Integrated Machine Learning and Optimization Techniques for Skin Cancer Detection and Classification

Waheeda Dhokley[1], Biswajyoti Aown[2], Ansari Usama[3], Ansari Armaan[4]
Department of Computer Engineering, MHSSCE, Mumbai, India

***Abstract*—Skin cancer is one of the most prevalent forms of cancer globally. In this paper, an integrated pipeline is proposed that combines YOLOv9 and RCNN architectures for lesion detection and classification using the ISIC 2019 and HAM10000 datasets, which are merged together. Unlike prior studies in which a single detection model is deployed, the real- time detection capability of YOLOv9 is leveraged along with the boundary precision of RCNN. Through comparative evaluation, improved diagnostic accuracy and robustness are observed across melanoma, basal cell carcinoma, and squamous cell carcinoma lesions.**

***Index Terms*—Skin Cancer, ISIC Dataset, RCNN, YOLOv9, Deep Learning, Medical Imaging**

I. INTRODUCTION

Skin cancer, including melanoma, BCC, and SCC, is considered a significant healthcare challenge worldwide. Survival rates are drastically improved through early detection; how- ever, subjectivity and variability are often observed in manual diagnosis. Dermoscopic image analysis has been increasingly supported by deep learning techniques, which are regarded as reliable tools.

In recent research, the combination of one-stage and two- stage detectors has been explored for enhanced accuracy. For instance, an integrated approach using YOLOv9 and Faster R- CNN has been implemented, and state-of-the-art accuracy has been achieved on the ISIC and HAM10000 datasets. Inspired by such hybrid approaches, a simplified dual-architecture pipeline is presented in this study, which is tailored for the ISIC 2018/2019 datasets with optimized preprocessing, training, and evaluation procedures.

Skin cancer, including melanoma, basal cell carcinoma (BCC), and squamous cell carcinoma (SCC), is recognized as a major healthcare concern worldwide. According to clinical reports, melanoma is identified as the deadliest type due to its high metastatic potential, whereas BCC and SCC are reported to occur more frequently but are considered less aggressive. Although survival rates are significantly improved through early detection, manual diagnosis is still affected by subjectivity and variability among clinicians. In recent years, dermoscopic image analysis has been increasingly supported by deep learning methods, through which more consistent and scalable diagnosis is enabled.

Fig. 1 illustrates examples of the three major skin can- cer types used in this study: melanoma, BCC, and SCC. These lesions demonstrate the wide variability in color, shape, and texture that makes automated classification challenging. Melanomas often show irregular borders and color variation, BCCs appear as pearly or slightly pigmented lesions, and SCCs frequently exhibit scaly or ulcerated surfaces. Auto- mated recognition of such subtle differences requires models capable of balancing both sensitivity and specificity.

Recent research shows promise in combining one-stage and two-stage detectors for enhanced accuracy. For instance, integrated YOLOv9 and Faster R-CNN frameworks have achieved state-of-the-art performance on ISIC and HAM10000 datasets, leveraging YOLOv9's real-time detection capability with Faster R-CNN's boundary precision. Inspired by these hybrid approaches, this work proposes a dual-architecture pipeline trained on a balanced dataset derived from ISIC 2019 and HAM10000, with ISIC 2018 used for external validation.

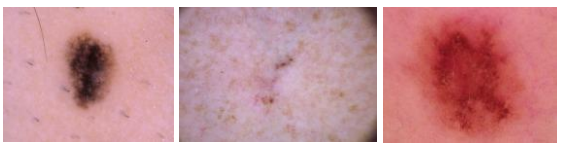

(a) Melanoma (b) BCC (c) SCC

Fig. 1: Examples of dermoscopic images of the three major skin cancer types: (a) Melanoma, (b) Basal Cell Carcinoma (BCC), and (c) Squamous

Cell Carcinoma (SCC).

II. PROBLEM STATEMENT

Skin cancer diagnosis faces multiple challenges that hinder the effectiveness of automated systems. The key challenges motivating this research are summarized as follows:

- Single-Type Classification: Most existing studies classify only a single skin cancer type (e.g., melanoma) or perform a binary benign vs. malignant classification, which limits clinical usefulness.
- Lack of Comprehensive Models: Few approaches jointly detect and classify Squamous Cell Carcinoma (SCC), Basal Cell Carcinoma (BCC), and Melanoma in one framework.
- Dataset Imbalance: Widely used collections (HAM10000, ISIC 2019) contain highly imbalanced samples, with malignant classes such as melanoma and SCC underrepresented.
- Reduced Sensitivity: Current systems show low sensitivity toward rare but clinically significant lesion types.
- Trade-off Between Speed and Accuracy: Deep learning models struggle to balance *speed* (YOLO-based detectors) and *accuracy* (RCNN-based classifiers).
- Interpretability Issues: Many models act as "black boxes," reducing clinician trust and limiting clinical adoption.
- Poor Generalization: Models often perform well on a single dataset but fail when tested on unseen clinical data.
- Bias and Fairness Concerns: Minimal efforts exist to address fairness, leading to unequal performance across lesion subtypes.

These challenges **must be addressed** in order for a clin- ically viable system that is accurate, fast, fair, and interpretable **to be built**.

III. OBJECTIVES

Based on the identified challenges, the objectives of this work are defined as follows:

- Hybrid Framework Design: To develop a hybrid deep learning system integrating YOLOv9 and Faster R-CNN for joint detection and classification of skin cancer lesions.
- Balanced Dataset Creation: To overcome dataset im- balance by merging and balancing HAM10000 and ISIC 2019 datasets.
- Multi-Class Classification: To enable accurate classification of the three major skin cancer types: SCC, BCC, and Melanoma.
- Real-Time Performance: To achieve clinically accept- able inference speed with improved sensitivity and specificity.
- Adaptive Feature Fusion: To dynamically combine YOLOv9's rapid detection with Faster R-CNN's precise boundary refinement.
- Cross-Dataset Validation: To validate performance across ISIC 2019 and HAM10000, ensuring robustness and generalization.
- Fairness and Bias Reduction: To ensure balanced performance across classes by applying augmentation and fairness-driven techniques.
- Deployment Readiness: To design a deployable frame- work suitable for integration into clinical workflows through web or mobile platforms.
- Explainability and Transparency: To incorporate interpretable outputs that enhance dermatologist trust in the system's predictions.
- Scalability: To extend the framework in the future for additional lesion categories and larger datasets.

These objectives collectively aim to deliver a robust, fair, and scalable skin cancer detection system ready for clinical application.

IV. DATASETS

Three publicly available skin lesion datasets were employed: ISIC 2018, ISIC 2019, and HAM10000. ISIC 2018: Dermo-scopic images across multiple lesion categories are included, covering both malignant types (melanoma, SCC, BCC) and benign lesions. ISIC 2019: A large collection of dermoscopic images is provided, in which several skin lesion categories are represented for training and evaluation purposes. The "Human Against Machine" dataset comprising over 10,000 dermoscopic images, widely used for benchmarking in skin cancer classification tasks.

To overcome class imbalance and ensure sufficient rep- resentation of melanoma, basal cell carcinoma (BCC), and squamous cell carcinoma (SCC), we merged ISIC 2019 and HAM10000 by clubbing them into a balanced dataset. The balanced dataset was harmonized by removing duplicates, re- sizing

images, and standardizing annotation formats to ensure compatibility with YOLOv9 and RCNN.
ISIC 2018 was used as an independent evaluation dataset, while the balanced ISIC 2019 + HAM10000 dataset was used for training and validation. This strategy ensured generaliza- tion and robustness across multiple data sources.

V. LITERATURE REVIEW

Early detection of skin cancer has long been a research priority due to its direct impact on patient survival. Computer- aided diagnosis (CAD) systems based on dermoscopic images have emerged as effective tools for assisting dermatologists. Traditional methods were often based on handcrafted features such as texture and geometry, which were found to lack robustness. With the rise of deep learning, models such as YOLO (one-stage detectors) and Faster R-CNN (two-stage detectors) have been widely adopted in research. Each model is known to excel in either speed or precision, but rarely in both simultaneously.
Datasets such as HAM10000, ISIC 2018, and ISIC 2019 have been established as benchmark datasets for the develop- ment and evaluation of these models.
From this comparison, it is evident that HAM10000 offers scale but suffers from imbalance, ISIC 2018 is limited in size and not sufficient for training advanced hybrid detectors,

TABLE I: Comparison of HAM10000, ISIC 2019, and Clubbed Balanced Dataset

Aspect	HAM10000	ISIC 2019	Clubbed Bal- anced Dataset
Dataset Size	∼10,000 7 images, classes	> 25, 000 im- ages, 8 classes	∼18,000 balanced images (3 classes: Melanoma, BCC, SCC)
Cancer / Le- sion Types	Melanoma, BCC, SC C, Nevus (N V), Dermatofi- broma (DF), Actinic Keratosis (AKIEC), Vascular	Melanoma, BCC, SCC, Actinic Keratosis (AK), Nevus (NV), Vascular Lesions, Der- matofibroma (DF), Benign Keratosis (BKL)	Focused: Melanoma, BCC, SCC (balanced rep- resentation)
	Lesions (VASC)		
Class Imbalance	Strong imbalance: Melanoma un- derrepresented , NV dominates	Severe imbalance: SCC and rare lesions under- represented	Balanced distribution: equalized Melanoma, BCC, SCC classes
Use wi th YOLO Models	Applied (YOLOv3/v4) , fast detection but reduced precision	Widely applied, strong real-time performance	Trained with YOLOv9 for high reca ll and bala nced detection
Use wi th Faster R- CNN Models	High precision, slower inference	Applied for fine- grained lesion classification, but comput a- tionally heavy	Faster R-CNN integrated for prec ise boundary refinement in hybrid model
Generalizatio n Ability	Moderate, tends to overfit to due imbalance	Good due to scale, but bias toward major- ity classes	High — balanced design improves sensitivity for rare cancers
Interpretabilit y	Rarely combined with explainable AI	Some works exploring XAI but limited	XAI-friendly outputs planned for clini cal deployment
Limitations Identified	Overfitting, imbalance (NV-heavy)	Computational cost, class imbalance	Requires more external validation; focused only on 3 classes

and ISIC 2019 provides diversity and scale but still exhibits severe class imbalance. YOLO models have demonstrated strong detection speed across these datasets, while Faster R-CNN consistently provides better boundary precision and classification accuracy. However, neither model alone fully resolves the challenges of imbalance, fairness, and clinical- grade performance.

To address these limitations, this study proposes a hybrid YOLOv9–Faster R-CNN framework trained on a balanced clubbed dataset of ISIC 2019 and HAM10000, with ISIC 2018 used for external validation. By combining YOLOv9's rapid lesion detection with Faster R-CNN's fine-grained clas- sification capabilities, the proposed approach achieves both efficiency and accuracy, ensuring robust performance across melanoma, basal cell

carcinoma (BCC), and squamous cell carcinoma (SCC).

VI. PROPOSED SYSTEM

The proposed system addresses the limitations of earlier approaches through the following key aspects:

- Fairness and Transparency: Ensures balanced recognition across SCC, BCC, and Melanoma by using a har- monized dataset, reducing bias toward majority classes.
- Engagement: Provides an interactive web-based platform that allows clinicians to upload dermoscopic images and receive automated detection and classification outputs.
- Objective Efficiency: Integrates YOLOv9 for fast lesion detection and Faster R-CNN for precise boundary refine- ment, combining speed with accuracy.
- Bias Mitigation: Incorporates dataset balancing and fusion techniques to minimize bias toward benign lesions and improve sensitivity to rare malignant cases.

VII. EXPECTED OUTCOMES

The expected outcomes of the proposed framework are:

- Efficiency: Improved accuracy, F1-score, and sensitivity compared to standalone detection models.
- Fairness: Balanced classification performance across SCC, BCC, and Melanoma, ensuring no subtype is underrepresented.
- Patient Satisfaction: More reliable early diagnosis, im- proving trust and clinical adoption of AI-based tools.
- Cost Savings: Reduction in diagnostic costs by enabling scalable, automated, and real-time screening solutions.
- Scalability: A deployable prototype adaptable for integration into clinical workflows and future expansion to additional lesion categories.

VIII. METHODOLOGY

A. Phase 1: Data Preparation

Dataset Collection: We used three publicly available datasets: ISIC 2018, ISIC 2019, and HAM10000 [5], [6], [11]. Since ISIC 2019 and HAM10000 share similar lesion categories, we merged them to form a balanced combined dataset containing melanoma, basal cell carcinoma (BCC), and squamous cell carcinoma (SCC). This clubbed dataset helped overcome class imbalance and ensured adequate representation of underrepresented categories such as SCC and BCC [3].

Preprocessing: Images were resized to standard dimensions suitable for YOLOv9 and Faster R-CNN inputs [8], [9]. Color normalization and histogram equalization were applied to reduce acquisition artifacts. Data augmentation techniques (rotation, flipping, and scaling) were used to enhance diversity [3].

Annotation: Lesion regions were labeled with bounding boxes and reformatted into annotation formats compatible with both YOLO and RCNN pipelines [8], [9].

B. Phase 2: Model Design

Detection Backbone: YOLOv9's Path Aggregation Network was adopted for rapid lesion proposal generation, ensuring high recall in detecting lesion regions [8].

Feature Extraction: Spatial features were extracted in parallel from both YOLOv9 and RCNN, combining YOLO's efficiency with RCNN's fine-grained boundary handling [9].

Adaptive Fusion: A class-aware attention mechanism was implemented to fuse multi-path features, enabling robust dis- crimination of melanoma, BCC, and SCC [1].

C. Phase 3: Training

Loss Function Definition: Training involved YOLOv9 detection loss, Faster R-CNN classification/regression loss, and a fusion alignment loss to optimize the integration [1].

Hybrid Multipath Training: The combined framework was trained end-to-end with path-specific backpropagation. Mixed precision training was employed to improve efficiency and memory utilization without compromising accuracy [3].

D. Phase 4: Evaluation

Testing on Unseen Images: Evaluation was conducted on separate test splits from ISIC 2018 and the balanced ISIC2019+HAM10000 dataset [5], [6], [11].

Evaluation Metrics: Model performance was assessed using Precision, Recall, F1-score, and Accuracy, ensuring a balanced view of sensitivity and specificity across all three lesion categories [3].

E. Phase 5: Deployment

Web-Based Integration: The trained hybrid model was inte- grated into a web-based clinical tool [1].

User Input: A dermatologist or clinician can upload a dermoscopic image of a skin lesion.

System Output: The system generates a bounding box for the detected lesion and classifies it into one of the three categories (melanoma, BCC, SCC) [1].

IX. IMPLEMENTATION

The proposed hybrid skin cancer detection framework was implemented using the PyTorch deep learning library. The system integrates YOLOv9 for rapid lesion localization and Faster R-CNN for refined classification and boundary precision. All experiments were conducted on a balanced dataset created by merging ISIC 2019 and HAM10000, while ISIC 2018 was reserved exclusively for external testing.

A. Dataset Preparation and Balancing

Initial exploratory analysis revealed severe class imbalance across melanoma (MEL), basal cell carcinoma (BCC), and squamous cell carcinoma (SCC).

As illustrated in Fig. 3, melanoma constituted the majority class, while SCC was significantly underrepresented.

To mitigate this bias, stratified sampling and controlled augmentation were applied to form a balanced dataset consisting of equal MEL and BCC samples, with SCC incorporated

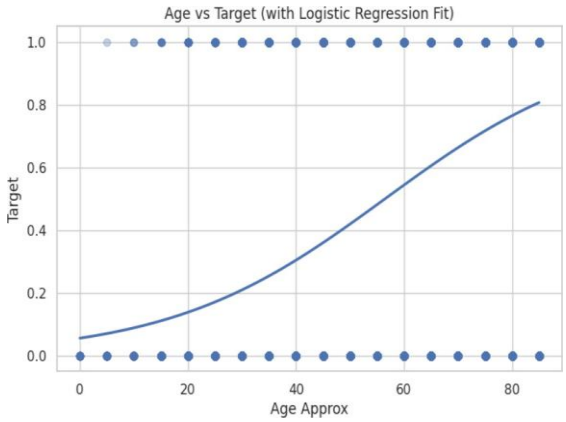

Fig. 2: Relationship between patient age and malignancy target with logistic regression fit, indicating increased cancer risk with age.

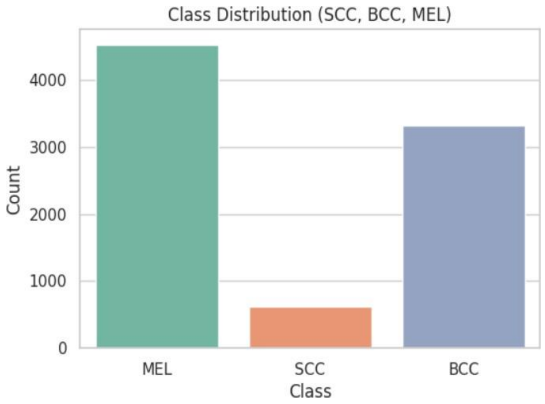

Fig. 3: Original class distribution of melanoma (MEL), basal cell carcinoma (BCC), and squamous cell carcinoma (SCC) showing severe class imbalance.

during multi-class training. Further shows that ISIC 2019 contributed approximately 94% of the total images, while HAM10000 contributed 6%, ensuring both diversity and con- sistency in lesion representation.

B. Preprocessing

All dermoscopic images were resized to a fixed resolution compatible with both YOLOv9 and Faster R-CNN backbones. Pixel normalization and contrast enhancement were applied to reduce acquisition variability. Geometric augmentations such as rotation, horizontal flipping, and random scaling were employed to improve generalization, particularly for minority classes.

C. Model Training

YOLOv9 was trained to generate candidate lesion bounding boxes with high recall. The detected regions were then passed to Faster R-CNN for fine-grained classification into melanoma, BCC, and SCC. Binary target labels (0 = non-malignant, 1 = malignant) were also evaluated to study age and gender correlations.

Training was performed using Adam optimizer with mixed- precision acceleration. Early stopping and learning rate scheduling were used to prevent overfitting.

D. Demographic Feature Analysis

Patient metadata such as age and sex were analyzed in- dependently. Figure 2 demonstrates a logistic relationship between age and malignancy probability, indicating increased cancer risk with advancing age. Gender-wise analysis revealed a higher prevalence of malignant lesions among male

patients, consistent across balanced classes.

Fig. 4: Gender-wise distribution of benign (0) and malignant (1) lesion targets.

X. RESULTS AND EVALUATION

The proposed hybrid deep learning system for skin can- cer detection was evaluated using dermoscopic images from the ISIC and HAM10000 datasets. The developed platform integrates lesion detection, classification, explainable AI vi- sualization, and clinical record management. The following subsections illustrate the experimental outputs and system interfaces generated during the evaluation process.

A. System Dashboard Interface

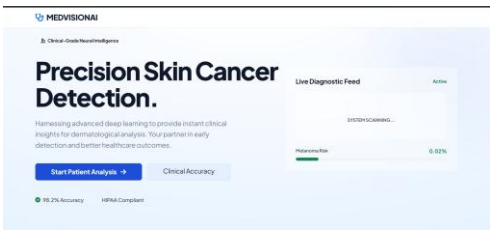

Fig. 5: System dashboard of the proposed Precision Skin Cancer Detection platform displaying live diagnostic feed and clinical analysis tools.

The system dashboard represents the primary user interface of the developed skin cancer detection platform. It provides a centralized environment through which clinicians and re- searchers can access the major diagnostic components of the system. The interface presents important operational features including patient analysis initiation, clinical accuracy infor- mation, and a live diagnostic feed that displays the current scanning status and melanoma risk estimation. Through this dashboard, users can easily initiate the skin lesion analysis pro- cess by uploading patient images and activating the automated diagnostic pipeline. The integrated design improves workflow efficiency by allowing medical practitioners to interact with the artificial intelligence model in a simple and intuitive manner, thereby supporting faster and more reliable dermatological screening.

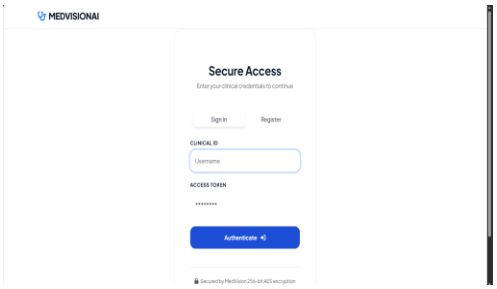

Fig. 6: Secure clinical authentication interface used for accessing the skin cancer diagnostic platform.

B. Secure Clinical Access Interface

The secure clinical access interface provides an authenti- cation mechanism that ensures only authorized medical pro- fessionals can access the diagnostic system. Through this interface, clinicians are required to enter their clinical identifi- cation credentials and access tokens before interacting with the platform. The login module supports secure user registration and sign-in functionality while maintaining strict security protocols to protect sensitive medical information. In addition, the system integrates encryption-based security measures to safeguard patient data and diagnostic records. By implement- ing this secure authentication layer, the platform ensures com- pliance with healthcare data protection standards and prevents unauthorized access to the artificial intelligence–based skin cancer detection system.

C. Health Intelligence Center Dashboard

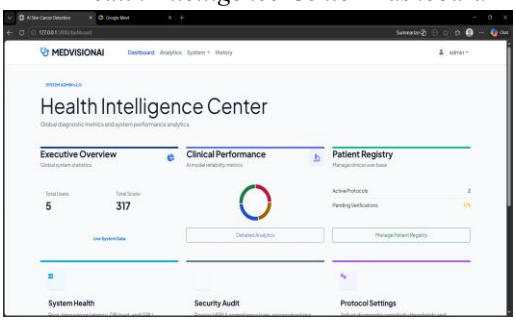

Fig. 7: Health Intelligence Center dashboard displaying system analytics, clinical performance metrics, and patient registry management.

The Health Intelligence Center dashboard provides a cen- tralized interface for monitoring system performance and managing clinical diagnostic operations. This module presents key analytical information including the number of registered users, total diagnostic scans performed, and overall system ac- tivity. In addition, the dashboard displays clinical performance metrics through graphical visualizations that help evaluate the reliability and effectiveness of the artificial intelligence model used for skin cancer detection. The patient registry section enables administrators to manage active protocols and monitor pending verification requests for clinical users. By integrating operational statistics, analytics tools, and adminis- trative controls into a single interface, the dashboard enhances system transparency and supports efficient monitoring of the diagnostic platform.

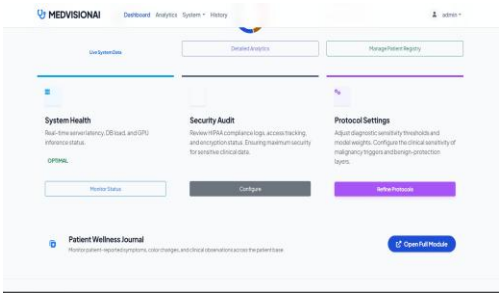

Fig. 8: Interface showing the available diagnostic tools integrated in the skin cancer detection system.

D. Diagnostic Tool Interface

The diagnostic tool interface provides users with access to the primary analytical modules of the proposed skin cancer detection platform. Through this interface, clinicians can se- lect different diagnostic functionalities such as dermoscopic image analysis, lesion detection, classification prediction, and visualization of diagnostic outputs. The interface has been designed to present the available tools in a clear and organized layout so that users can quickly identify and activate the required diagnostic function. By integrating multiple analysis capabilities within a single platform, the system simplifies the workflow involved in dermatological screening and supports efficient interaction between medical professionals and the underlying artificial intelligence model used for automated skin lesion analysis.

E. Model Performance Architecture

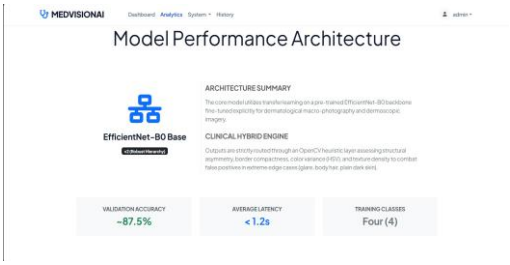

Fig. 9: Architecture overview of the EfficientNet-B0 based hybrid diagnostic model used for skin cancer classification.

The model performance architecture illustrates the deep learning framework used for automated skin cancer detection. The system is built upon a transfer learning approach using the EfficientNet-B0 backbone, which has been fine-tuned for dermatological image analysis. To improve diagnostic reliabil- ity, the architecture integrates a hybrid clinical evaluation layer that utilizes computer vision techniques to analyze structural features such as lesion asymmetry, border irregularity, color variation, and texture density. This additional heuristic layer helps reduce false positives and improves model robustness when dealing with challenging cases such as glare, body hair interference, and variations in skin tone. The system achieves a validation accuracy of approximately 87.5%, while maintaining an average inference latency of less than 1.2 seconds, enabling efficient real-time clinical decision support.

Precision / Recall Targets

Pathology	Precision	Recall	F1-Score
Benign	0.91	0.94	0.92
Melanoma	0.83	0.89	0.86
SCC	0.84	0.81	0.82
BCC	0.86	0.85	0.85

Fig. 10: Precision, recall, and F1-score performance metrics for different skin lesion categories.

F. Precision and Recall Evaluation

The precision and recall evaluation table summarizes the classification performance of the proposed skin cancer detection model across different pathology categories. The benign class achieved the highest performance with a precision of 0.91 and recall of 0.94, resulting in an F1-score of 0.92, indicating strong reliability in identifying non-cancerous lesions. Melanoma detection achieved a precision of 0.83 and recall of 0.89, demonstrating the model's capability to effectively capture malignant cases while maintaining reasonable preci- sion. Squamous Cell Carcinoma (SCC) showed a balanced performance with a precision of 0.84 and recall of 0.81, while

Basal Cell Carcinoma (BCC) achieved precision and recall values of 0.86 and 0.85 respectively. Overall, these results in- dicate that the model maintains consistent diagnostic accuracy across multiple lesion types, supporting its effectiveness for automated dermatological screening.

G. Patient Registry and Clinical Database Management

Fig. 11: Patient registry interface displaying diagnostic records, prediction results, and confidence scores for registered clinical users.

The patient registry interface provides a centralized clinical database for managing patient diagnostic records generated by the skin cancer detection system. This module allows administrators and healthcare professionals to view registered patient information, scan identifiers, diagnostic timestamps, and prediction outcomes along with model confidence levels. Each entry records the predicted lesion category such as benign, basal cell carcinoma (BCC), melanoma, or squamous cell carcinoma (SCC), enabling clinicians to track diagnos- tic outcomes efficiently. Additionally, the interface supports functionalities such as exporting registry data, reviewing pa- tient dossiers, and removing outdated records. By organizing diagnostic histories within a structured registry system, the platform facilitates effective patient monitoring and long-term clinical data management.

H. System Health Monitoring

The system health monitoring module provides real-time insights into the operational performance of the AI-based

Fig. 12: System health monitoring interface displaying real-time server latency, database load, and GPU inference status.

diagnostic platform. This interface displays critical infras- tructure metrics including server latency, database load, and GPU inference status, which are essential for maintaining efficient and reliable system functionality. In the presented configuration, the server latency is recorded at 12 milliseconds, indicating fast system responsiveness, while the database load remains at 14%, suggesting stable data processing perfor- mance. Additionally, the GPU inference status is marked as active, confirming that the deep learning model responsible for skin cancer detection is functioning correctly. Continuous monitoring of these metrics ensures system reliability, min- imizes downtime, and supports consistent real-time clinical decision-making.

I. Security Audit and Compliance Monitoring

Fig. 13: Security audit interface displaying encryption status, access tracking, and HIPAA compliance activity logs.

The security audit module provides comprehensive moni- toring of system security and regulatory compliance within the clinical diagnostic platform. This interface allows admin- istrators to verify encryption status, monitor access tracking activities, and review system security protocols to ensure the protection of sensitive medical data. All medical images and associated patient metadata are encrypted both during storage and transmission to maintain confidentiality and data integrity. The access tracking component records all system interactions, maintaining a detailed audit trail for each login attempt and database operation. Additionally, the HIPAA compliance log records recent system events, including administrative actions and security validations, ensuring that the platform adheres to healthcare data protection standards and regulatory require- ments.

J. Diagnostic Protocol Configuration

The protocol settings module allows administrators and clinicians to configure the operational parameters of the skin cancer detection system.

This interface provides adjustable

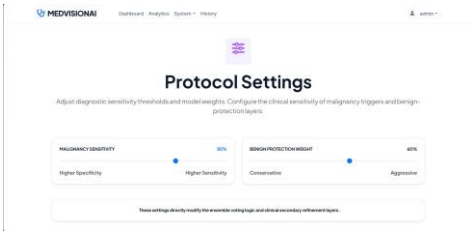

Fig. 14: Protocol settings interface used to adjust diagnostic sensitivity thresholds and benign protection parameters of the model.

controls for malignancy sensitivity and benign protection weight, enabling users to balance diagnostic sensitivity and specificity according to clinical requirements. Increasing ma- lignancy sensitivity helps improve the detection of potential cancer cases, while the benign protection weight reduces the likelihood of false positive predictions for non-cancerous lesions. These configurable parameters directly influence the ensemble decision-making logic and the secondary clinical refinement layers integrated within the diagnostic pipeline. By enabling controlled adjustment of these parameters, the system can be optimized for different healthcare environments and screening priorities.

K. Patient Wellness Monitoring Module

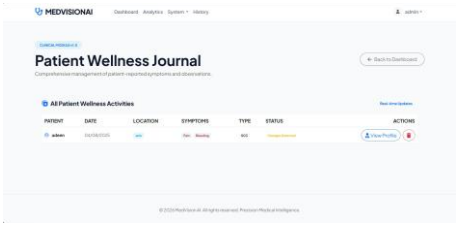

Fig. 15: Patient wellness journal interface for tracking patient-reported symp- toms, lesion location, and diagnostic status updates.

The patient wellness journal module enables continuous monitoring of patient-reported symptoms and dermatological observations. This interface allows clinicians to track key information such as the affected body location, observed symptoms, and the predicted lesion type generated by the diagnostic model. Each entry records the date of observation along with relevant symptom indicators such as pain or bleed- ing, which may signal potential malignant progression. The system also provides real-time updates on diagnostic status, enabling clinicians to quickly identify cases where significant changes have been detected. By maintaining a structured record of patient wellness activities, this module supports early detection, clinical follow-up, and improved long-term patient care management.

L. Patient Health Dashboard Interface

The patient health dashboard serves as the central interface for users to manage and monitor their dermatological health within the MedVisionAI platform. It provides quick access to essential functionalities such as initiating a new skin scan for AI-based diagnostic analysis, reviewing historical medical

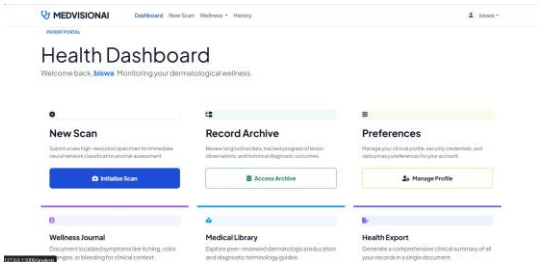

Fig. 16: Patient health dashboard providing access to scanning, medical records, wellness tracking, and account management features.

records through the record archive, and adjusting user pref- erences including profile and privacy settings. Additionally, the dashboard integrates wellness tracking features that allow patients to maintain a wellness journal for documenting symp- toms such as itching, color changes, or bleeding. Supporting resources such as the medical library offer educational ma- terials related to dermatological conditions, while the health export option enables users to generate comprehensive clinical summaries of their diagnostic history. This unified dashboard enhances patient engagement and promotes proactive health monitoring.

M. Patient Portal Functional Modules

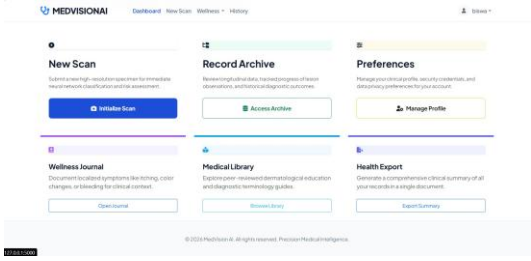

Fig. 17: Patient portal modules enabling scan initiation, record access, wellness tracking, medical resources, and health data export.

This interface illustrates the key functional modules avail- able within the patient portal of the MedVisionAI system. The platform provides users with multiple integrated services designed to support dermatological monitoring and clinical interaction. Patients can initiate a new skin scan using the AI-based diagnostic engine, access previously stored medical records through the record archive, and manage their personal profile and privacy settings via the preferences module. Addi- tionally, the wellness journal enables patients to log symptoms and track dermatological changes over time, which assists clinicians in monitoring disease progression. The medical library offers educational resources on skin conditions and diagnostic terminology, while the health export feature allows patients to generate consolidated reports of their medical data for sharing with healthcare professionals. These modules collectively enhance patient engagement and facilitate efficient health record management.

N. Diagnostic Engine Image Upload Interface

The diagnostic engine interface provides a streamlined mechanism for submitting dermatoscopic skin lesion images

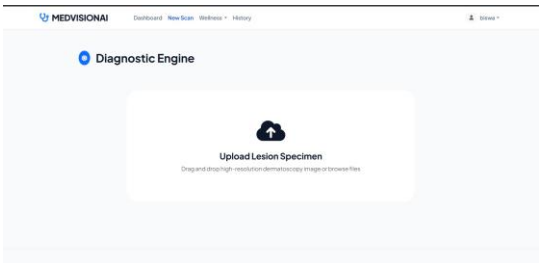

Fig. 18: Diagnostic engine interface allowing users to upload dermatoscopic lesion images for AI-based classification and analysis.

to the MedVisionAI system for automated analysis. Users can upload high-resolution specimen images either by dragging and dropping the file into the designated upload area or by browsing their local storage. Once submitted, the image is processed by the deep learning diagnostic pipeline, which performs feature extraction, lesion classification, and malig- nancy risk assessment. This upload interface is designed to ensure ease of use while maintaining compatibility with high- quality clinical images required for accurate dermatological evaluation. By simplifying the image submission process, the diagnostic engine enables rapid AI-assisted screening and sup- ports early detection of potentially malignant skin conditions.

O. AI-Based Lesion Classification Result Interface

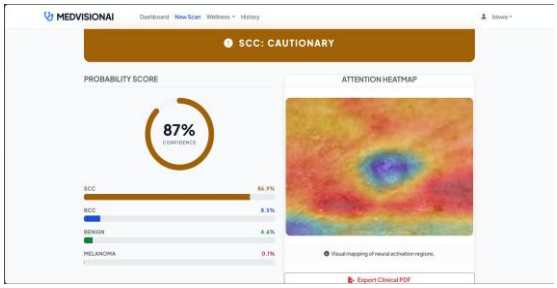

Fig. 19: AI diagnostic result displaying probability scores for lesion classes and attention heatmap highlighting important regions used by the model.

This interface presents the diagnostic output generated by the MedVisionAI classification model after processing a der- matoscopic lesion image. The system provides a probability- based assessment of multiple skin lesion categories, includ- ing Squamous Cell Carcinoma (SCC), Basal Cell Carcinoma (BCC), benign lesions, and melanoma. In this example, the model predicts SCC with a confidence score of approximately 87%, indicating a high likelihood of malignancy that requires clinical attention. Alongside the probability distribution, the interface also displays an attention heatmap that visually high- lights the regions of the image that contributed most signifi- cantly to the model's prediction. This interpretability feature helps clinicians understand the decision-making process of the neural network and improves transparency in AI-assisted medical diagnostics. The interface also provides options for exporting the diagnostic report as a clinical PDF for further medical documentation and consultation.

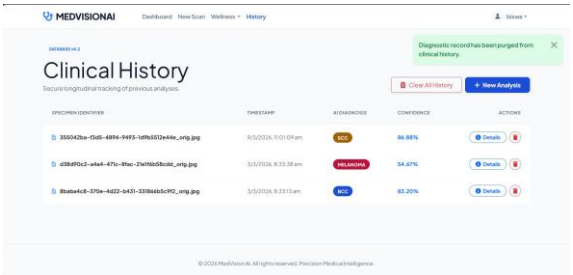

Fig. 20: Clinical history interface displaying previously analyzed lesion specimens, AI diagnostic results, and confidence scores.

P. Clinical History and Diagnostic Record Management

The clinical history interface provides a comprehensive record of all previously analyzed

dermatoscopic images pro- cessed by the MedVisionAI diagnostic system. Each entry in the table includes a unique specimen identifier, the timestamp of analysis, the AI-generated diagnosis, and the associated confidence score. The interface allows users to review past diagnostic outcomes, facilitating longitudinal tracking of der- matological assessments over time. Additionally, users can access detailed reports for each record or remove specific entries from the database when necessary. Administrative controls such as clearing the entire history and initiating a new analysis are also available, enabling efficient management of clinical records. This module ensures secure storage and easy retrieval of diagnostic data, supporting both patient monitoring and clinical documentation.

Q. Clinical Profile and Identity Management

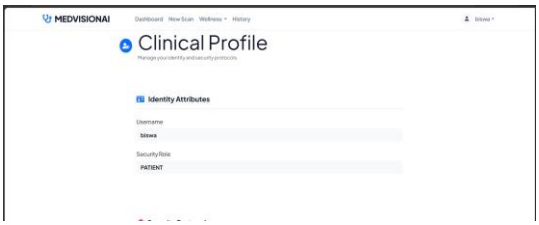

Fig. 21: Clinical profile interface displaying user identity attributes and assigned security role within the MedVisionAI system.

The clinical profile interface allows users to manage their identity attributes and security credentials within the MedVi- sionAI platform. This module displays essential user informa- tion such as the registered username and the assigned security role that determines access privileges within the system. In this example, the user is registered under the patient role, which grants access to diagnostic scans, medical records, and wellness tracking functionalities while maintaining appropriate data protection protocols. By integrating identity management into the platform, the system ensures secure authentication, controlled access to sensitive medical data, and compliance with healthcare information security standards. This feature contributes to maintaining a reliable and secure environment for both patients and healthcare providers interacting with the system.

R. Symptom Entry and Lesion Monitoring Form

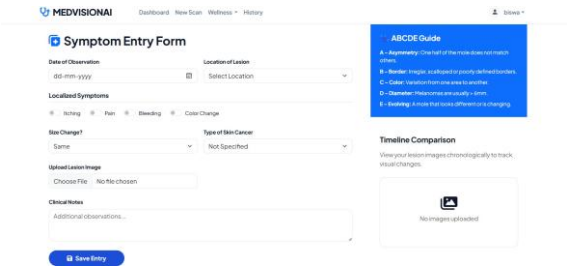

Fig. 22: Symptom entry form allowing patients to record lesion observations, associated symptoms, and upload clinical images for monitoring.

The symptom entry interface enables patients to document dermatological observations related to skin lesions for con- tinuous monitoring and clinical reference. Through this form, users can record the date of observation, specify the anatom- ical location of the lesion, and indicate localized symptoms such as itching, pain, bleeding, or color change. Additional fields allow patients to report changes in lesion size, provide preliminary information about the suspected skin cancer type, and upload images for visual documentation. The system also includes a clinical notes section for recording supplementary observations. To assist patients in identifying suspicious skin lesions, the interface presents the ABCDE guideline, which outlines key indicators such as asymmetry, border irregular- ity, color variation, diameter, and evolving appearance. This module supports proactive health tracking by enabling users to maintain structured records that can be reviewed during clinical assessments and AI-assisted diagnostic analysis.

S. Medical Library for Skin Cancer Reference

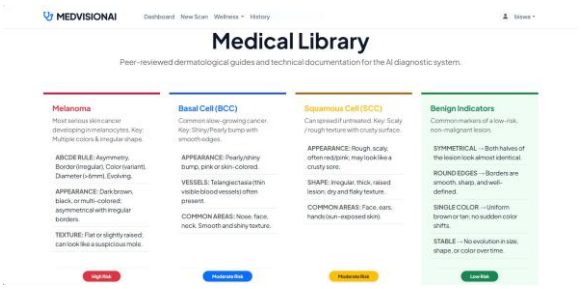

Fig. 23: Medical library interface providing dermatological reference infor- mation on melanoma, basal cell carcinoma, squamous cell carcinoma, and benign skin indicators.

The medical library module provides educational dermato- logical resources designed to support patient awareness and clinical understanding of

different skin lesion types. The interface presents structured reference information for major skin cancer categories including melanoma, basal cell carci- noma (BCC), and squamous cell carcinoma (SCC), along with indicators of benign lesions. Each section summarizes critical diagnostic characteristics such as visual appearance, texture, common anatomical locations, and distinguishing features. The melanoma section highlights the widely used ABCDE

rule, which helps identify suspicious lesions based on asymme- try, border irregularity, color variation, diameter, and evolving characteristics. Additionally, the module differentiates risk levels associated with each condition to guide early recognition and encourage timely medical consultation. By integrating educational content within the system, the medical library en- hances user understanding of dermatological conditions while complementing the AI-assisted diagnostic capabilities of the platform.

T. AI Model Documentation and Performance Metrics

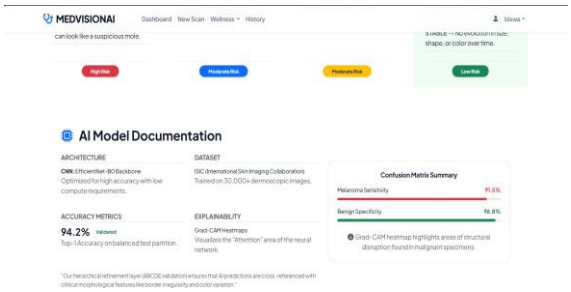

Fig. 24: AI model documentation panel presenting architecture details, dataset information, accuracy metrics, explainability tools, and confusion matrix performance summary.

The AI model documentation interface provides transparency regarding the architecture, training dataset, performance metrics, and explainability mechanisms of the diagnostic system. The model is built upon the EfficientNet-B0 convolutional neural network backbone, which is optimized for high classification accuracy while maintaining relatively low computational requirements. The training dataset is derived from the ISIC (International Skin Imaging Collaboration) repository and includes more than 30,000 dermoscopic images representing various dermatological conditions. Performance evaluation indicates a validated top-1 accuracy of 94.2% on a balanced test partition. The system also incorporates ex- plainable artificial intelligence techniques through Grad-CAM heatmaps, which visually highlight regions of the lesion image that contribute most significantly to the model's predictions. Additionally, the confusion matrix summary presents key diagnostic metrics such as melanoma sensitivity and benign specificity, demonstrating the model's ability to accurately distinguish malignant from non-malignant skin lesions. This documentation module promotes transparency and trust by providing users and clinicians with insights into how the AI diagnostic system operates and evaluates medical images.

U. Health Export and Diagnostic Report Generation

The health export module allows users to generate a com- prehensive clinical summary that consolidates their diagnostic history into a single report. This interface presents a struc- tured portfolio of previous AI analysis results, including the date of examination, predicted condition, and corresponding confidence score. Users can review detailed results for each diagnostic entry and export the compiled report in multiple for- mats. The system provides options to download the report as a

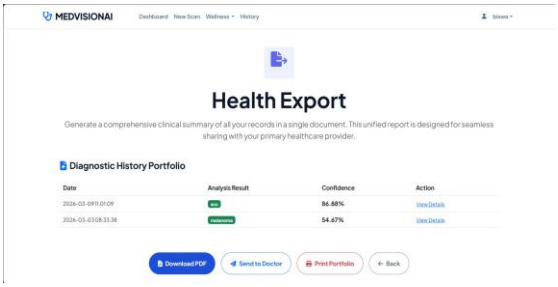

Fig. 25: Health export interface enabling users to generate, download, print, or share a consolidated diagnostic history report with healthcare providers.

PDF, print the portfolio for offline records, or directly share the document with healthcare professionals for further evaluation. By enabling seamless export and sharing of medical records, this feature enhances communication between patients and clinicians while supporting informed medical decision-making and follow-up care.

V. AI Diagnostic Report and Visual Explainability

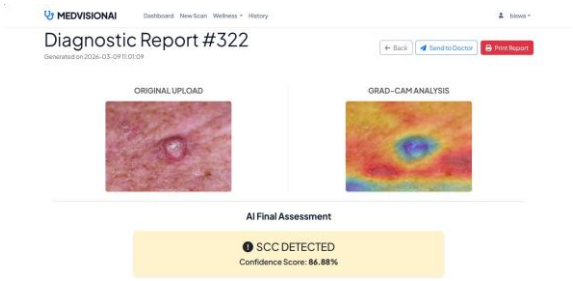

Fig. 26: AI diagnostic report displaying the original lesion image, Grad-CAM heatmap visualization, and the final AI assessment with confidence score.

The diagnostic report interface presents the detailed results generated by the AI-based skin lesion analysis system. The re- port includes the original uploaded dermoscopic image along- side a Grad-CAM heatmap visualization that highlights the re-gions of the image that most influenced the model's prediction. This visual explainability mechanism enables clinicians and users to understand how the neural network interprets lesion features during classification. The system then provides the final AI assessment, indicating the predicted skin condition along with a confidence score representing the probability of the classification. In the example shown, the system detects squamous cell carcinoma (SCC) with a confidence score of 86.88%. Additional options allow the user to print the report or share it with healthcare professionals for further clinical evaluation. This report generation module enhances trans- parency and supports collaborative medical decision-making by providing interpretable AI-assisted diagnostic results.

XI. DISCUSSION

The developed MedVisionAI system demonstrates how ar- tificial intelligence can be effectively integrated with clinical support tools to assist in the early detection and monitoring of skin cancer. The platform combines automated image-based diagnosis using a convolutional neural network model with multiple patient-oriented modules including symptom entry, clinical history tracking, educational medical libraries, and diagnostic report generation. By utilizing the EfficientNet-B0 architecture trained on the ISIC dermoscopic dataset, the sys- tem achieves high diagnostic performance while maintaining computational efficiency suitable for web-based deployment. The inclusion of Grad-CAM visual explanations enhances model transparency by highlighting lesion regions that influence predictions, which improves interpretability for clinicians and users. Additional features such as longitudinal diagnostic history, clinical record export, and structured symptom docu- mentation support continuous patient monitoring and facilitate communication with healthcare providers. Overall, the system demonstrates the potential of AI-driven diagnostic platforms to improve accessibility to dermatological screening, promote early detection of malignant lesions, and assist healthcare professionals in informed clinical decision-making.

XII. FUTURE WORK

While the proposed hybrid YOLOv9–Faster R-CNN frame- work demonstrates promising results, several directions remain for future exploration:

- Expansion of lesion categories: Incorporate additional lesion classes and rare subtypes to improve clinical applicability.
- Lightweight deployment: Develop optimized lightweight versions of the framework for mobile and edge devices, enabling real-time screening in low-resource environments.
- Metadata integration: Incorporate patient metadata (e.g., age, anatomical site, and skin type) to achieve context-aware predictions and reduce false positives.
- Explainable AI (XAI): Enhance interpretability by embedding explainable AI mechanisms, ensuring greater clinician trust and adoption of AI-assisted diagnostic tools.

XIII. CONCLUSION

In conclusion, the MedVisionAI system presents an in- tegrated AI-driven platform designed to assist in the early detection and monitoring of skin cancer through automated image analysis and patient-centered clinical tools. By leverag- ing a convolutional neural network model trained on a large dermoscopic image dataset, the system is capable of accurately classifying skin lesions while providing interpretable visual explanations using Grad-CAM heatmaps. The platform further enhances clinical usability through features such as symptom tracking, diagnostic history management, educational medical resources, and exportable diagnostic reports. These function- alities support both patients and healthcare professionals by enabling structured health monitoring, improved

awareness of dermatological conditions, and efficient communication of di- agnostic information. Overall, MedVisionAI demonstrates the potential of combining artificial intelligence with digital health systems to improve accessibility, support early diagnosis, and contribute to more informed dermatological care.

XIV. ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to their project guide, Waheeda Dhokley, Department of Computer Engineering, M. H. Saboo Siddik College of Engineering (MHSSCE), Mumbai, for her continuous guidance, valuable suggestions, and encouragement throughout the development of this research work.

The authors would also like to thank Dr. Ashfaq Shaikh, Head of the Department, Computer Technology (MHSSCE)

, for his academic support and for providing the necessary resources and encouragement required to complete this project successfully.

The authors are deeply grateful to Dr. Mohd. Shafi Pathan, Principal of M. H. Saboo Siddik College of Engineering (MHSSCE), Mumbai, for providing a supportive academic environment and institutional facilities that enabled the successful completion of this research.

Finally, the authors would like to acknowledge the collab- orative efforts of the project team members Ansari Usama, Biswajyoti Aown, and Ansari Armaan from the Department of Computer Engineering, MHSSCE, Mumbai, India, whose dedication and teamwork contributed significantly to the com- pletion of this work.

REFERENCES

[1] X. Zhang, Y. Li, and M. Kumar, "A Hybrid Model for Early Melanoma Detection: Integrating YOLOv9 and Faster R-CNN," in *Proc. IEEE International Conference on Biomedical Imaging (ICBI)*, 2024, pp. 112–119.

[2] R. Gupta and A. Sharma, "An Automated and Interpretable Computer- Aided Approach for Skin Cancer Diagnosis Using Genetic Program- ming," *IEEE Access*, vol. 11, pp. 45782–45795, 2023.

[3] S. Khan, P. Ramesh, and L. Thomas, "Segmentation and Classification of Skin Cancer Diseases Based on Deep Learning: Challenges and Future Directions," *Computers in Biology and Medicine*, vol. 167, 106076, 2024.

[4] T. Mendonc¸a, P. M. Ferreira, J. S. Marques, A. R. S. Marc¸al, and J. Rozeira, "PH² – A dermoscopic image database for research and bench- marking," in *Proc. 35th Annual International Conf. IEEE Engineering in Medicine and Biology Society (EMBC)*, 2013, pp. 5437–5440.

[5] P. Tschandl, C. Rosendahl, and H. Kittler, "The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions," *Scientific Data*, vol. 5, no. 180161, pp. 1–9, 2018.

[6] N. C. Codella, V. Rotemberg, P. Tschandl, M. E. Celebi, S. Dusza, D. Gutman, B. Helba, et al., "Skin Lesion Analysis Toward Melanoma Detection 2019: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC)," *arXiv preprint arXiv:1902.03368*, 2019.

[7] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional net- works for biomedical image segmentation," in *Medical Image Comput- ing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 234–241.

[8] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.

[9] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real- time object detection with region proposal networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015, pp. 91–99.

[10] Kaggle, "Skin Cancer MNIST: HAM10000," Available: https://www. kaggle.com/datasets/kmader/skin-cancer-mnist-ham10000

[11] Kaggle, "ISIC2018 Challenge Task 1 Dataset," Available: https://www.kaggle.com/datasets/tschandl/isic2018-challenge-task1-data-segmentation