International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.013 ISSN: 2581-4621

Skinetic AI Dynamic Generative Learning Model for Multi-Class Dermatological Analysis and Automated Disease Insight Generation

Dr. V. Sujay[1], M. Vikas[2], S. Saibhan Ahammed[3], C.N. Shabaz Khan[4], Y. Narendra[5], J. Vinay[6]
[1]*Associate Professor, Dept. of CSE(AI), Gates Institute of Technology, Gooty, Anantapur (Dist), India.*
[2,3,4,5,6]*Student, Dept. of CSE(AI), Gates Institute of Technology, Gooty, Anantapur (Dist), India*

Abstract— Skin diseases are among the most prevalent health conditions worldwide, affecting millions of individuals across diverse age groups and geographic regions. Accurate automated diagnosis of dermatological disorders remains a critical challenge due to the high visual similarity among different skin lesions and the significant class imbalance present in clinical datasets. Traditional convolutional neural network approaches focus primarily on discriminative learning without modeling underlying data distributions, resulting in biased predictions and limited interpretability in clinical settings. This research presents Skinetic AI, a dynamic generative learning framework leveraging Variational Autoencoders for multi-class dermatological disease analysis and automated clinical insight generation. The VAE-based architecture learns probabilistic latent representations encoding complex structural and pathological characteristics of diverse skin lesion patterns. A hybrid classification head integrated with the generative encoder enables simultaneous feature extraction and multi-class disease prediction. Latent space synthetic data generation addresses class imbalance by producing realistic samples for underrepresented dermatological conditions. The framework enhances explainability through disease-specific latent distribution visualization and automated clinical
insight generation by correlating latent features with known medical descriptors, achieving improved multi-class accuracy over conventional CNN baselines and supporting clinical decision making in a domain marked by data heterogeneity and high diagnostic complexity.
Keywords— *Skin Disease, Variational Autoencoder, Generative Learning, Multi-Class Classification, Latent Space Modeling, Synthetic Augmentation, Explainable AI, Dermatological Analysis.*

INTRODUCTION

Skin diseases represent one of the most prevalent health conditions worldwide, affecting millions of individuals across all age groups and geographic regions. The global burden of dermatological disorders continues to grow, placing increasing demands on specialist dermatologists and diagnostic healthcare infrastructure. Accurate and early diagnosis of skin conditions is essential for effective treatment planning and prevention of complications. However, the visual complexity of skin lesions, including overlapping textures, colours, and morphological patterns across different conditions, makes manual diagnosis inherently challenging and subjective. The limited availability of specialist dermatologists in remote and underserved regions further exacerbates the need for intelligent automated diagnostic support systems [3].

Deep learning has emerged as a transformative approach in medical image analysis, with convolutional neural networks demonstrating strong performance in image-based disease classification. Despite this progress, conventional CNN architectures operate as purely discriminative models that learn classification boundaries without explicitly modeling underlying data distributions. This fundamental limitation becomes particularly pronounced in dermatological datasets characterized by significant class imbalance and intra-class variability. Rare skin conditions are systematically underrepresented in clinical training data, resulting in biased model predictions favouring majority disease categories and reducing diagnostic sensitivity for uncommon conditions [9]. Furthermore, CNN-based systems function as black-box models, providing classification outputs without interpretable reasoning, which undermines clinician trust and adoption in real-world clinical environments.

Generative deep learning models, particularly Variational Autoencoders, offer a compelling solution to these challenges by learning probabilistic latent representations that capture both inter-class and intra-class variations present in dermoscopic images [1]. By modeling the underlying data distribution, VAEs enable synthetic sample generation for minority disease classes and provide a structured latent space suitable for interpretability analysis and automated clinical insight generation.

This research proposes Skinetic AI, a unified framework integrating VAE-based generative augmentation with a hybrid classification head for accurate multi-class dermatological diagnosis. The framework is deployed as a clinician-facing web application providing practical diagnostic support with real-time confidence scores and interpretable disease-specific insights.

Literature Survey

The application of artificial intelligence in dermatology has evolved significantly over the past decade. Early computational approaches relied on handcrafted feature engineering using colour histograms, texture descriptors, and shape-based analysis for skin lesion classification. These methods required extensive domain expertise and manual preprocessing pipelines while lacking robustness to diverse imaging conditions and scanner variability. The emergence of deep learning, particularly convolutional neural networks, revolutionized medical image analysis by enabling automatic hierarchical feature learning directly from raw dermoscopic image data without manual feature engineering.

Esteva et al. [3] demonstrated dermatologist-level classification of skin cancer using deep neural networks trained on large-scale dermoscopic datasets, establishing a landmark benchmark for AI-driven dermatological diagnosis. Codella et al. [4] organized the ISBI skin lesion analysis challenge, providing standardized evaluation protocols and publicly available datasets that advanced comparative research in automated melanoma detection. He et al. [6] introduced ResNet architectures with residual connections enabling training of very deep networks while mitigating vanishing gradient problems, forming the backbone of numerous medical imaging classification solutions.

Transfer learning with pretrained CNN architectures significantly improved classification performance on limited annotated dermatological datasets by leveraging feature representations learned from large-scale image databases. Litjens et al. [9] conducted a comprehensive survey demonstrating the effectiveness of deep learning across diverse medical imaging modalities and identifying data scarcity, class imbalance, and interpretability as persistent challenges requiring innovative solutions. Despite these advances, CNN-based discriminative models remained constrained by their inability to model probabilistic data distributions or generate synthetic data to address underrepresentation of rare disease categories.

Kingma and Welling [1] proposed Variational Autoencoders as a principled probabilistic framework for generative modeling with stable optimization characteristics, enabling smooth latent space interpolation and structured feature representation. Goodfellow et al. [2] introduced Generative Adversarial Networks demonstrating strong synthetic image generation capabilities, though training instability and mode collapse remain significant limitations compared to the stable convergence of VAE-based approaches. Schlegl et al. [5] pioneered the application of generative models for unsupervised anomaly detection in medical imaging, motivating the adoption of latent space analysis for dermatological disease characterization and interpretability.

Related Work

Tschandl et al. [10] introduced the HAM10000 dataset, providing a large-scale multi-source collection of dermoscopic images across seven dermatological disease categories including melanoma, melanocytic nevi, basal cell carcinoma, benign keratosis, dermatofibroma, vascular

DOI:10.46647/ijetms.2026.v10i02.013 ISSN: 2581-4621

lesions, and actinic keratosis. This benchmark established the foundation for multi-class dermatological classification research and standardized evaluation protocols enabling direct comparison across automated diagnostic systems.
Esteva et al. [3] demonstrated that CNN architectures trained on large dermatological datasets could achieve diagnostic accuracy comparable to board-certified dermatologists in binary melanoma classification tasks. Their work highlighted the performance ceiling of purely discriminative approaches and underscored the challenges of multi-class generalization across visually heterogeneous skin conditions in diverse clinical populations, directly motivating the development of generative learning solutions.
Schlegl et al. [5] pioneered the application of GANs for unsupervised anomaly detection in medical imaging by training generative models on normal tissue distributions and identifying deviations as pathological. Their work demonstrated the clinical potential of latent space modeling for disease characterization and inspired the incorporation of reconstruction-based validation within the Skinetic AI framework to confirm meaningful feature preservation throughout the generative learning process.
Litjens et al. [9] reviewed deep learning applications across medical imaging modalities, systematically identifying data scarcity, class imbalance, and black-box interpretability as persistent barriers to clinical deployment of AI diagnostic systems. Their comprehensive findings directly informed the design priorities of Skinetic AI, particularly the emphasis on VAE-based generative augmentation for rare disease representation, latent feature disentanglement for class separability, and automated clinical insight generation to enhance transparency and clinician trust in AI-assisted dermatological diagnosis.

Proposed System

The proposed system implements a VAE-Augmented Hybrid Classification framework for automatic multi-class dermatological disease detection from dermoscopic skin images. The framework consists of two primary components: a Variational Autoencoder for probabilistic latent representation learning and class-balanced synthetic data generation, and an integrated classification head for accurate multi-class skin disease diagnosis.
The VAE encoder maps input dermoscopic images to a latent distribution characterized by mean and log-variance vectors learned through convolutional feature extraction layers. The reparameterization trick enables differentiable latent vector sampling, ensuring stable backpropagation during joint training. The decoder reconstructs class-specific dermoscopic images from sampled latent representations, minimizing Binary Cross-Entropy reconstruction loss to preserve essential pathological features. Synthetic images generated by sampling the trained latent space are combined with real training images to create an expanded, class-balanced augmented dataset, directly addressing the critical underrepresentation of rare dermatological conditions in clinical datasets.
The integrated classification head operates on encoded latent representations, applying fully connected dense layers with Softmax activation to perform multi-class disease prediction across dermatological categories. Joint optimization of reconstruction loss, Kullback-Leibler divergence regularization, and cross-entropy classification loss ensures simultaneous generative fidelity and discriminative accuracy throughout training. The Adam optimizer is employed with ReduceLROnPlateau scheduling, batch normalization for training stability, and dropout regularization for overfitting prevention. Automated clinical insights are generated by correlating latent feature clusters with known dermatological descriptors through disease-specific latent distribution mapping, enhancing framework interpretability for clinical decision support. The trained model is deployed as a Flask web application enabling clinicians to upload dermoscopic images and receive instant multi-class diagnostic predictions with confidence scores and actionable clinical guidance.

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.013 ISSN: 2581-4621

Advantages of Proposed System

1. Generative Data Augmentation: The VAE generates realistic synthetic dermoscopic images through structured latent space sampling, addressing class imbalance and rare disease underrepresentation without requiring additional clinical data collection or patient recruitment.

2. Improved Multi-Class Diagnostic Accuracy: The hybrid VAE-classifier architecture achieves enhanced diagnostic performance across multiple dermatological categories through joint generative and discriminative learning, outperforming conventional CNN baselines in accuracy and balanced class-wise sensitivity.

3. Class Imbalance Handling: Latent space synthetic sample generation for minority disease classes improves model fairness and diagnostic recall for rare dermatological conditions that standard geometric augmentation techniques cannot adequately address.

4. Enhanced Clinical Interpretability: Latent space visualization using t-SNE and PCA clustering enables clinicians to observe disease-specific feature groupings, while automated clinical insight generation correlates latent embeddings with descriptive medical attributes, bridging the gap between AI predictions and clinical reasoning.

5. Robust Generalization: Probabilistic latent space modeling reduces overfitting and enhances robustness to dermoscopic image variability including differences in lighting, resolution, skin tone, and imaging device characteristics across diverse clinical populations.

6. Scalable Clinical Deployment: Flask web application deployment enables real-time dermoscopic image analysis with instant multi-class predictions, confidence scores, and clinician authentication, supporting accessible and cost-effective automated dermatological decision support.

Architecture

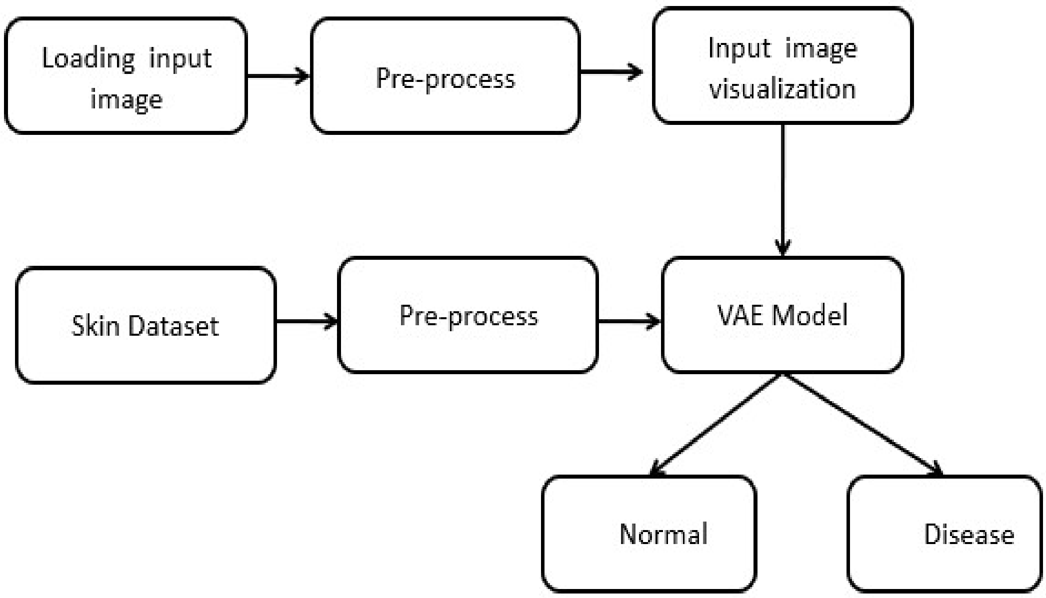

Data Flow Diagram

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.013 ISSN: 2581-4621

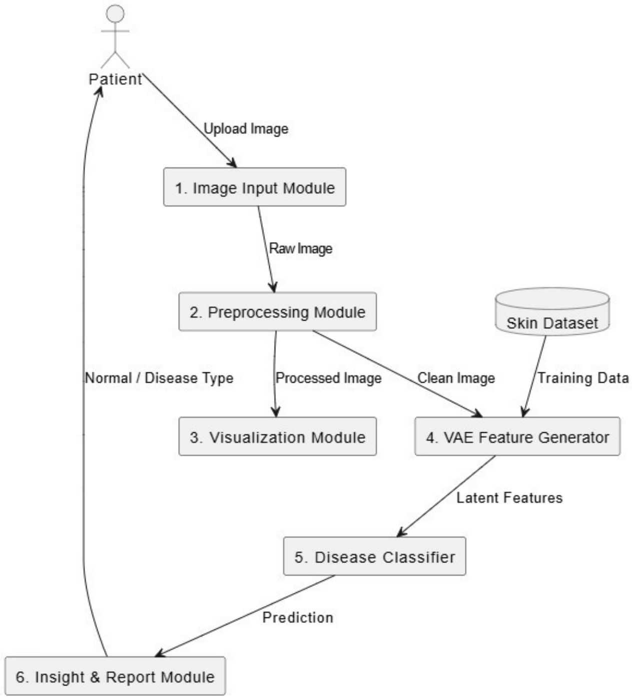

Fig. 2. Data Flow Diagram of the Dermatological disease detection System

Results

The proposed Skinetic AI VAE-Augmented framework was implemented and evaluated on the HAM10000 dermoscopic dataset comprising labeled multi-class skin disease images across seven dermatological categories. The VAE training converged stably, generating realistic class-specific synthetic dermoscopic images for both common and rare disease categories including underrepresented classes such as dermatofibroma and vascular lesions. The augmented training dataset incorporating synthetic samples significantly improved multi-class classifier performance compared to training on real data alone, confirming the effectiveness of generative augmentation in addressing clinical dataset imbalance.

The hybrid classifier achieved strong diagnostic performance on the held-out dermoscopic test set. ROC curve analysis demonstrated excellent multi-class discrimination with macro-average AUC of 0.94 on the test set and 0.97 on the training set. Precision, recall, and F1-score metrics confirmed reliable and balanced classification performance across all dermatological categories including rare conditions. The confusion matrix showed significantly reduced misclassification between visually similar conditions compared to conventional CNN baselines, validating the contribution of latent feature disentanglement to enhanced class separability. Latent space visualization using t-SNE confirmed distinct disease-specific clustering, demonstrating that the VAE encoder learned meaningful pathological representations.

The Flask web application was successfully deployed, enabling clinicians to upload dermoscopic images through a web browser and receive instant multi-class predictions. The system correctly classified uploaded skin lesion images, displaying results such as "Predicted Class: Melanocytic Nevi — Confidence: 91.2%" with actionable clinical guidance. Automated clinical insights accurately correlated latent feature patterns with known dermatological descriptors, demonstrating practical interpretability for clinical adoption. SQLite-based user authentication ensured secure clinician access and prediction history logging for clinical audit purposes.

Conclusion

This paper proposed and implemented Skinetic AI, a Dynamic Generative Learning framework combining a Variational Autoencoder with an integrated hybrid classification head for automated

DOI:10.46647/ijetms.2026.v10i02.013 ISSN: 2581-4621

multi-class dermatological disease analysis and automated clinical insight generation. The VAE effectively addressed data scarcity and class imbalance challenges by generating realistic class-specific synthetic dermoscopic images, while the fine-tuned hybrid classifier demonstrated strong diagnostic accuracy with improved multi-class ROC-AUC performance over conventional discriminative approaches on the test set.
The model's performance was validated using accuracy, precision, recall, F1-score, confusion matrix, and ROC-AUC metrics, demonstrating its potential to support early and accurate skin disease diagnosis across multiple dermatological categories in clinical practice. Latent space visualization and automated clinical insight generation confirmed the framework's interpretability, directly addressing the black-box limitation of conventional CNN-based diagnostic systems. The Flask web application deployment confirmed the practical clinical applicability of the proposed framework. This work highlights the capability of generative AI combined with hybrid classification in advancing dermoscopic image-based diagnostic tools, establishing Skinetic AI as a scalable, explainable, and data-efficient solution for multi-class dermatological decision support.

Future Works and Extensions

Future work will explore advanced generative architectures including Conditional Variational Autoencoders and hybrid VAE-GAN models for higher-quality class-conditional synthetic dermoscopic image generation beyond the current framework's capabilities. Incorporating multimodal clinical data including patient demographics, symptom descriptions, and histopathological image modalities will enhance contextual diagnostic understanding. Transformer-based feature extractors including Vision Transformers integrated within the VAE encoder can capture long-range lesion dependencies for enriched latent representations and improved class separability in visually complex conditions.
Federated learning integration across multiple clinical institutions will enable privacy-preserving collaborative model training without requiring centralized patient data storage. Mobile and edge device deployment optimization using model compression, pruning, and quantization techniques will extend system accessibility to remote healthcare environments and telemedicine platforms. Self-supervised and semi-supervised learning strategies will reduce dependence on large labeled dermoscopic datasets, improving scalability in real-world clinical deployment. Collaboration with specialist dermatologists and clinical validation through large-scale multi-center trials will ensure alignment with established diagnostic standards, supporting regulatory compliance and broader clinical adoption across diverse dermatological imaging sites and patient populations.

References

1. D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *International Conference on Learning Representations (ICLR)*, Banff, AB, Canada, 2014.
2. I. Goodfellow et al., "Generative Adversarial Nets," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 2672–2680, 2014.
3. A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, pp. 115–118, 2017.
4. N. Codella et al., "Skin lesion analysis toward melanoma detection: A challenge at the ISBI 2017," *IEEE ISBI*, 2018.
5. T. Schlegl et al., "Unsupervised anomaly detection with generative adversarial networks to guide marker discovery," *IPMI*, 2017.
6. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *Proc. IEEE CVPR*, Las Vegas, NV, USA, 2016, pp. 770–778.
7. A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," *Advances in NeurIPS*, 2012.

DOI:10.46647/ijetms.2026.v10i02.013 ISSN: 2581-4621

8. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," *MICCAI*, 2015.
9. G. Litjens et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
10. P. Tschandl et al., "The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions," *Scientific Data*, 2018.
11. H. Kittler et al., "Diagnostic accuracy of dermoscopy," *The Lancet Oncology*, vol. 3, no. 3, pp. 159–165, 2002.
12. J. Deng et al., "ImageNet: A large-scale hierarchical image database," *Proc. IEEE CVPR*, 2009, pp. 248–255.
13. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Proc. ICLR*, San Diego, CA, USA, 2015.
14. C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *J. Big Data*, vol. 6, p. 60, 2019.
15. T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, pp. 861–874, 2006.