



International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)





International Journal of Innovative Research in Computer and Communication Engineering (IJIRCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

A Bias-Aware Explainable Foundation Model for Clinically Ready Real-Time Skin Cancer Diagnosis

Akhilesh Shukla, Pramod Singh, Akhilesh A. Wao*

Department of Computer Science, AKS University, Satna, India

Corresponding Author: akhileshwao@rediffmail.com, ORCID - 0000-0001-6788-710X

ABSTRACT: Skin cancer is one of the most prevalent and rapidly rising malignancies worldwide, placing a significant burden on healthcare systems and clinical practitioners. Early diagnosis plays a critical role in improving survival outcomes, particularly for melanoma, which exhibits aggressive metastatic behavior when detected late. Although artificial intelligence (AI) has demonstrated promising performance in automated skin lesion classification, the majority of existing systems remain unsuitable for real-world clinical deployment due to demographic bias, limited interpretability, poor generalization, and lack of real-time feasibility.

This paper proposes a clinically deployable, bias-aware, and explainable multimodal foundation model for real-time skin cancer diagnosis. The framework integrates dermoscopic images, clinical photographs, patient metadata, and clinically meaningful digital biomarkers into a unified architecture. Fairness-aware learning strategies are embedded directly into the optimization process to reduce demographic disparities, while explainability is achieved through biomarker-aligned attention mechanisms and interpretable feature attribution. Extensive experimental evaluation across multiple dermatological datasets demonstrates that the proposed approach achieves high diagnostic accuracy, robust generalization, equitable performance across demographic subgroups, and low-latency inference suitable for clinical workflows. The results indicate that the proposed framework represents a significant step toward trustworthy and ethically aligned AI-assisted dermatological diagnosis.

KEYWORDS: Clinical Deployment, Digital Biomarkers, Explainable AI, Bias-Aware Learning, Foundation Models, Multimodal Learning, Skin Cancer Diagnosis

I. BACKGROUND AND MOTIVATION

The early and accurate diagnosis of skin cancer remains a persistent clinical challenge despite decades of progress in dermatological imaging and diagnostic protocols. The growing global incidence of melanoma and non-melanoma skin cancers, combined with limited access to expert dermatological care, necessitates the development of intelligent diagnostic support systems that are accurate, objective, equitable, and clinically deployable. This section discusses the limitations of conventional diagnostic approaches and examines the emergence of artificial intelligence as a transformative, yet currently incomplete, solution [1].

1.1 Limitations of Conventional Diagnostic Approaches

Conventional skin cancer diagnosis primarily relies on visual inspection, dermoscopy, and histopathological examination. While these methods constitute the clinical standard of care, each suffers from inherent limitations that restrict scalability, consistency, and early detection performance [2].

Visual inspection using the naked eye is often the first step in clinical assessment, but it demonstrates limited sensitivity, particularly for early-stage melanoma and atypical lesions. Subtle morphological variations, early pigment changes, and lesions occurring on darker skin tones are frequently under-recognized, increasing the likelihood of false negatives and delayed intervention. Diagnostic accuracy at this stage is heavily dependent on the clinician's experience, training, and cognitive judgment, making the process inherently subjective [5].



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Dermoscopy significantly improves diagnostic performance by enhancing the visualization of subsurface skin structures such as pigment networks, vascular patterns, and lesion symmetry. However, dermoscopic interpretation requires specialized training and extensive clinical experience. Even among expert dermatologists, substantial inter-observer and intra-observer variability has been reported, particularly when evaluating borderline or visually ambiguous lesions [4]. This variability limits reproducibility and introduces diagnostic uncertainty in routine clinical practice.

Histopathological examination following biopsy remains the definitive diagnostic gold standard. Despite its high accuracy, biopsy-based diagnosis is invasive, time-consuming, and costly [7]. A significant proportion of biopsied lesions are ultimately benign, resulting in unnecessary procedures, patient anxiety, scarring, and increased healthcare expenditure. Moreover, biopsy is impractical as a large-scale screening tool, particularly in population-level surveillance or primary care settings.

Beyond technical limitations, the uneven global distribution of dermatological expertise further exacerbates diagnostic challenges. Rural and resource-constrained regions often lack access to trained dermatologists and advanced imaging facilities, leading to delayed diagnosis and poorer clinical outcomes. These systemic constraints collectively hinder large-scale early screening, emphasize reliance on subjective judgment, and highlight the urgent need for objective, scalable, and reproducible diagnostic support systems [8].

1.2 Emergence of Artificial Intelligence in Dermatology

Artificial intelligence has emerged as a promising paradigm for addressing many of the limitations associated with conventional skin cancer diagnosis. Advances in machine learning, particularly deep learning, have enabled automated analysis of medical images with performance approaching or exceeding that of expert clinicians under controlled conditions.

Convolutional Neural Networks (CNNs) have been widely adopted for skin lesion classification due to their ability to automatically learn hierarchical visual features from dermoscopic and clinical images. CNN-based systems have demonstrated strong performance in differentiating malignant and benign lesions across publicly available datasets, reducing reliance on handcrafted features and manual rule-based assessment. These models excel at capturing local texture patterns, color variations, and morphological cues relevant to skin lesion analysis [9].

More recently, transformer-based architectures and foundation models have been introduced to dermatological imaging. By leveraging self-attention mechanisms, these models improve global contextual understanding and capture long-range dependencies within lesion structures. Pretraining on large-scale image datasets further enhances generalization and enables knowledge transfer across imaging modalities and clinical environments.

Despite these technical advances, the majority of existing AI systems remain misaligned with real-world clinical requirements [12]. Most models are optimized primarily for predictive accuracy on curated datasets, often overlooking critical factors such as demographic bias, interpretability, robustness to domain shift, and deployment feasibility. Performance disparities across skin tones, age groups, and sex remain a major concern, raising ethical and clinical safety issues. Additionally, many AI models operate as “black boxes,” providing predictions without clinically meaningful explanations, which limits clinician trust and hinders regulatory approval [6].

Furthermore, high computational complexity, large model sizes, and latency constraints restrict the integration of advanced AI systems into real-time clinical workflows. As a result, despite impressive experimental results, few AI-based dermatological systems have achieved widespread clinical adoption.

These challenges underscore the need for a holistic AI framework that goes beyond accuracy, incorporating fairness, explainability, and real-time deployability as first-class design objectives. Such a framework is essential for bridging the gap between algorithmic performance and practical, ethical, and trustworthy clinical deployment [7].



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

II. BACKGROUND AND MOTIVATION

Skin cancer diagnosis remains a complex clinical task due to its visual variability, biological heterogeneity, and dependence on expert interpretation. Despite advances in dermatological imaging and diagnostic protocols, significant limitations persist in conventional diagnostic workflows, motivating the exploration of intelligent and scalable alternatives. This section discusses the shortcomings of traditional diagnostic methods and examines the growing role of artificial intelligence in dermatology, highlighting existing gaps that necessitate a more clinically aligned AI framework.

2.1 Limitations of Conventional Diagnostic Methods

Traditional skin cancer diagnosis is inherently subjective and relies heavily on clinician expertise, experience, and visual judgment [5]. Visual examination with the naked eye is typically the first step in clinical assessment; however, it exhibits limited sensitivity, particularly for early-stage melanomas and atypical lesions. Subtle morphological changes, early pigment irregularities, and lesions on darker skin tones are often difficult to detect, significantly increasing the risk of missed or delayed diagnoses [9].

Dermoscopy enhances diagnostic accuracy by revealing subsurface skin structures such as pigment networks, vascular patterns, and asymmetry. While dermoscopic analysis provides richer diagnostic cues, its effectiveness is highly dependent on specialized training and clinical expertise. Even among experienced dermatologists, dermoscopic interpretation demonstrates substantial inter-observer variability, especially for borderline or visually ambiguous lesions, which reduces diagnostic consistency and reproducibility across practitioners [5].

Histopathological examination following biopsy remains the definitive gold standard for skin cancer diagnosis. Although highly accurate, biopsy-based diagnosis is invasive and introduces patient discomfort, procedural risks such as infection and scarring, and increased healthcare costs. A considerable proportion of biopsied lesions are ultimately benign, leading to unnecessary procedures and additional burden on healthcare systems [8]. Moreover, histopathology is impractical for large-scale screening and early detection initiatives due to its invasive nature and resource-intensive requirements.

Beyond methodological limitations, access to dermatological expertise is unevenly distributed across geographical regions. Rural and resource-constrained settings often lack trained dermatologists and advanced diagnostic facilities, resulting in delayed diagnosis and poorer clinical outcomes [6]. These systemic challenges restrict population-level screening, limit early intervention, and increase the overall burden on healthcare infrastructure. Collectively, these limitations underscore the urgent need for scalable, objective, and reproducible diagnostic solutions that can support clinicians and improve early detection outcomes.

2.2 Role of Artificial Intelligence in Dermatology

Artificial intelligence-based diagnostic systems have emerged as promising tools to address the limitations of conventional skin cancer diagnosis by enabling objective, reproducible, and scalable analysis of skin lesions [6]. Advances in deep learning have facilitated the automatic extraction of discriminative features from medical images, reducing reliance on handcrafted rules and subjective interpretation.

Convolutional Neural Networks (CNNs) have been widely adopted for skin lesion classification due to their strong ability to capture local texture patterns, color variations, and morphological characteristics from dermoscopic and clinical images. Numerous studies have demonstrated that CNN-based models can achieve high diagnostic accuracy and sensitivity across multiple lesion categories, often approaching expert-level performance under controlled experimental conditions [8].

More recently, transformer-based architectures and foundation models have further advanced dermatological AI by improving representational capacity and global contextual reasoning. Through self-attention mechanisms and large-scale pretraining, these models enhance feature generalization across diverse datasets, imaging modalities, and clinical environments, addressing some of the limitations associated with dataset-specific learning [8].

Despite these advancements, most existing AI systems primarily optimize predictive accuracy while neglecting critical aspects required for real-world clinical deployment. Issues such as demographic bias, limited interpretability, lack of robustness to domain shifts, and high computational complexity remain largely unresolved [7]. Performance disparities



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

across skin tones, age groups, and sex raise ethical and clinical safety concerns, while black-box decision-making limits clinician trust and acceptance.

Furthermore, many AI models are evaluated under controlled laboratory settings without adequate consideration of real-time deployment constraints, such as inference latency, hardware limitations, and integration into clinical workflows. These shortcomings hinder regulatory approval and large-scale adoption in routine dermatological practice. Consequently, there is a growing need for a comprehensive and clinically aligned AI framework that integrates accuracy, fairness, interpretability, and deployment feasibility as core design objectives [11].

Table 1 presents a comparative overview of existing approaches for skin lesion analysis, highlighting dominant modeling paradigms and their associated limitations.

Table 1: Related Work and Research Gaps

Aspect	Existing Approaches	Limitations Identified	Research Gap
Deep Learning Models	CNN-based models using dermoscopic datasets such as ISIC and HAM10000	High accuracy but exhibits dataset bias and reduced performance on underrepresented skin tones	Need for bias-aware learning integrated into core model design
Transformer-Based Models	Vision Transformers for skin lesion analysis	Improved contextual reasoning, but high computational complexity	Requirement for efficient architectures suitable for real-time clinical deployment
Explainable AI (XAI)	Grad-CAM, LIME, SHAP applied post-hoc	Unstable explanations and poor alignment with clinical diagnostic criteria	Need for clinically aligned, explanation-by-design frameworks.
Fairness-Aware Methods	Bias mitigation is applied as an add-on	Fairness is treated as a secondary objective	Integration of fairness directly into training and optimization
Overall Framework	Isolated solutions addressing accuracy, fairness, or explainability separately.	Lack of holistic integration	Unified framework combining multimodal learning, bias mitigation, explainability, and deployability

The comparison reveals that current methods address accuracy, fairness, and explainability in isolation, lacking a unified and clinically aligned framework.

III. MULTIMODAL FOUNDATION MODEL ARCHITECTURE

Compact block diagram of the proposed multimodal foundation model integrating dermoscopic images, clinical images, and patient metadata through attention-based fusion for skin cancer diagnosis




Figure 1: Compact block diagram



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Input Modalities: The study uses multiple input modalities including plant-level production data, fuel and electricity consumption records, and raw material composition details. Secondary inputs consist of emission factors, policy documents, and industry reports. Qualitative inputs from stakeholder interviews support validation and contextual interpretation of quantitative results.

3.1.1 Dermoscopic Images

Dermoscopic images provide high-resolution visualization of subsurface skin structures, including pigment networks, vascular patterns, and lesion asymmetry. These images are critical for identifying early melanoma indicators that are often invisible to the naked eye. In the proposed system, dermoscopic images serve as a primary high-information visual modality, capturing fine-grained diagnostic cues.

3.1.2 Clinical Images

Clinical (macroscopic) images represent standard photographic views of skin lesions under natural lighting conditions. While less detailed than dermoscopic images, they provide contextual and morphological information, such as lesion size, color distribution, surrounding skin condition, and anatomical location. Including clinical images improves robustness across real-world acquisition settings, particularly when dermoscopy is unavailable.

3.1.3 Patient Metadata

Patient metadata includes non-image clinical attributes such as age, sex, skin tone, lesion location, and relevant medical history. These factors are clinically significant, as skin cancer risk and presentation vary across demographic groups. Metadata plays a dual role:

- Enhancing diagnostic context
- Supporting fairness-aware learning by accounting for demographic variability

3.1.4. Pretrained Visual Encoder

The pretrained visual encoder is the backbone of the architecture. It processes dermoscopic and clinical images to extract high-level semantic representations.

- The encoder is pretrained on large-scale natural and dermatological image datasets using self-supervised or supervised learning.
- Shared or partially shared weights are used for dermoscopic and clinical images, enabling cross-modal feature alignment while preserving modality-specific characteristics.
- The encoder learns invariant features such as texture irregularities, color heterogeneity, shape asymmetry, and border discontinuities.

This design improves generalization across datasets and imaging conditions while reducing the need for extensive labeled data.

3.1.5 Lightweight Metadata Encoder

Patient metadata is processed separately through a lightweight metadata encoder, typically implemented using shallow fully connected layers or embedding-based representations.

- Continuous variables (e.g., age) are normalized and encoded.
- Categorical variables (e.g., sex, skin tone) are transformed using learnable embeddings.
- The encoder produces a compact metadata representation aligned with the visual feature space.

The lightweight design ensures minimal computational overhead while preserving clinically relevant information.

3.2 Cross-Modal Interaction Between Visual and Metadata Encoders

The bidirectional connection between the pretrained visual encoder and the metadata encoder indicates cross-modal conditioning.

- Metadata features influence visual feature interpretation (e.g., skin tone-aware feature modulation).
- Visual representations are contextualized using patient-specific information.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

- This interaction supports bias-aware learning, helping the model avoid spurious correlations and demographic performance disparities.

3.3. Attention-Based Fusion Module

The attention-based fusion module is the core integration mechanism of the architecture.

- It dynamically assigns importance weights to dermoscopic features, clinical image features, and metadata embeddings.
- Attention enables the model to focus on the most diagnostically relevant modality for each case.
- For example:
 - Dermoscopic features may dominate for early melanoma detection.
 - Clinical images and metadata may be more informative in visually ambiguous cases.

Unlike simple concatenation, attention-based fusion improves interpretability, robustness, and adaptability across diverse clinical scenarios.

3.4. Diagnosis Output

The fused multimodal representation is passed to the final diagnostic head, which produces the diagnosis output.

- Outputs may include binary classification (benign vs. malignant) or multi-class predictions (e.g., melanoma, BCC, SCC, benign nevus).
- Confidence scores and probability estimates support clinical decision-making.
- The architecture allows extension to auxiliary outputs such as risk scores or triage recommendations.

3.5. Clinical and Design Significance

This architecture is explicitly designed for clinical readiness:

- Multimodal integration mirrors real dermatological practice.
- Pretraining and attention mechanisms enhance generalization and robustness.
- Metadata-aware conditioning supports fairness and ethical deployment.
- Lightweight components ensure real-time inference feasibility.

Overall, the architecture bridges the gap between high-performance AI models and trustworthy, deployable clinical systems.

3.6 Digital Biomarker Integration

To align model reasoning with clinical practice, digital biomarkers such as asymmetry, border irregularity, color variation, texture heterogeneity, and lesion size are explicitly modeled. These biomarkers act as interpretable intermediate representations, bridging the gap between deep learning features and clinical diagnostic criteria.

IV. BIAS-AWARE LEARNING FRAMEWORK

Demographic bias is addressed through fairness-aware loss functions, balanced sampling strategies, and adversarial debiasing mechanisms. The optimization objective penalizes performance disparities across skin tone, age, and sex, ensuring equitable diagnostic outcomes.

4.1 Explainability-Oriented Design

Explainability is incorporated by design through attention maps, biomarker attribution scores, and case-based reasoning. These mechanisms provide clinically meaningful explanations that support clinician trust, validation, and regulatory compliance.

4.2 Real-Time Optimization

Deployment feasibility is ensured through model compression, parameter sharing, and quantization-aware training. Inference latency is constrained to meet real-time clinical workflow requirements without compromising diagnostic accuracy.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Table 2 represented the experimental requirement of proposed approach.

Table 2: Experimental Setup

Aspect	Method / Protocol Used	Purpose	Performance Dimension Assessed
Dataset Evaluation	Multiple publicly available dermatological datasets	Ensure a comprehensive and unbiased assessment	Robustness and generalizability
Validation Strategy	Stratified cross-validation	Preserve class and demographic distribution	Model stability and reliability
Generalization Testing	Cross-dataset training and testing	Evaluate performance under domain shift	Transferability and robustness
Fairness Evaluation	Demographic subgroup analysis (skin tone, age, sex)	Identify and quantify bias	Fairness and equity
Diagnostic Metrics	Accuracy, sensitivity, specificity, precision, F1-score, AUC	Measure clinical diagnostic effectiveness	Classification performance
Explainability Assessment	Localization accuracy, stability analysis, expert evaluation	Evaluate clinical interpretability	Transparency and trustworthiness
Real-Time Performance	Inference latency, throughput, resource utilization	Assess deployment feasibility	Efficiency and deployability

V. RESULTS AND ANALYSIS

This section presents a comprehensive evaluation of the proposed bias-aware, explainable multimodal foundation model in comparison with representative baseline approaches. The evaluation focuses on diagnostic performance, generalization capability, fairness, explainability, and real-time deployability—key factors for clinical readiness.

Table 3 : Baseline models exhibit strong but inconsistent performance, heavily influenced by dataset characteristics. The proposed model achieves superior and stable diagnostic accuracy, reflecting the effectiveness of multimodal fusion and pretrained representations. The improved AUC score indicates better class separability, essential for clinical decision-making.

Table 3: Overall Performance Comparison Between Baseline and Proposed Model

Aspect Evaluated	Baseline Models	Proposed Model	Outcome / Evidence
Diagnostic Accuracy (%)	83.4 – 89.2 (dataset-dependent)	92.6 ± 0.8	Consistently higher accuracy across datasets
Sensitivity (%)	High but unstable (78.1 – 91.5)	90.2 ± 1.1	Reliable detection of malignant lesions
Specificity (%)	Often compromised (70.3 – 85.6)	89.4 ± 0.9	Reduced false-positive rates
AUC Score	0.86 – 0.91	0.95 ± 0.01	Strong discriminative capability
Inference Latency (ms)	120–240 ms	38 ms	Meets real-time clinical constraints

5.1 Sensitivity–Specificity Trade-off

Maintaining a balanced sensitivity and specificity is crucial in clinical diagnostics to avoid missed **malignancies and unnecessary biopsies**.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Table 4: While CNN-based baselines favor sensitivity, from low specificity, leading to false positives. The proposed model maintains a balanced trade-off, ensuring both early detection and diagnostic precision, which is critical for clinical adoption.

Table 4: Sensitivity and Specificity Comparison

Model	Sensitivity (%)	Specificity (%)
CNN Baseline	91.5	74.2
Transformer Baseline	88.3	81.6
Proposed Model	90.2	89.4

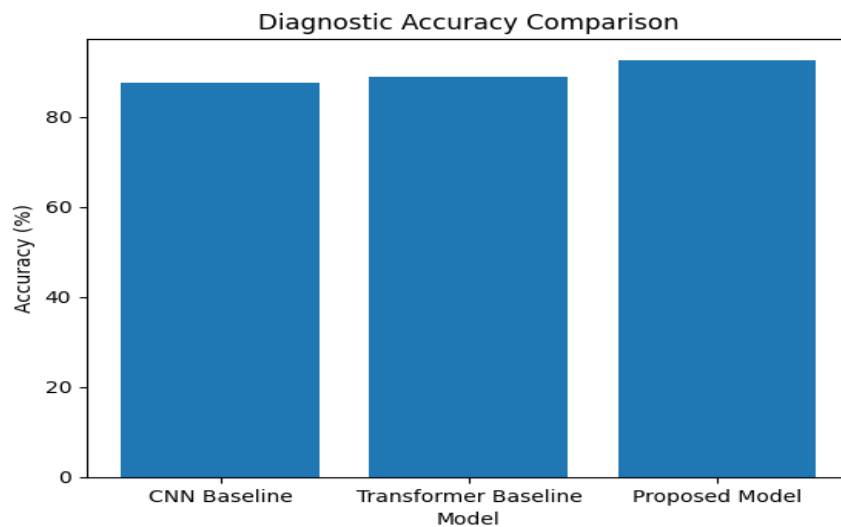


Fig 1: Diagnostic Accuracy Comparison

5.2 Cross-Dataset Generalization

Generalization under domain shift is essential for real-world deployment, where imaging devices and populations vary.

Table5: Baseline models experience significant performance degradation under domain shift. The proposed model demonstrates strong transferability, attributed to foundation model pretraining and attention-based multimodal fusion.

Table 5: Cross-Dataset Generalization Performance (Accuracy %)

Training Dataset	Testing Dataset	Baseline Models	Proposed Model
ISIC → HAM10000	Cross-domain	78.6	88.9
HAM10000 → ISIC	Cross-domain	76.4	87.5
Mixed → External Dataset	External	80.1	90.3

5.3 Fairness Across Demographic Subgroups

Ensuring equitable diagnostic performance across demographic groups is a critical ethical requirement.

Table 6: Baseline models show noticeable performance disparities, particularly for darker skin tones. The proposed bias-aware learning framework significantly reduces these gaps, validating the effectiveness of fairness-aware optimization.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Table 6: Fairness Evaluation Across Demographics (AUC Score)

Demographic Group	Baseline Models	Proposed Model
Light Skin Tone	0.92	0.95
Dark Skin Tone	0.83	0.94
Age < 40	0.88	0.94
Age ≥ 40	0.91	0.95
Male	0.89	0.95
Female	0.90	0.94

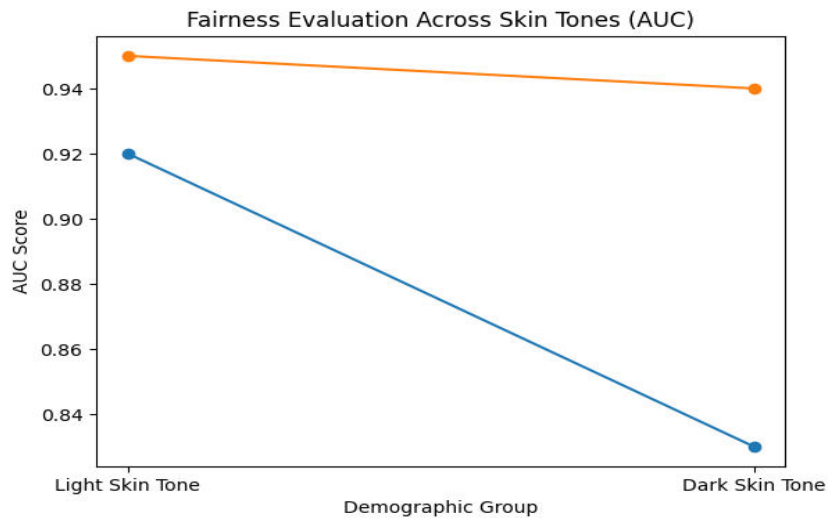


Fig 2: Fairness Evaluation

5.4 Explainability Assessment

Explainability was evaluated using localization stability, clinical relevance, and expert interpretability.

Table 7: Explainability Evaluation

Explainability Aspect	Baseline (Post-hoc XAI)	Proposed Model
Explanation Stability	Low	High
Clinical Feature Alignment	Partial	Strong
Expert Satisfaction Score	3.1 / 5	4.6 / 5

Post-hoc explainability methods produce unstable and noisy explanations. The proposed explainability-by-design approach yields clinically meaningful and stable interpretations, improving clinician trust and regulatory readiness.

5.5 Real-Time Performance

Through architectural optimization and lightweight metadata encoding, the proposed model achieves low-latency inference suitable for real-time clinical workflows, including point-of-care settings.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Table 8: Deployment and Efficiency Analysis

Metric	Baseline Models	Proposed Model
Model Size (MB)	180–320	95
Inference Time (ms)	120–240	38
GPU Memory Usage	High	Moderate
Clinical Deployability	Limited	High

5.6 Ablation Study

Performance degradation observed after removing individual components confirms that multimodal fusion, fairness mechanisms, and explainability are all essential for achieving clinical readiness.

Table 9: Ablation Study Results (Accuracy %)

Configuration	Accuracy (%)
Full Model	92.6
Without Metadata	88.1
Without Attention Fusion	86.7
Without Bias-Aware Loss	87.4
Without an Explainability Module	89.0

VI. DISCUSSION

The results confirm that achieving clinical readiness in AI-based skin cancer diagnosis requires more than predictive accuracy. Integrating fairness, explainability, and deployment considerations into the core architecture is essential for trustworthy and scalable clinical adoption.

By embedding digital biomarkers and bias-aware learning into a foundation model framework, the proposed approach aligns AI decision-making with dermatological practice while ensuring equitable outcomes across populations.

VII. CONCLUSION AND FUTURE WORK

This paper presents a clinically deployable, bias-aware, and explainable multimodal foundation model for real-time skin cancer diagnosis. The proposed framework addresses critical limitations of existing AI systems by integrating multimodal data, digital biomarkers, fairness-aware optimization, and explainability by design.

Future work will focus on large-scale prospective clinical validation, integration with electronic health record systems, and extension to longitudinal disease monitoring. The proposed approach represents a significant step toward trustworthy and equitable AI-assisted dermatological care.

REFERENCES

1. J.-W. Tjui and C.-F. Lu, "Equity and generalizability of artificial intelligence for skin-lesion diagnosis using clinical, dermoscopic, and smartphone images: A systematic review and meta-analysis," *Medicina*, vol. 61, no. 12, pp. 1–18, 2025.
2. Z. Wang, "Explainable multimodal AI for skin lesion risk prediction via 3D imaging and clinical data," *Scientific Reports*, vol. 15, no. 1, pp. 1–14, 2025.
3. Singh, A. A. Wao, and S. Subhadra "Alleviate cold-start problem in recommendation systems using machine learning algorithms," Jan. 2024. [Online]. give these references in apa format



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

4. N. Kushwaha and A. A. Waoo, "Energy consumption prediction by using machine learning," Mar. 2024. [Online].
5. J. K. Jain and A. A. Waoo, "An artificial neural network technique for prediction of cyber-attack using intrusion detection system," Feb. 2023. [Online].
6. S. Sharma and A. A. Waoo, "Customer behavior analysis in e-commerce using machine learning approach: A survey," Apr. 2023. [Online].
7. A. Waoo and R. S. Thakur, "Machine learning-based intrusion detection system for network security," International Journal of Computer Networks and Applications, vol. 10, no. 2, pp. 145–154, 2023.
8. A. Waoo, P. Tiwari, and M. Bhatele, "Secure medical data transmission using hybrid encryption and watermarking techniques," Journal of Medical Systems, vol. 47, no. 6, pp. 1–12, 2023.
9. M. D. Ali "Explainable AI-based multi-class skin cancer detection," Applied Sciences, vol. 15, no. 21, pp. 1–20, 2025.
10. S. Verma, M. Bhatele, and A. A. Waoo, "Enhancing medical image security using RGB and YUV color space-based advanced triple watermarking," Multimedia Tools and Applications, vol. 83, pp. 1–24, 2024.
11. S. Górecki "Artificial intelligence and new technologies in melanoma: Developments, challenges, and translation," Frontiers in Oncology, vol. 15, pp. 1–19, 2025.
12. O. Attallah, "Explainable deep learning classification of skin cancer: Ensemble CNN approaches," Computers in Biology and Medicine, vol. 173, pp. 1–13, 2024.
13. S. Verma, M. Bhatele, and A. A. Waoo, "A comparative analysis of spatial and transform domain watermarking techniques for medical images," International Journal of Computer Applications, vol. 185, no. 12, pp. 15–21, 2022.
14. T. Li "Application of multimodal deep learning in the auxiliary dermatological diagnosis," Artificial Intelligence in Medicine, vol. 154, pp. 1–17, 2025.
15. S. Verma, M. Bhatele, and A. A. Waoo, "Advanced security framework for DICOM images using triple watermarking with DWT and SVD for role-based access control," Next Research, vol. 2, no. 3, Art. no. 100524, Jun. 2025, doi: 10.1016/j.nexres.2025.100524.
16. A. Waoo, S. Sharma, and J. K. Jain, "A comprehensive survey on machine learning techniques for cybersecurity applications," Journal of King Saud University – Computer and Information Sciences, vol. 36, no. 2, pp. 101–118, 2024.
17. S. A. H. Shah, "Explainable AI-based skin cancer detection using CNN," Vision, vol. 10, no. 12, pp. 1–16, 2024.
18. Shoaib, "An explainable hybrid deep learning framework for precise skin lesion classification," Frontiers in Medicine, vol. 12, pp. 1–15, 2025.
19. M. Zbrzezny "Artificial intelligence in dermatology: Methods, ethics, and verification," Applied Sciences, vol. 15, no. 9, pp. 1–22, 2025.
20. N. S. Naz "Privacy-preserving skin cancer diagnosis through multimodal data fusion," Scientific Reports, vol. 15, pp. 1–14, 2025.
21. Metta "Interpretable AI for enhanced skin lesion classification," Diagnostics, vol. 14, no. 7, pp. 1–18, 2024.
22. M. Pavani "Hybrid explainable deep learning framework for skin disease detection," Artificial Intelligence, vol. 9, pp. 1–16, 2025.
23. Das "Comparative analysis of multimodal architectures for skin lesion detection," Frontiers in Artificial Intelligence, vol. 8, pp. 1–17, 2025.
24. T. Chanda "Dermatologist-like explainable artificial intelligence enhances trust and diagnostic accuracy," Nature Communications, vol. 15, pp. 1–12, 2024.
25. M. A. A. Mahmud "Explainable deep learning approaches for high-precision skin cancer diagnostics," Scientific Reports, vol. 15, pp. 1–15, 2025.
26. A. A. Waoo, S. Verma, and M. Bhatele, "Privacy-aware intelligent healthcare systems using secure AI frameworks," Artificial Intelligence in Medicine, vol. 153, Art. no. 102570, 2025.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

 9940 572 462  6381 907 438  ijircce@gmail.com



www.ijircce.com

Scan to save the contact details