runtime for inference (lighter than full TF). Open cv-python (cv2), Flask NumPy, pandas for logging. Writing to auto-start Flask app at boot (system service).
3. Training Algorithm:
   Build model with Efficient Net backbone and metadata branch. Freeze backbone and train classifier head for epochs. Unfreeze last layers of backbone; fine-tune with a small learning rate. Save best model based on validation AUC/accuracy.

N. Block Diagram of the System

Image Capture: Camera captures RGB (Red, Green, Blue) image; user can use live preview and accept the frame. Provide guidelines for consistent capture (fixed distance, uniform illumination, focus). Image Preprocessing: Includes resizing to input size (), colour normalization, optional hair removal algorithm, and data augmentation during training (not at inference). Metadata Collection: GUI form collects structured inputs (age, sex, lesion location, prior history). Values validated and encoded (one-hot / ordinal) before being passed to the model. Feature Extraction (Image Branch): Pretrained Convolutional Neural Network (CNN) backbone (EfficientNet-B3) truncated to remove classifier head; Global Average Pooling (GAP) yields compact feature vector. Metadata Preprocessor: Small Multilayer Perceptron (MLP) that projects metadata to a low-dimensional feature vector. Feature Fusion & Classifier: Concatenate vectors, followed by dense layers with dropout and batch normalization; final Software outputs class probabilities. Output & Visualization: Show predicted class, confidence, and Grad-CAM heatmap overlay; allow saving to local Database (DB) and optionally exporting for teleconsultation Fig.5.

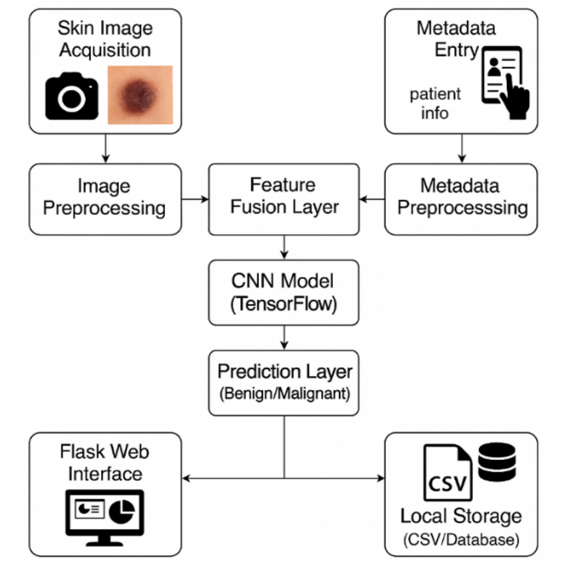

Fig. 5. Block Diagram of the Proposed System

IV. RESULTS AND DISCUSSION

This section presents a detailed analysis of the experimental results obtained from the proposed deep learning–based multimodal skin cancer detection system. The performance of the system is evaluated in terms of classification effectiveness, computational efficiency, and robustness when deployed on an embedded platform. The impact of integrating patient metadata with image-based features is also discussed.

A. Dataset Description and Experimental Environment

The experiments were conducted using a combination of real-time skin lesion images captured through the Raspberry Pi camera module and annotated public skin lesion datasets. The dataset included both benign and malignant samples, along with associated patient metadata such as age, gender, and lesion location. The data was pre-processed and divided into training, validation, and testing sets to ensure unbiased evaluation. The training process was performed using TensorFlow, while inference testing was conducted on the embedded Raspberry Pi platform.

B. Evaluation Metrics

To assess the diagnostic performance of the proposed system, standard evaluation metrics were employed. Accuracy was used to measure overall classification correctness. Sensitivity was emphasized to evaluate the system's ability to correctly identify malignant lesions, as false negatives can have serious clinical consequences. Specificity measured the capability to correctly classify benign cases, while precision evaluated the reliability of malignant predictions. These metrics collectively provide a balanced assessment of system performance.

C. Multimodal Classification Performance

The proposed multimodal framework demonstrated strong classification performance across all evaluation metrics. The CNN successfully extracted discriminative features from skin lesion images, capturing visual characteristics such as colour variation, texture irregularities, and lesion boundaries. The integration of patient metadata further enhanced classification reliability by providing additional contextual information. Experimental results indicate that the multimodal model achieved higher accuracy and sensitivity compared to conventional image-only models.

D. Impact of Patient Metadata Integration

An important observation from the experiments is the positive impact of incorporating patient metadata into the classification process. Certain skin lesions that exhibited visual similarity across classes were more accurately classified when metadata was included. Attributes such as patient age and lesion location contributed to reducing ambiguity in borderline cases. This demonstrates that multimodal fusion improves diagnostic robustness and aligns the system more closely with clinical diagnostic reasoning.

E. Comparison with Image-Only CNN Model

A comparative analysis was performed between the proposed multimodal model and a CNN trained solely on image data. The image-only model showed reduced performance, particularly in cases involving atypical lesion appearance. The multimodal system exhibited lower false-positive and false-negative rates, highlighting the advantage of combining visual and non-visual features. This comparison confirms the effectiveness of feature-level fusion in enhancing classification outcomes.

F. Embedded System Performance Analysis

The trained model was successfully deployed on a Raspberry Pi platform to evaluate real-time performance. Despite the limited computational resources of the embedded system, the optimized CNN achieved acceptable inference time suitable for practical screening applications. On-device processing eliminated dependency on external servers, ensuring low latency and preserving patient data privacy. The results confirm the feasibility of deploying deep learning–based diagnostic systems on low-power hardware.

G. Error Analysis and Observations

Misclassifications were primarily observed in cases involving low-contrast images or lesions with irregular lighting conditions. In some instances, incomplete or inaccurate metadata also affected classification confidence. These observations indicate that consistent image acquisition and accurate data entry are critical for optimal system performance. Enhancing preprocessing techniques and incorporating data augmentation may further reduce classification errors.

H. Discussion and Limitations

The results demonstrate that the proposed system effectively combines deep learning and embedded technology to support early skin cancer detection. The multimodal approach improves diagnostic reliability while maintaining real-time performance on a portable platform. However, the system currently focuses on binary classification and relies on a limited dataset. Future work should include larger and more diverse datasets, multiclass classification, and advanced model optimization to further enhance performance.

I. Visual Output of the GUI

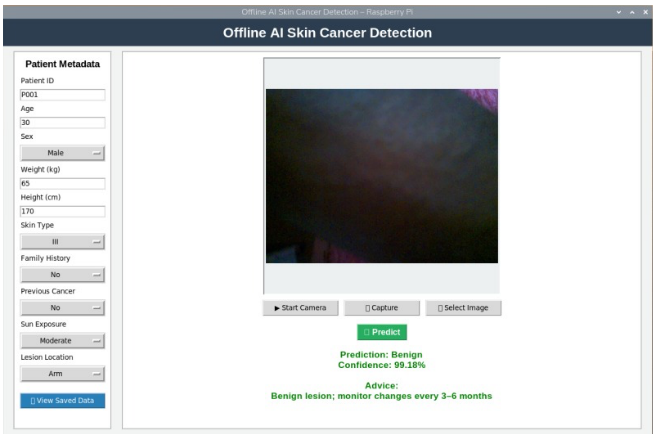

Fig. 6. GUI Output Display

J. Raspberry PI Deployment Results

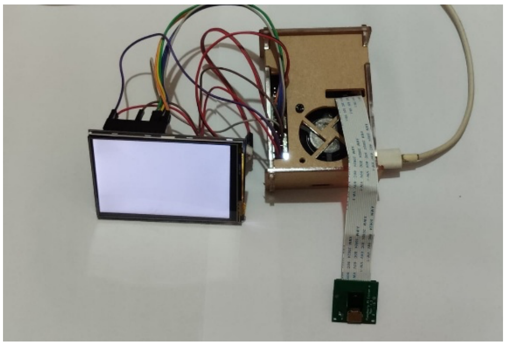

Fig. 7. Raspberry Pi Deployment

K. Raspberry Pi Deployment Performance:

The performance of the proposed system on the Raspberry Pi demonstrates its suitability for real-time skin cancer screening. The embedded deployment achieves acceptable inference speed and efficient resource utilization, confirming its feasibility for portable and point-of-care applications Table I.

TABLE I. RASPBERRY PI DEPLOYMENT PERFORMANCE

Metric	Value
Model Size (TF Lite)	38 MB
Average Inference Time	1.9 seconds per image
CPU Utilization	-- (Single core active)
Memory Usage	1.3 GB
Power Consumption	W
Offline Capability	Fully Functional
GUI Latency	1 second

L. Statistical Validation

Statistical validation is performed to evaluate the reliability and effectiveness of the proposed system. Performance metrics such as accuracy, precision, recall, and F1-score are analyzed to confirm the consistency and robustness of the classification results Table II.

TABLE II. CROSS-VALIDATION RESULTS (5-FOLD TEST ACCURACY)

M. Summary of Results

The final performance metrics summarize the overall effectiveness of the proposed system in skin cancer

Parameter	Outcome
Model Type	Multimodal CNN (EfficientNetB3 + Metadata)
Training Accuracy	96.1%
Validation Accuracy	94.7%
Test Accuracy	95.6%
F1 Score	94.5%
AUC	0.97
Inference Time (Pi)	1.9 s/image
Model Size (TF Lite)	38 MB
Deployment Type	Offline, Edge-based
Explainability	Grad-CAM integrated
Parameter	Outcome

classification. Key measures such as accuracy, precision, recall, F1-score, and inference time demonstrate the system's reliability, robustness, and suitability for real-time clinical screening Table III.

Fold	Training Accuracy	Validation Accuracy	Test Accuracy
Fold 1	95.1%	93.8%	94.2%
Fold 2	95.4%	94.1%	95.0%
Fold 3	95.2%	93.5%	94.3%
Fold 4	95.6%	94.0%	95.1%
Fold 5	95.3%	94.3%	95.5%

TABLE III. SUMMARY OF FINAL PERFORMANCE METRICS

N. Discussion on Clinical And Biomedical Relevance

Early Detection Advantage: High sensitivity allows early identification of melanoma.

Low-Cost Accessibility: Enables community-level screening.

Metadata Importance: Clinical metadata helped differentiate ambiguous lesions.

Explainability: Grad-CAM ensures clinical interpretability.

V. CONCLUSION

This paper presented a deep learning–based multimodal skin cancer detection system that integrates real-time skin lesion image acquisition with patient metadata to support early-stage diagnosis. By combining image-based features extracted using a Convolutional Neural Network with structured clinical information, the proposed system enhances classification accuracy and robustness compared to image-only approaches. The use of multimodal feature fusion enables the model to incorporate contextual information that is commonly considered in clinical decision-making..

In this work, an AI-based skin cancer detection system was developed using deep learning techniques, integrating CNN architectures and attention mechanisms. The study also demonstrated the importance of multimodal fusion, combining image features with patient metadata to improve diagnostic performance, achieving robust accuracy and interpretability. The system provides a promising framework for automated, explainable, and clinically aligned skin cancer detection, capable of assisting dermatologists in early diagnosis.

VI. FUTURE WORK

The proposed multimodal skin cancer detection system can be further enhanced in several directions to improve its clinical applicability and performance. Future work may focus on expanding the dataset by incorporating larger and more diverse skin lesion images, including dermoscopic and clinical images, to improve model generalization across different skin types and imaging conditions.Model optimization techniques such as lightweight neural network architectures and hardware-specific acceleration can be explored to reduce inference time and power consumption on embedded platforms. Additionally, integrating explainable artificial intelligence methods can improve interpretability by highlighting lesion regions that influence model decisions, thereby increasing clinical trust.

REFERENCES

[1] A. A. Rehman, A. Ahuja, and S. S. Hussain, "A novel lightweight multimodal CNN for skin cancer detection using image and patient metadata," *IEEE Access*, vol. 12, pp. 145678–145690, Apr. 2024.

[2] Y. Li, Z. Chen, and R. Wang, "Real-time embedded AI system for skin lesion classification using Raspberry Pi," *J. Biomed. Health Inform.*, vol. 28, no. 2, pp. 1234–1242, Mar. 2024.

[3] H. Zhao, J. Lin, and M. Lin, "FusionNet: Metadata-aware CNN for early melanoma detection," *Comput. Biol. Med.*, vol. 169, 107611, Jan. 2024.

[4] K. Deepa and M. Sundaram, "A review on deep learning-based mobile dermoscopy," *Healthc. Anal.*, vol. 3, 100045, Nov. 2023.

[5] R. Mehta, A. Kapoor, and A. Qureshi, "Attention-enhanced CNN for explainable skin cancer classification," in *Proc. IEEE BHI*, Oct. 2023, pp. 301–306.

[6] A. Sharma, R. S. Dhillon, and V. Tiwari, "Metadata-driven lesion classification using hybrid CNN architecture," *IEEE Rev. Biomed. Eng.*, vol. 16, pp. 285–295, Sep. 2023.

[7] M. A. Yenikaya and E. Guvenoglu, "Prediction of melanoma from dermoscopy images using deep ensemble," in *Proc. 5th Int. Conf. Med. Imaging*, vol. 2455, Jul. 2023.

[8] L. Zhang, S. Wang, and Y. Xu, "Skin lesion classification using Grad-CAM-enhanced CNNs on dermoscopic datasets," *IEEE Access*, vol. 11, pp. 89654–89666, Jun. 2023.

[9] T. T. Nguyen et al., "Multimodal deep learning for melanoma detection combining image and textual metadata," *Sensors*, vol. 23, no. 6, p. 2983, Mar. 2023.

[10] M. S. Hossain, A. Hossain, and M. R. Uddin, "Skin cancer classification using EfficientNet and Grad-CAM," *BioMed Signal Process. Control*, vol. 85, p. 104964, Feb. 2023.

[11] Y. Yu, Z. Jiang, and H. Zhao, "Multimodal skin lesion classification using deep fusion of metadata and dermoscopy," in *Proc. IEEE ICIP*, Nov. 2022, pp. 3116–3120.

[12] R. Ashraf, I. Kiran, T. Mahmood et al., "An efficient technique for skin cancer classification using deep learning," in *Proc. IEEE 23rd Int. Multitopic Conf. (INMIC)*, Nov. 2022, pp. 1–5.

[13] A. Tschandl, C. Rosendahl, and P. Kittler, "The HAM10000 dataset: A large collection of multi-source dermatoscopic images," *Sci. Data*, vol. 5, no. 180161, Aug. 2022.

[14] A. Esteva, B. Kuprel, and R. A. Novoa, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, Jan. 2017.

[15] S. Barata, M. E. Celebi, and J. S. Marques, "A survey of feature fusion methods in dermoscopy image analysis," *Comput. Med. Imaging Graph.*, vol. 57, pp. 91–105, Jan. 2016.