

Explainable Melanoma Diagnosis with Contrastive Learning and LLM-based Report Generation

Junwen Zheng¹, Xinran Xu¹, Li Rong Wang^{1,2}, Chang Cai¹, Lucinda Siyun Tan³, Dingyuan Wang³, Hong Liang Tey^{1,3*}, Xiuyi Fan^{1*}

¹Nanyang Technological University, Singapore

²Centre for Frontier AI Research, A*STAR, Singapore

³National Skin Centre, National Healthcare Group, Singapore

{JUNWEN003, XINRAN007, LIRONG002, CHANG023}@e.ntu.edu.sg,

{lucindatan, dywang}@nsc.com.sg, {teyhongliang, xyfan}@ntu.edu.sg

Abstract

Deep learning has demonstrated expert-level performance in melanoma classification, positioning it as a powerful tool in clinical dermatology. However, model opacity and the lack of interpretability remain critical barriers to clinical adoption, as clinicians often struggle to trust the decision-making processes of black-box models. To address this gap, we present a Cross-modal Explainable Framework for Melanoma (CEFM) that leverages contrastive learning as the core mechanism for achieving interpretability. Specifically, CEFM maps clinical criteria for melanoma diagnosis—namely Asymmetry, Border, and Color (ABC)—into the Vision Transformer embedding space using dual projection heads, thereby aligning clinical semantics with visual features. The aligned representations are subsequently translated into structured textual explanations via natural language generation, creating a transparent link between raw image data and clinical interpretation. Experiments on public datasets demonstrate 92.79% accuracy and an AUC of 0.961, along with significant improvements across multiple interpretability metrics. Qualitative analyses further show that the spatial arrangement of the learned embeddings aligns with clinicians’ application of the ABC rule, effectively bridging the gap between high-performance classification and clinical trust.

Code — <https://eattt-wen.github.io/CEFM/>

Introduction

The global incidence of malignant melanoma continues to rise steadily, with this aggressive skin cancer accounting for over 80% of skin cancer-related deaths (Mallardo, Basile, and Vitale 2025). Traditional diagnosis predominantly relies on physicians’ visual assessments and dermoscopic inspection. To support systematic evaluation, clinicians often refer to the well-established **ABCD** rule, which assesses key dermoscopic features: *Asymmetry*, *Border irregularity*, *Color variation*, *Differential structure*. However, the accuracy of such evaluations is highly dependent on the clinician’s expertise and is susceptible to subjective variability.

Recent advances in convolutional neural networks (CNNs) have led to remarkable improvements in melanoma

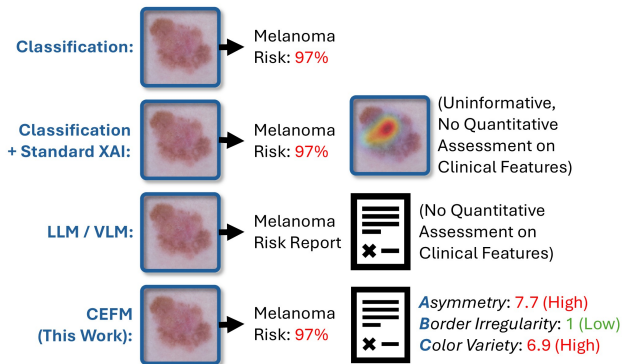


Figure 1: Advantages of the Cross-modal Explainable Framework (CEFM) over existing frameworks.

classification performance (Almufareh et al. 2024; Hussien and Alasadi 2024). Despite achieving high diagnostic accuracy ($\geq 90\%$), most CNN-based systems suffer from a lack of interpretability, which limits their adoption in clinical practice. These models are often regarded as “black boxes,” with opaque decision-making processes that undermine clinicians’ trust in AI-assisted diagnosis (Xu and Shuttleworth 2024). Enhancing transparency through clinically aligned interpretable representations is essential for building physician trust and integrating AI into clinical workflows (Hossain et al. 2025). In this context, explainable AI (XAI) approaches that generate clinically meaningful, context-aware explanations are increasingly recognized as critical to bridging the gap between algorithmic performance and real-world clinical utility (Toh et al. 2025).

Existing XAI methods, such as Grad-CAM (Selvaraju et al. 2017), highlight regions of interest but fail to integrate domain-specific ABCD diagnostic priors, limiting their clinical depth and trustworthiness.

To bridge this trust gap between black-box AI models and clinical deployment, we propose a **Cross-modal Explainable Framework for Melanoma diagnosis (CEFM; Figure 1)**. Our system is designed to produce semantically rich explanations aligned with established diagnostic criteria, thereby enhancing the interpretability, credibility, and

*Corresponding authors

usability of AI in dermatological settings. The core contributions of this work are as follows:

1. **Clinically-Grounded Feature Extraction Pipeline:** We introduce a coarse-to-fine segmentation module that leverages UltraLight VM-UNet and SAM2 to extract three key dermoscopic features—**Asymmetry**, **Border irregularity**, and **Color variation**—based on the ABCD rule, enabling interpretable lesion characterization.
2. **Cross-Modal Contrastive Alignment:** We develop a contrastive learning architecture that aligns visual representations from a Vision Transformer with clinically features, enabling the classifier to learn semantically meaningful embeddings that reflect dermatological reasoning.
3. **Hybrid Diagnostic Report Generation:** We propose a novel explanation module integrating CLIP-based medical concept activation with a domain-adapted language model (DeepSeek) to generate structured, clinically accurate diagnostic reports supporting clinical decision-making.

Our ablation study shows that each module in the explanation pipeline contributes uniquely to generating coherent and clinically meaningful melanoma explanations, further supported by a structured user study with three board-certified dermatologists. Experts highlighted the framework’s strong interpretability (avg. score: 4.60/5), endorsed the ABC feature analysis, and affirmed its clinical utility for early triage, junior clinicians, and longitudinal lesion tracking. Furthermore, our framework demonstrates strong quantitative performance across segmentation and classification tasks.

Related Work

Automatic Classification of Melanoma

In recent years, deep learning has made remarkable progress in the Automatic Classification of Melanoma in dermoscopic images. Sabir and Mehmood (2024) compared a baseline CNN, ResNet-18, and EfficientNet-B0, and demonstrated that EfficientNet-B0 performed best, achieving a remarkable accuracy of 97%. In our work, we reproduce their comparison of convolutional backbones and extend it with a Vision Transformer (ViT), which outperformed EfficientNet-B0. In parallel, Abir et al. (2024) focused on image preprocessing and feature engineering, applying noise removal and lesion segmentation strategies that significantly improved ResNet-based performance.

Multimodal Learning in Dermatological Diagnosis

Beyond image-based unimodal models, multimodal learning approaches have gained increasing attention in dermatological diagnosis due to their ability to integrate complementary clinical information and improve diagnostic accuracy. Previous studies performed this integration using traditional algorithms such as Random Forest, Logistic Regression, and Support Vector Machines for multi-type skin disease classification and risk prediction (Almustafa 2025). However, these works predominantly focus on classification alone, without addressing interpretability.

XAI for Melanoma Detection

Existing work aimed at improving interpretability with explainable AI (XAI) includes Thomas et al. (2021), who applied a U-Net with a ResNet-50 encoder–decoder to segment skin-tissue images into 12 semantic classes, producing tissue-distribution maps that enable pathology-style explanations. By contrast, post-hoc saliency methods such as Grad-CAM merely highlight regions of model attention without linking them to diagnostic structures (Gamage et al. 2024). Another line of work generates clinically relevant textual descriptions using LSTMs based on visual attention maps after classification (Zhang et al. 2019). However, these approaches still do not generate clinically meaningful explanations grounded in established diagnostic criteria.

In the area of Clinical Explanation of Melanoma Classification, the work by Chanda et al. (2024) is most closely related to ours. They introduced a guided attention mechanism (CompA), which aligns model-generated feature maps with clinician-annotated regions for more precise focus, and combines it with Grad-CAM to provide interpretable diagnoses. However, their approach requires extensive manual annotation for model construction, presenting scalability challenges. Additionally, highlighting model attention regions is perceived as less useful for clinical decision-making than textual explanations (Kayser et al. 2024), driving interest in explanation methods beyond visual attention maps.

LLM-Based Report Generation

One strategy to producing textual explanations is to employ large language models (LLM) to generate structured reports. Wang et al. (2023) proposed R2GenGPT, a lightweight framework that maps image features into the embedding space of a frozen LLM via a small projection module, enabling efficient end-to-end generation. Meanwhile, other studies proposed a prompt-guided framework incorporating anatomical priors and clinical context (Li et al. 2024). Key anatomical regions are detected via object detection, and region-level descriptions are generated by a transformer decoder. Structured reports are then generated by the LLM with both the region-level descriptions and the clinical prompts. While this improves report structure and interactivity, the approach offers limited diagnostic guidance for generating region-level descriptions. Additionally, the clinical prompts, based on history, indications, and reasons for examination, do not explicitly reflect diagnostic criteria. As a result, the generated reports may fall short in supporting accurate diagnosis.

Methodology

Our proposed **Cross-modal Explainable Framework for Melanoma diagnosis (CEFM)** is designed to integrate accurate image-based classification with clinically meaningful explanations, enhancing the interpretability of image classification models while preserving high performance. As illustrated in Figure 2, CEFM comprises four components:

1. a *Deep Neural Network (DNN) Classification Pipeline*,
2. a *Clinical Explanation Pipeline*,

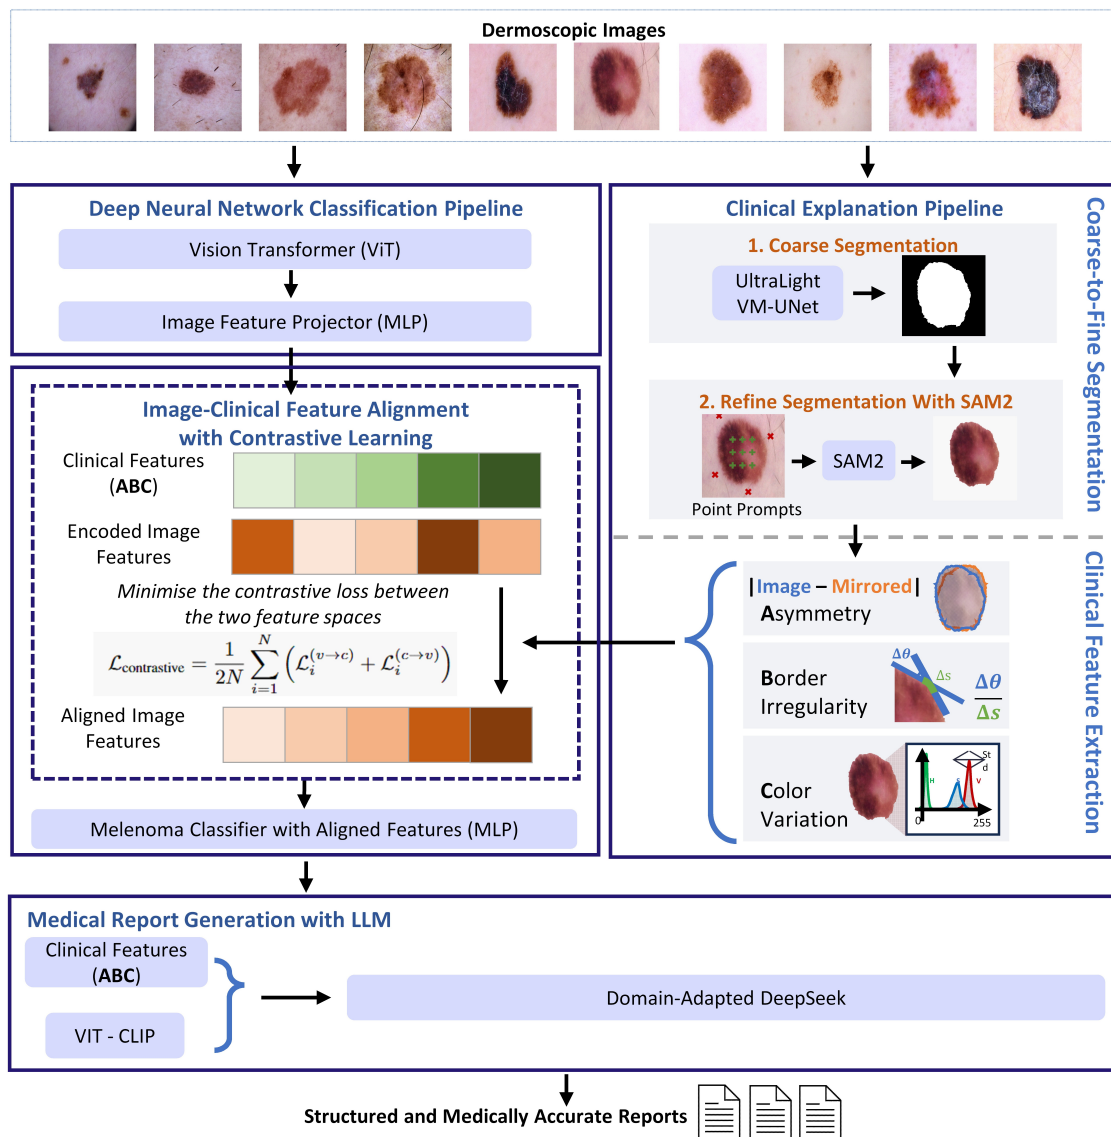


Figure 2: We propose a cross-modal melanoma diagnosis framework integrating four interconnected pipelines: a ViT-based classification pipeline that extracts semantic image features, a clinical explanation pipeline that segments lesions and quantifies ABC criteria, a contrastive module that aligns visual and clinical representations for interpretability, and a report generation module that produces structured diagnostic reports using CLIP descriptors and a domain-adapted LLM (DeepSeek).

3. a *Contrastive Learning Module*, and
4. a *Medical Report Generation Module*.

The DNN Classification Pipeline extracts semantic image embeddings from dermoscopic images using a ViT, followed by an image feature projector for dimensional alignment. In parallel, the Clinical Explanation Pipeline performs coarse-to-fine lesion segmentation using UltraLight VM-UNet and SAM2, from which three clinically relevant features—*Asymmetry*, *Border Irregularity*, and *Color Variation*—are computed. These image and clinical features are then fused via a cross-modal contrastive learning mechanism that encourages alignment between visual embeddings

and diagnostic priors, thereby enhancing interpretability. Finally, to facilitate clinician understanding and trust, CEFM includes a report generation module that leverages CLIP-based concept activation and a domain-adapted large language model to produce structured and medically grounded diagnostic reports. We introduce them as follows.

Deep Neural Network Classification Pipeline

The **DNN Classification Pipeline** in CEFM is responsible for extracting high-level semantic features from dermoscopic images to support melanoma classification and enable alignment with clinically derived features. We evaluated several backbone architectures for latent feature ex-

Feature	Clinical Concept	Calculation Formula	Interpretation
Asymmetry (A)	The two halves of the mole are asymmetrical	$A = \frac{\sum I(x, y) - I_{\text{mirror}}(x, y) }{\sum M(x, y)}$	Differential pixels in lesion area / Total lesion area pixels
Border (B)	The border of the mole is irregular, jagged, or blurred	$B_2 = \frac{1}{N} \sum_{i=1}^N \kappa_i \text{ where } \kappa_i = \frac{\Delta\theta_i}{\Delta s_i}$	Mean curvature of the lesion border with κ_i representing local curvature at point i
Color (C)	The mole has uneven colour, which may include brown, black, and sometimes pink, white, or blue spots	$\sigma_H = \sqrt{\frac{1}{N} \sum_{i=1}^N (H_i - \mu_H)^2}$ $\sigma_S = \sqrt{\frac{1}{N} \sum_{i=1}^N (S_i - \mu_S)^2}$ $\sigma_V = \sqrt{\frac{1}{N} \sum_{i=1}^N (V_i - \mu_V)^2}$	Standard deviations of HSV color channels

Table 1: Quantitative computation of ABC features from dermoscopic images. $I(x, y)$ and $I_{\text{mirror}}(x, y)$ are pixel values in the original and mirrored images, $M(x, y)$ is the lesion mask; κ_i is the local curvature calculated from the angle change ($\Delta\theta_i$) and arc length (Δs_i), N is the number of valid curvature sample points; H_i , S_i , and V_i are the hue, saturation, and value; μ_H , μ_S , and μ_V are their respective mean values.

traction, including ResNet50, DenseNet121, Inception-V3, EfficientNet-B2, and VGG16. Among these, the ViT consistently achieved the best classification performance and was therefore selected as the core feature extractor for our framework.

Given an input dermoscopic image, the ViT encodes spatial and texture information into a fixed-dimensional latent representation. These representations are then passed through a **Multilayer Perceptron (MLP)** head, which acts as a feature projector to transform the ViT-encoded image embeddings into a common space suitable for contrastive alignment with clinical features. This MLP consists of two fully connected layers with non-linear activations, followed by L_2 normalization to ensure consistent feature scales.

The projected image embeddings serve dual purposes: they are used both for melanoma classification through a separate classification head and as input to the contrastive learning module, which aligns these features with clinically extracted ABC features. This modular design ensures that the learned representations are both discriminative and semantically aligned with dermatological reasoning.

Clinical Explanation Pipeline

The **Clinical Explanation Pipeline** is designed to extract human-interpretable, clinically relevant features from dermoscopic images, enabling transparent reasoning aligned with dermatological diagnostic practices. This pipeline complements the DNN-based classification by introducing domain knowledge through a two-stage process: *lesion segmentation* and *clinical feature extraction*. First, a coarse-to-fine segmentation module isolates the lesion region from surrounding skin using lightweight and accurate deep segmentation models. Then, clinically meaningful features—specifically **Asymmetry**, **Border Irregularity**, and **Color Variation**—are computed from the segmented lesion mask. These features correspond to the “ABC” components of the ABCD rule and serve as structured diagnostic priors. The extracted features are later aligned with image representations computed in the DNN Classification Pipeline to ensure the **internal consistency** between the two pipelines.

Lesion Segmentation Lesion segmentation serves as the foundational step in the Clinical Explanation Pipeline, aiming to accurately delineate the melanoma region from the surrounding skin. This ensures that all downstream computations are based on the relevant lesion area. To achieve this, we employ a two-stage segmentation strategy for accurate lesion masks with minimal manual intervention. Coarse segmentation uses a pretrained UltraLight VM-UNet (trained on ISIC 2018), efficiently delineating lesion regions but yielding rough boundaries. These coarse masks serve as pseudo-labels for subsequent refinement. In fine segmentation, we leverage SAM (Ravi et al. 2024): foreground and background points are sampled automatically from the coarse masks to prompt SAM2, which generates multiple candidates per image; we then select the mask with highest IoU against the pseudo-label. This yields high-precision lesion contours for downstream ABC feature computation.

Clinical Feature Extraction From the segmented lesion images, we designed computational functions to automatically extract clinical features that are routinely used in dermatological practice to assess the malignancy risk of skin lesions. Our analysis focuses on three criteria that can be reliably and robustly computed: *Asymmetry*, *Border*, and *Colour*. The *Differential structure* feature is excluded due to the lack of fine-grained expert annotations, as the identification and delineation of these structures typically require precise, region-specific labeling by dermatology specialists. For each of the retained criteria, we design quantitative feature descriptors derived from the refined segmentation masks (Table 1).

Feature Spaces Alignment with Contrastive Learning

To bridge the semantic gap between deep image features and clinically meaningful descriptors, we introduce a **cross-modal contrastive learning** approach that aligns high-level visual representations, computed in the DNN Classification Pipeline, with structured clinical features computed in the Clinical Explanation Pipeline. This alignment grounds abstract image embeddings in interpretable dermatological concepts, enhancing both model transparency and robust-

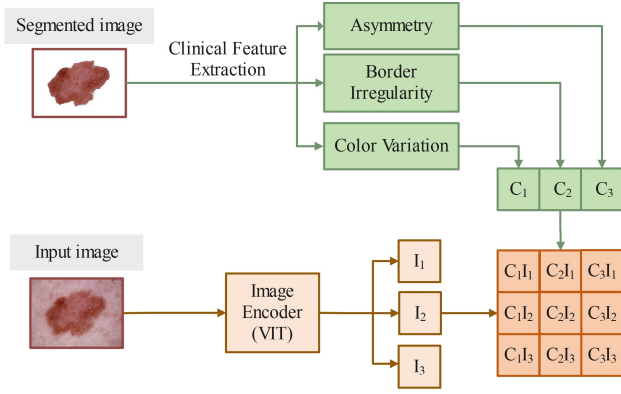


Figure 3: Overview of the cross-modal alignment process. Clinical features extracted from lesion segmentations are projected into a shared latent space, while image features from a ViT encoder are jointly aligned via element-wise interactions to supervise contrastive learning.

ness. An overview of the cross-modal interaction is illustrated in Figure 3.

Inspired by SimCLR (Chen et al. 2020) and BYOL (Grill et al. 2020), we construct two modality-specific branches for contrastive learning—one for dermoscopic image features and another for clinical feature vectors. Each branch contains a projection head that maps modality-specific embeddings into a shared latent space where semantic alignment is enforced through contrastive objectives. Formally, let x denote an input image and c the corresponding clinical feature vector. A pretrained image encoder $f_v(\cdot)$ extracts visual features $v = f_v(x) \in \mathbb{R}^{d_v}$, while a lightweight MLP $f_c(\cdot)$ encodes the clinical features as $u = f_c(c) \in \mathbb{R}^{d_c}$. These embeddings are then projected into a common embedding space using two multilayer perceptrons (MLPs), h_v and h_c :

$$z_v = h_v(v), \quad z_c = h_c(u), \quad z_v, z_c \in \mathbb{R}^d. \quad (1)$$

To facilitate contrastive learning based on cosine similarity, we apply ℓ_2 normalization:

$$\tilde{z}_v = \frac{z_v}{\|z_v\|_2}, \quad \tilde{z}_c = \frac{z_c}{\|z_c\|_2}. \quad (2)$$

We optimize the alignment using normalized temperature-scaled cross-entropy loss (NT-Xent). The loss for a positive image-clinical pair $(\tilde{z}_{v_i}, \tilde{z}_{c_i})$ is defined as:

$$\mathcal{L}_i = -\log \frac{\exp(\text{sim}(\tilde{z}_{v_i}, \tilde{z}_{c_i})/\tau)}{\sum_{j=1}^N \exp(\text{sim}(\tilde{z}_{v_i}, \tilde{z}_{c_j})/\tau)}, \quad (3)$$

where $\text{sim}(a, b) = a^\top b$ is the cosine similarity and τ is the temperature hyperparameter. We compute a bidirectional contrastive loss by averaging over the batch:

$$\mathcal{L}_{\text{contrastive}} = \frac{1}{2N} \sum_{i=1}^N \left(\mathcal{L}_i^{(v \rightarrow c)} + \mathcal{L}_i^{(c \rightarrow v)} \right). \quad (4)$$

During training, we freeze the parameters of the pretrained image encoder and optimize only the projection

heads, preserving the semantic structure of visual features while enforcing cross-modal alignment. After contrastive pretraining, the projection head is frozen, and a lightweight classification head is trained on the projected image embeddings for downstream melanoma classification. This design enables clinically grounded visual representations that support both accurate prediction and meaningful explanation.

Report Generation with Large Language Model

To improve interpretability and clinical usability, we introduce a medical report generation module that translates model outputs into structured, natural language summaries. We first discretize the computed ABC features into five severity levels. In parallel, we use CLIP-ViT-B/16 to retrieve top-ranked melanoma-related descriptors (e.g., “asymmetric shape,” “blue-gray areas”) based on image features. CLIP is a vision-language model trained via contrastive learning to align image and text embeddings within a shared latent space, enabling the retrieval of semantically meaningful textual descriptions by computing similarity scores in this multimodal space. These quantitative and semantic cues are then combined into prompts and fed into DeepSeek, a domain-adapted large language model, to generate concise, medically grounded diagnostic reports. Table 2 summarizes the CLIP descriptors, and Table 3 shows a representative report.

Model Performance Quantitative Evaluation

Datasets To develop and evaluate our framework, we curated two benchmark dermoscopic datasets from the International Skin Imaging Collaboration (ISIC) Archive—the largest publicly available repository of quality-controlled skin lesion images. We used the ISIC 2018 dataset to train and evaluate the coarse segmentation model, as it provides expert-annotated lesion masks. For the classification and fine segmentation tasks, we selected 3,130 high-quality melanoma and nevus dermoscopic images from ISIC 2020 (excluding cases with ambiguous boundaries or severe artifacts) and supplemented them with 448 unique melanoma images from ISIC 2018 to mitigate class imbalance—resulting in 3,578 images (646 melanoma, 2,932 nevi) for training—and evaluated performance on the official ISIC 2020 validation ($n = 573$) and test ($n = 714$) splits.

Segmentation Performance As the first stage in our segmentation pipeline, we employ the UltraLight VM-UNet—an existing lightweight architecture—as the backbone for coarse lesion mask generation. To validate its suitability for generating reliable pseudo-labels, we benchmarked its segmentation performance on the ISIC2018 dataset under consistent evaluation settings with several representative classical and lightweight models. UltraLight VM-UNet achieves a Dice Similarity Coefficient (DSC) of 0.8909, accuracy of 95.56%, and specificity of 0.9746, outperforming most baselines in overall performance.

Comparison of Vision Models on ISIC2020 ViT and EfficientNet-B2 achieve the highest accuracy (94.26%), with ViT further exhibiting the best precision (88.19%), AUC (0.972), and specificity (97.15%). Given its strong

Lesion Type	CLIP-Matched Clinical Feature
Benign Lesion	Regular symmetric shape, Smooth borders, Uniform color, Light brown color, Small size, Well-defined edges, No ulceration, No bleeding, No satellite lesions, Dome-shaped structure, Single color tone
Melanoma	Asymmetric shape, Irregular borders, Uneven color, Dark brown areas, Black areas, Blue-gray areas, Red areas, White areas, Multiple mixed colors, Fuzzy edges, Ulceration, Bleeding, Crusting, Scaling, Satellite lesions

Table 2: CLIP-matched dermatological attributes used for interpretable lesion classification. These attributes are curated to align with clinical features of benign nevi and malignant melanoma for visual-language modeling.

AI Diagnosis	Suspicious for Melanoma
Quantitative Features	Asymmetry (A): 11.69 – severe asymmetry, strong malignancy indicator Border Irregularity (B2): 2.71 – highly irregular and serrated borders Color Variation (H): 1.02 – minimal hue variation; Saturation Variation (S): 2.17 – mild; Brightness Variation (V): 2.53 – mild
Visual Features	Satellite lesions – indicative of possible metastasis Ulceration – high-risk prognostic feature Dark brown pigmentation – associated with melanocytic proliferation
Risk Assessment	High suspicion due to: severe asymmetry, border irregularity, satellite lesions, ulceration. Low hue variation does not exclude malignancy.

Table 3: Excerpt from AI-generated diagnostic report. Bolded features indicate high-risk indicators.

overall performance and suitability for contrastive learning frameworks like CLIP, ViT is adopted as the baseline for subsequent experiments.

Classification Performance after Feature Alignment To evaluate the effectiveness of our cross-modal contrastive learning framework, we froze the image encoder and projection head, then fine-tuned lightweight classification heads on their embeddings. All experiments were repeated 10 times, with performance reported as mean \pm standard deviation. The ViT-based model achieved the highest overall performance, with an accuracy of $92.79\% \pm 0.57\%$ and an AUC of 0.961 ± 0.004 . Given its high performance and architectural compatibility with the CLIP-based explainability module, ViT is selected as the backbone of our diagnostic pipeline.

Feature Alignment via Contrastive Learning We evaluate feature alignment by analyzing the cosine similarity distributions between paired images and clinical embeddings after contrastive learning. As shown in Fig. 4, positive pairs exhibit significantly higher similarity values, concentrated near 0.8, while negative pairs are distributed near zero. This clear separation demonstrates that contrastive learning effectively aligns cross-modal features and improves discriminative capability in the shared embedding space.

LLM Report Comparison We constructed a single prompt to generate dermatology reports—structured into quantitative feature analysis, visual feature description, consistency and complementarity evaluation, and risk assessment with treatment recommendations—and applied it to **Anthropic Claude, Google Gemini Base, and DeepSeek**. We then used BERTScore (RoBERTa-Large, layer 9) to compare DeepSeek’s outputs against those of Claude and Gemini, reporting Precision, Recall, and F1 to quantify semantic similarity. Given the comparable semantic per-

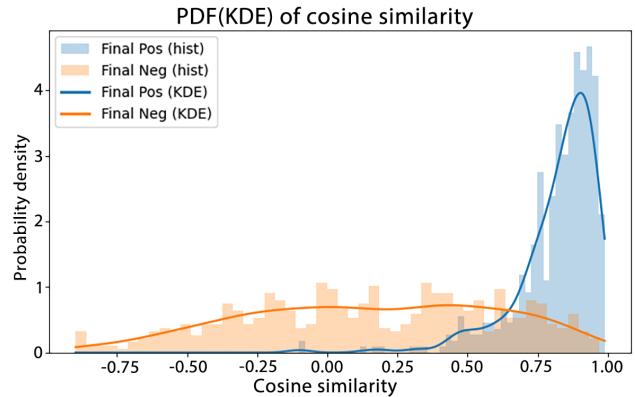


Figure 4: The figure shows cosine similarity distributions of positive and negative pairs after contrastive learning. Positive pairs cluster in the high-similarity region (> 0.75), whereas negative pairs occupy lower values, indicating effective cross-modal alignment.

formance across models and DeepSeek’s superior cost-effectiveness, we selected DeepSeek for subsequent analyses.

Ablation Study To assess the contribution of each component in our explanation pipeline, we conducted an ablation study comparing outputs with individual modules removed (Figure 5). The **CLIP module** supplements the report with additional visual cues that complement the quantified ABC features. When CLIP is removed, the report centers on ABC features and their numerical values, lacking additional contextual factors that may affect the diagnosis. Removing the **clinical explanation module** eliminates critical quantification of ABC features, resulting in text that relies

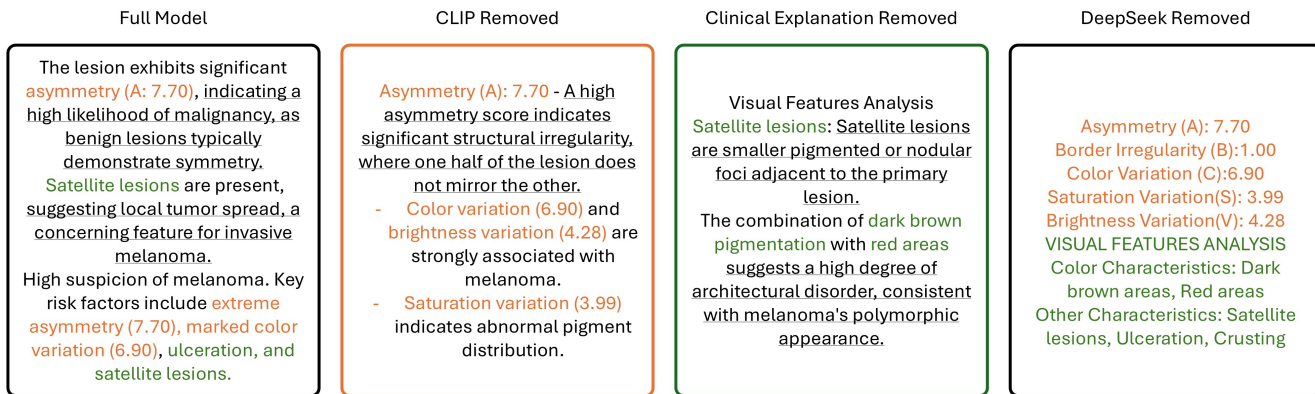


Figure 5: Comparison of the risk assessments generated by models with components removed. Text highlighted in orange are resulted from the clinical explanation component; green are resulted from CLIP with direct prompts; the underline is generated by deepseek. With DeepSeek removed, the report is more fragmented.

Evaluation Dimension	Expert 1	Expert 2	Expert 3	Average Score
1. Consistency with clinical judgment	4.55	4.18	3.81	4.18
2. Usefulness of ABC feature analysis	5.00	4.20	4.00	4.40
3. Report clarity and readability	4.50	4.50	4.50	4.50
4. Interpretability of AI decision-making	5.00	4.40	4.40	4.60
5. Clinical applicability and decision support	4.50	3.50	4.00	4.00

Table 4: Quantitative Expert Evaluation (Likert scale: 1 = strongly disagree; 5 = strongly agree)

solely on visual descriptions without numerical interpretation. The **DeepSeek module** plays a central role in synthesizing these elements into coherent and well-structured narrative reports, ensuring both fluency and contextual integration. Without DeepSeek, the report becomes fragmented, as it lacks the structured presentation that supports narrative coherence. These results underscore the complementary roles of each module in generating comprehensive, interpretable, and clinically meaningful melanoma risk assessments.

User Study with Expert Dermatologists

To comprehensively evaluate the interpretability, usability, and clinical applicability of our proposed multimodal explainable melanoma diagnostic framework, we conducted a structured user study involving three board-certified dermatologists with extensive clinical experience.

Each expert independently reviewed one system-level report and five lesion-level reports. For the system-level report, a broad questionnaire evaluated presentation clarity, AI model trustworthiness, clinical impact, and application. The questionnaire consists of Likert-scale questions assessing each dimension, followed by open-ended and multiple-choice questions to explore rationale, limitations, and applications. For the lesion-level reports, a focused questionnaire with Likert-scale and open questions assessed model accuracy and interpretability. Quantitative responses will be summarized using descriptive statistics (means) across the six evaluated reports. For qualitative responses, thematic analysis will be conducted using an inductive coding approach. Table 4 summarizes expert scores. Experts confirmed that

structured ABC quantification offers clinically relevant and interpretable outputs consistent with physician intuition, establishing ABC reference ranges and diagnostic formats to complete the ABCD framework. They noted that its diagnostic consistency, transparency, and utility can enhance clinician trust and support decision-making.

Conclusion

We introduce **CEFM**, a cross-modal framework for explainable melanoma diagnosis that integrates ViT-based classification with clinically grounded ABC feature extraction and contrastive alignment. A CLIP-driven LLM report module produces structured, transparent diagnostic summaries. Expert evaluation confirms the system’s interpretability and clinical utility, yielding a mean rating of 4.6/5. Additionally, Our experiments on the ISIC dataset demonstrate strong performance, with 92.79 % accuracy and an AUC of 0.961.

Despite its promising results, CEFM has two limitations. First, the exclusion of the **differential structure** criteria from the ABCD rule—due to the lack of expert-level annotated data—limits the completeness of assessment. Second, the current system assumes high-quality dermoscopic inputs and may underperform in cases involving poor imaging conditions or rare lesion subtypes not represented in the training data. Future work will integrate differential structure analysis, support multi-temporal lesion tracking, and validate the framework in prospective clinical studies to enhance robustness across diverse lesion types and imaging conditions.

Acknowledgements

This research was supported by the Ministry of Education, Singapore (Grant ID: RS15/23).

References

- Abir, S. I.; Shoha, S.; Hossain, M. M.; Rahman, S. M.; Saimon, S. I.; Islam, I.; Mamun, M. A. I.; and Khan, N. I. 2024. Deep Learning-Based Classification of Skin Lesions: Enhancing Melanoma Detection through Automated Preprocessing and Data Augmentation. *Journal of Computer Science and Technology Studies*, 6(5): 152–167.
- Almufareh, M. F.; Tariq, N.; Humayun, M.; and Khan, F. A. 2024. Melanoma identification and classification model based on fine-tuned convolutional neural network. *Digital Health*, 10: 20552076241253757.
- Almustafa, K. M. 2025. Predictive modeling and optimization in dermatology: Machine learning for skin disease classification. *Computers in Biology and Medicine*, 189: 109946.
- Chanda, T.; Hauser, K.; Hobelsberger, S.; Bucher, T.-C.; Garcia, C. N.; Wies, C.; Kittler, H.; Tschandl, P.; Navarrete-Dechent, C.; Podlipnik, S.; et al. 2024. Dermatologist-like explainable AI enhances trust and confidence in diagnosing melanoma. *Nature Communications*, 15(1): 524.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PmLR.
- Gamage, L.; Isuranga, U.; Meedeniya, D.; De Silva, S.; and Yogarajah, P. 2024. Melanoma skin cancer identification with explainability utilizing mask guided technique. *Electronics*, 13(4): 680.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33: 21271–21284.
- Hossain, M. I.; Zamzmi, G.; Mouton, P. R.; Salekin, M. S.; Sun, Y.; and Goldgof, D. 2025. Explainable AI for medical data: current methods, limitations, and future directions. *ACM Computing Surveys*, 57(6): 1–46.
- Hussien, M. A.; and Alasadi, A. H. H. 2024. Classification of melanoma skin cancer using deep learning approach. *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, 22(1): 129–137.
- Kayser, M.; Menzat, B.; Emde, C.; Bercean, B.; Novak, A.; Espinosa, A.; Papiez, B. W.; Gaube, S.; Lukasiewicz, T.; and Camburu, O.-M. 2024. Fool me once? contrasting textual and visual explanations in a clinical decision-support setting. *arXiv preprint arXiv:2410.12284*.
- Li, H.; Wang, H.; Sun, X.; He, H.; and Feng, J. 2024. Prompt-guided generation of structured chest X-ray report using a pre-trained LLM. In *2024 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Mallardo, D.; Basile, D.; and Vitale, M. G. 2025. Advances in Melanoma and Skin Cancers.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Sabir, R.; and Mehmood, T. 2024. Classification of melanoma skin Cancer based on Image Data Set using different neural networks. *Scientific Reports*, 14(1): 29704.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Thomas, S. M.; Lefevre, J. G.; Baxter, G.; and Hamilton, N. A. 2021. Interpretable deep learning systems for multi-class segmentation and classification of non-melanoma skin cancer. *Medical Image Analysis*, 68: 101915.
- Toh, S. Y.; Cai, C.; Wang, L. R.; Bai, X.; Ngeow, J.; and Fan, X. 2025. The Effect of Explainable AI and Uncertainty Quantification on Medical Students' Perspectives of Decision-Making AI: A Cancer Screening Case Study. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 1–13.
- Wang, Z.; Liu, L.; Wang, L.; and Zhou, L. 2023. R2gengpt: Radiology report generation with frozen llms. *Meta-Radiology*, 1(3): 100033.
- Xu, H.; and Shuttleworth, K. M. J. 2024. Medical artificial intelligence and the black box problem: a view based on the ethical principle of “do no harm”. *Intelligent Medicine*, 4(1): 52–57.
- Zhang, Z.; Chen, P.; McGough, M.; Xing, F.; Wang, C.; Bui, M.; Xie, Y.; Sapkota, M.; Cui, L.; Dhillon, J.; et al. 2019. Pathologist-level interpretable whole-slide cancer diagnosis with deep learning. *Nature Machine Intelligence*, 1(5): 236–245.