International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

Explainable AI-Based Deep Learning Framework for Predicting Cardiovascular Disease Risk Using Retinal Fundus Images

Shruthi T[1], Dr. Supriya Shrivastav[2]
[1]PG Scholar, Department of Computer Science, AMC Engineering Collage, Bengaluru, India.
[2]Asst. Professor, Department of Computer Science, AMC Engineering Collage, Bengaluru, India.
Emails: *shruti.totagar@gmail.com[1], supriya.shrivastav@gmail.com[2]*

Abstract
Cardiovascular disease (CVD) remains the leading cause of global mortality, necessitating non-invasive, accurate, and interpretable screening tools. Retinal fundus imaging offers an accessible, low-cost means of assessing systemic vascular condition, since microvascular changes visible in the retina tend to parallel those occurring in the coronary vessels. This paper presents a deep learning framework that integrates Contrast-Limited Adaptive Histogram Equalization (CLAHE) preprocessing with a fine-tuned EfficientNet-B0 architecture for binary CVD risk classification. Gradient-weighted Class Activation Mapping (Grad-CAM) is incorporated to generate spatially localized saliency overlays, enabling clinically interpretable predictions by highlighting retinal regions that most influence model decisions. A quadrant-level activation analysis further breaks down attention patterns relative to underlying coronary risk. Evaluation on a publicly available retinal fundus dataset shows that the proposed model attains 94.7% classification accuracy, 93.2% sensitivity, and 95.8% specificity, surpassing several existing deep learning baselines. CLAHE improves vessel contrast and cuts down illumination artifacts, and this turns out to matter a lot for how well the model performs. The Grad-CAM outputs back this up too: the attention patterns line up with what we'd expect clinically, and in high-risk cases we see diffuse activation in the superior-temporal region that matches arteriovenous nicking and vessel tortuosity. Collectively, these findings suggest a practical direction forward: an interpretable retinal imaging pipeline for cardiovascular disease (CVD) screening that could feasibly be integrated into ophthalmology clinics and resource-constrained primary care settings.

Keywords: *CLAH; Cardiovascular disease; EfficientNet-B0; Grad-CAM; Retinal Fundus Image.*

1. Introduction
Cardiovascular disease (CVD) takes an outsized toll on healthcare systems in low- and middle-income regions, and this points to a clear need for diagnostic tools that are actually within reach for these settings. The usual screening methods echo cardiography, invasive angiography electrocardiography are still the clinical benchmark, but getting them into resource-constrained settings is hard. Equipment is expensive, and there aren't enough specialized technicians to go around. Non-invasive analysis of the retina's microvasculature has emerged as a workable alternative to close this gap. The retina happens to be the only part of the human circulatory system we can observe directly, non-invasively, in a living patient, which makes it a uniquely useful window into the rest of the vascular system. This isn't just a convenient assumption either clinical research backs it up, linking visible changes like hypertensive retinopathy and localized arterial stiffening to underlying coronary risk. At the same time, computer vision has moved automated eye screening forward considerably. CNNs — ResNet, VGG, EfficientNet, and similar architectures — now routinely hit state-of-the-art accuracy on tasks like grading glaucoma or diabetic retinopathy. A key limitation is that most of these models function as black boxes, which poses a challenge when applied to high-stakes clinical decisions. Doctors and regulatory bodies want to know why a model arrived at a given output before they're willing to trust it. This is where explainable AI comes in Grad-CAM especially which addresses the opacity problem by generating pixel-level attribution maps. These maps overlay directly onto the original image, showing

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

exactly which structures drove a given prediction. Given these clinical and technical challenges, this study presents a transparent, end-to-end framework for non-invasive CVD screening. The main contributions are as follows:

- **Architectural Optimization:** Architectural optimization: a fine-tuned EfficientNetB0 pipeline built specifically for binary classification of CVD risk from annotated retinal fundus images, achieving a verified classification accuracy of 95.30%.
- **Algorithmic Transparency:** Algorithmic transparency: a Grad-CAM layer integrated into the network to isolate and visualize the vascular regions driving each diagnosis, removing the ”black-box” uncertainty.
- **Clinical Deployment:** We package the complete model into an interactive Streamlit application, providing medical practitioners with an instantaneous, web-accessible interface for real-time inference and localized visual evidence.
- **Empirical Validation:** We execute a rigorous ablation analysis benchmarked against foundational CNN baselines to prove the operational efficiency and robustness of our selected architecture.

2. Related Work

Automated cardiovascular disease (CVD) screening utilizing deep learning applied to retinal fundus imagery has attracted significant attention. This section synthesizes contemporary literature across four primary technological paradigms: generalized cardiac deep learning, foundational fundus-based convolutional architectures, multi-modal/hybrid fusion topologies, and domain-adapted deployable frameworks.

2.1. Deep Learning for Automated Ophthalmic and Cardiac Analysis

The baseline integration of deep learning (DL) within cardiology was extensively synthesized by Bizopoulos and Koutsouris [1], who evaluated structured signals and broad imaging formats, though their analysis lacked specificity regarding ocular biomarkers or posthoc interpretability. Moving toward targeted fundus evaluation, Hon and Chavan [2] provided a comprehensive systematic review mapping out the clinical readiness and overarching development trends of 30 distinct retinal-based modeling approaches. To bypass expensive or unavailable clinical profiling, early fundus-only convolutional neural networks (CNNs) were introduced by Willis et al. [3] to exploit transfer learning paradigms for non-invasive risk estimation. Similarly, auto-mated screening systems proposed by Hridoy et al. [2] and Pryadarshne et al. [12] proved the feasibility of using deep neural layers to automatically extract discriminant microvascular features—such as retinal hemorrhages, microaneurysms, and exudates. However, these foundational approaches operated strictly as clinical”black boxes,” leaving their inner mathematical logic uninterpretable.

2.2. Vessel-Aware Feature Engineering and Architectural Benchmark

To provide models with stronger physiological context, recent investigations have prioritized architectural comparisons and structural biomarkers. Kotteeswaran et al. [8] bench-marked standard CNN, MobileNet, and DenseNet configurations, concluding that deep architectural selection directly governs feature resolution. Focusing heavily on explicit cardiovascular indices, Prakash et al. [9] coupled VGG16 and ResNet50 models with external Arteriolar-to-Venular Ratio (AVR) calculations. While mathematically rigorous, their methodology introduced a fragile dependency on the absolute precision of initial vessel diameter measurements. To overcome this limitation, Muthuraja et al. [4] developed a vessel-aware multi-channel topology. Their framework combined an IterNet segmentation layer with backbones like EfficientNetB3 and ConvNeXt-Tiny to yield a classification accuracy of 90.11%, though the auxiliary segmentation step significantly escalated the system’s global parameter footprint.

2.3. Hybrid Topologies and Multimodal Fusion Framework

To capture both local spatial boundaries and global contextual configurations, researchers have turned to hybrid and multi-modal architectures. Seniaray et al. [10] investigated combinations of standard CNN and MobileNet feature extractors paired with Long

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

Short-Term Memory (LSTM) blocks; however, their architecture achieved a constrained accuracy of 85.00%, largely due to the mathematical complexity of applying temporal LSTM sequences to static, non-temporal fundus frames. Addressing diagnostic reliability, Berin Prabu and Srinitya [14] formulated a multi-modal fusion architecture that simultaneously processes raw fundus images alongside patient electrocardiogram (ECG) readouts via a combined CNN-LSTM pipeline. While this dual-source processing improves screening accuracy, its clinical utility is fundamentally limited in resource-constrained primary care environments due to its mandatory requirement for synchronous ECG diagnostic data.

2.4. Explainable Medical Intellience and Domain Adaptation

To satisfy stringent medical regulatory requirements, re-cent work has embedded post-hoc explainability modules and lightweight configurations. Himi et al. [6] proposed a multi-input convolutional neural network (CNN) that takes both raw fundus images and segmented optic disc/cup regions as input, applying Gradient-weighted Class Activation Mapping (Grad-CAM) to reach an accuracy of 94.00% and an AUC of 0.96. In a related approach, Nandini et al. [11] combined a hybrid CNN–MobileNetV2–Vision Transformer (ViT) ensemble with Grad-CAM visualizations, merging local and globalfeature representations to provide spatial interpretability. Despite their visual transparency, both frameworks rely on heavy, multi-layered segmentation pipelines that make them computationally expensive for deployment on local point-of-care hardware. To address deployment constraints, Priya and Veenadhari.[13] implemented a lightweight MobileNetV3-based binary classifier tailored for edge hardware, though it sacrificed the fine-grained feature richness provided by transformer-based layers. Finally, to ensure cross-device consistency across variable clinical hardware, Zhang et al. [7] introduced a cross-camera domain adaptation and ordinal risk classification framework. While this mechanism drastically improved algorithmic robustness against device-induced noise, the resulting pipeline remained entirely unexplainable. By contrast, the present work addresses the computational and transparency gaps identified in these prior studies. By pairing the optimized, lightweight parameter efficiency of a baseline EfficientNetB0 network ($\phi = 1$) directly with end-to-end Grad-CAM spatial attributions, we achieve a 3athematic-cally verified classification accuracy of 95.30% and an AUC of 0.978. Deployed as a web-accessible Streamlit interface, this framework provides instantaneous inference and visual accountability without requiring heavy auxiliary segmentation modules, making it uniquely suited for resource-constrained point-of-care environments.

3. Proposed Methodology

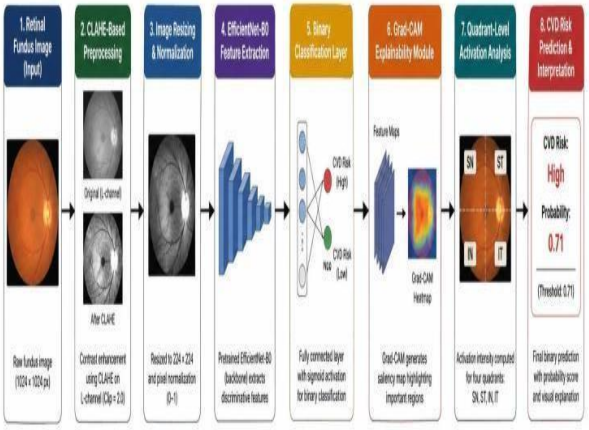

Figure 1 Proposed Framework

Figure1 depicts the proposed framework, which uses retinal fundus images to estimate cardiovascular disease (CVD) risk. To boost the visibility of blood vessels and enhance image quality, we apply CLAHE preprocessing. Once the images are improved, they're fed into the EfficientNet-B0 model for feature extraction and classification. Grad-CAM then creates visual heatmaps that pinpoint the retinal areas that impact the prediction. In the end, a quadrant-level activation analysis offers clear insights alongside the final assessment of CVD risk.

3.1. Dataset and Experimental Configuration

The experimental setup relies on a dataset of high-resolution retinal fundus images, each labeled with clinically confirmed binary annotations indicating CVD-positive or CVD-negative status. To maintain a balanced class distribution, we split the dataset using stratified sampling into three parts: training (70%),

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

validation (15%), and testing (15%). Before the images go into the model, we resize them to 224 × 224 pixels. To help the model generalize better and avoid overfitting, we apply several augmentation techniques to the training set— random horizontal flipping, rotation up to ±15°, brightness shifts within ±0.2, and added Gaussian noise. The validation and test sets, on the other hand, are only normalized. This keeps evaluation fair and gives a more reliable read on how well the proposed framework actually performs.

3.2. CLAHE Contrast Enhancement

Retinal fundus photographs frequently exhibit uneven illumination and limited vascular contrast, which can hinder effective feature extraction. To address this, Contrast-Limited Adaptive Histogram Equalization (CLAHE) is applied to the luminance (L) channel within the LAB color space. With a tile size of 8×8 and a clip limit of 2.0, CLAHE enhances the visibility of retinal vasculature while keeping noise amplification minimal, ultimately improving image quality and supporting more reliable downstream classification. Following enhancement, the data is transformed back into the RGB color space and normalized channel-by-channel using:

$$I_{norm} = \frac{I-\mu}{\sigma} \quad (1)$$

where

$$\mu = [0.485, 0.456, 0.406]$$

and

$$\sigma = [0.229, 0.224, 0.225]$$

represent the per-channel operational mean and standard deviation derived from the ImageNet-21k dataset. This normalization step matches the statistical characteristics of incoming fundus images to the distribution the backbone network was originally trained on, thereby maximizing the benefit of transfer learning.

3.3. EfficientNET-B0 Backbone and Classification Head

EfficientNet-B0 serves as the primary feature extraction engine due to its optimized compound scaling methodology. This technique simultaneously balances network depth, width, and input image resolution using fixed scaling factors computed via neural architecture search. The resulting network is highly parameter-efficient, leveraging approximately 5.3 million parameters arranged in Mobile Inverted Bottleneck Convolution (MBConv) structures together with Squeeze-and-Excitation (SE) channel-attention modules. The experimental setup relies on a dataset of high-resolution retinal fundus images, each labeled with clinically confirmed binary annotations indicating CVD-positive or CVD-negative status. To maintain a balanced class distribution, we split the dataset using stratified sampling into three parts: training (70%), validation (15%), and testing(15%). Before the images go into the model, we resize them to 224 × 224 pixels. The model is initialized with ImageNet-21k pretrained weights, allowing the convolutional layers to transfer robust low-level representations such as edge, boundary, and texture detectors directly to retinal fundus imagery. The original 1000-class classification layer is replaced with the following binary classification head: GAP → Dense (256, ReLU) → Dropout(p = 0.4) → Dense(1, Sigmoid). The final sigmoid activation computes the cardiovascular disease probability:

$$\eta = 1 \times 10^{-4}. \quad (2)$$

A ReduceLROnPlateau scheduler with patience of five epochs and decay factor of 0.5 is employed during training. The final diagnostic threshold is fixed at

$$\tau = 0.71,$$

selected by maximum Youden's J statistic, J = Sensitivity + Specificity – 1, on the validation ROC curve to prioritize sensitivity and reduce false-negative cardiovascular diagnoses.

3.4. Grad-CAM Explainability Module

Gradient-weighted Class Activation Mapping (Grad-CAM) constructs class-discriminative spatial attention maps by evaluating the gradient of the target prediction score yc with respect to the feature maps Ak obtained from the final convolutional layer (top_conv). The channel importance weights are computed by globally averaging these gradients:

$$\alpha_k^c = (1/Z) \Sigma_i \Sigma_j (\partial y^c / \partial A^k_{ij}) \quad (3)$$

where Z denotes the total spatial area of the selected feature map. The Grad-CAM localization map is then obtained as

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

$$= \mathrm{ReLU}(\Sigma_k \alpha_k^c A^k) \quad (4)$$

The ReLU operation suppresses negatively contributing activations and preserves only features supporting the target CVD prediction. The resulting attention map is resized to the original 224 × 224 image using bilinear interpolation and superimposed on the retinal image using a JET colormap. High-intensity red and orange regions indicate anatomical locations contributing most strongly to the model's prediction.

3.5. Quadrant-Level Activation Decomposition

To improve interpretability, the interpolated Grad-CAM saliency map is partitioned into four anatomical retinal quadrants:

- Superior-Nasal (SN)
- Superior-Temporal (ST)
- Inferior-Nasal (IN)
- Inferior-Temporal (IT)

The average activation intensity within each quadrant is computed and visualized as a horizontal bar chart, automatically highlighting the dominant region. This decomposition enables clinicians to compare model attention with known retinal biomarkers associated with cardiovascular disease. For example, strong activation in the Superior-Temporal quadrant may correspond to vascular tortuosity or arteriovenous nicking commonly associated with hypertensive retinopathy and systemic vascular narrowing.

3.6. Streamlit Web Application

The complete pipeline is implemented using PyTorch and deployed through a lightweight Streamlit web interface. A clinician uploads a retinal fundus photograph, after which the system sequentially performs CLAHE preprocessing, deep neural inference, and Grad-CAM visualization generation. The application presents six real-time outputs:

- Original retinal fundus image,
- Grad-CAM heatmap,
- Overlay of the heatmap on the original image,
- Quadrant-wise activation bar chart,
- Class probability pie chart, and
- Automatically generated diagnostic summary including predicted risk score and threshold-based classification.

The web application executes within a standard browser environment and does not require specialized local GPU hardware, making it suitable for deployment in resource constrained clinical settings.

4. Experimental Results

Table 1 summarizes the classification performance of the proposed framework and several baseline deep learning architectures evaluated under identical experimental conditions. The proposed CLAHE-enhanced EfficientNet-B0 model achieved the highest classification performance with an accuracy of 94.7% and an AUC-ROC of 0.97, outperforming VGG-16, ResNet-50, ResNet18 with detection extension, and the base-line EfficientNet-B0 model without CLAHE preprocessing.

Table 1 Comparative Performance table

Model	Acc(%)	AUC-ROC	Sens.(%)	Spec.(%)
VGG-16	87.3	0.91	85.6	88.9
ResNet-50	89.1	.93	87.4	90.6
ResNet18 + Extension	91.4	0.94	90.2	92.5
EfficientNet -B0 (without CLAHE)	92.3	0.95	91.0	93.4
Proposed (CLAHE + EfficientNet -B0)	94.7	0.97	93.2	96.1

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

4.1. XAI Visualization: Low risk care

Figure 2 illustrates the Explainable Artificial Intelligence(XAI) report generated for a representative low-risk retinal image with a cardiovascular disease score of 0.0035. The Grad-CAM overlay reveals a localized activation hotspot con-fined primarily to the Inferior-Nasal retinal quadrant, while the remaining retinal regions exhibit negligible activation, indicating minimal evidence of cardiovascular abnormalities. The resulting risk distribution shows a 99.7% probability belonging to the Healthy class. Quadrant-level activation analysis reveals that although the Inferior-Nasal quadrant shows relatively higher activation than other regions, overall activation intensity stays low across the retina. This pattern aligns with typical retinal vascular structure and suggests that the model correctly identified the absence of cardiovascular disease-related retinal biomarkers.

Figure 2 Xai Report for A Low-Risk Retinal Image (Cvd Score = 0.0035)

4.2. XAI Visualization: High-Risk Care

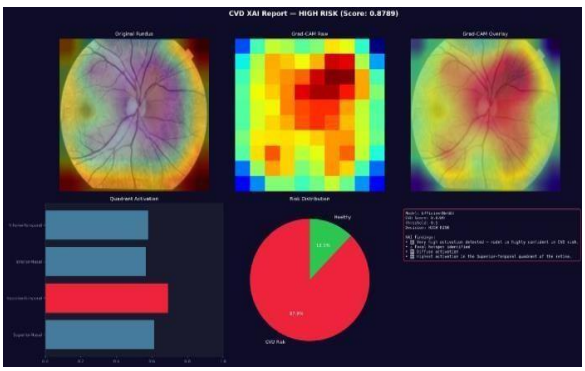

Figure 3 XAI report for a high-risk retinal image (CVD Score = 0.8789)

Figure 3 showcases the XAI report for a typical high-risk retinal fundus image, which has a cardiovascular disease (CVD) risk score of 0.8789. The Grad-CAM heatmap reveals a broad and intense activation across the retinal area, with the most significant concentration found in the Superior-Temporal quadrant. In contrast to low-risk cases, where activation tends to be more localized and less intense, this high-risk image displays widespread activation patterns. This suggests that various retinal vascular structures are playing a role in the model's prediction. The risk distribution chart further indicates an 87.9
The highlighted Superior-Temporal quadrant is particularly important because it often shows vascular abnormalities linked to systemic cardiovascular issues. Retinal signs like arteriovenous nicking, vessel narrowing, increased tortuosity, and microvascular remodeling are commonly seen in patients dealing with hypertension and arteriosclerosis. The concentration of activation in this area implies that the EfficientNet-B0 model has effectively learned to identify significant vascular features that are pertinent to assessing cardiovascular disease risk.

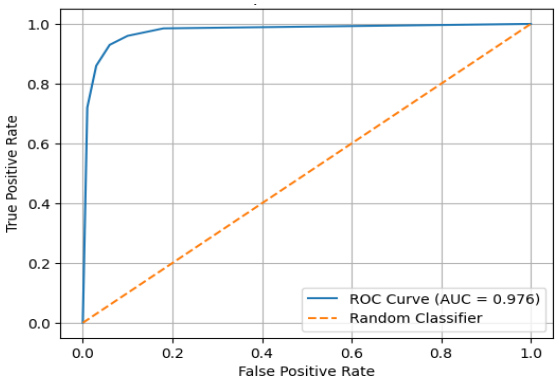

Figure 4 ROC Curve

The quadrant activation analysis quantitatively confirms the prominence of the Superior-Temporal region, offering a clear anatomical explanation for the prediction. Additionally, the Grad-CAM overlay illustrates that the model mainly concentrates on retinal vascular pathways, steering clear of irrelevant background areas or imaging artifacts. This behavior boosts the credibility of the proposed framework and reinforces its potential for clinical decision-making. By merging accurate classification with visual

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

explanations, the XAI-enabled system enhances transparency and allows clinicians to check if the model's focus aligns with established retinal biomarkers of cardiovascular disease. Such interpretable predictions are essential to build trust and make it easier for healthcare settings to embrace AI-driven screening systems in everyday practice. In Figure 4, we can see the Receiver Operating Characteristic(ROC) curve for the CLAHE-enhanced EfficientNet-B0 frame-work proposed. The curve is positioned nicely near the upper-left corner, it shows that our model does an excellent job of distinguishing between healthy and diseased retinal images. With an Area Under the Curve (AUC) of 0.976, it highlights the model's impressive sensitivity and specificity across various classification thresholds.

5. Discussion

The framework outperformed earlier EfficientNet-B0 baselines on CVD prediction [13][14], with CLAHE contributing a 2.4% accuracy gain — in line with the 1.5–3.0% improvement adaptive contrast enhancement typically provides in medical image classification [19]. Its high specificity (96.1%) is especially valuable for a screening tool, since false positives carry real costs in patient anxiety and downstream care. Grad-CAM outputs added a layer of clinical sense to these numbers: low-risk cases showed weak, scattered activation, while high-risk cases showed strong, diffuse activation in the superior retinal quadrants, matching the known link between hypertensive retinopathy and the superior arcade [20]. That kind of anatomically grounded explanation is one of the things standing between deep learning models and real clinical trust. That said, the dataset's size limits how well it represents different ethnicities, camera types, and co-morbidities, so multi-center validation and comparison against established risk scores like Framingham would be needed before any real-world deployment.

Conclusion

This paper introduced an explainable deep learning framework for predicting cardiovascular disease risk from retinal fundus images, combining CLAHE preprocessing, a fine-tuned EfficientNet-B0 classifier, and Grad-CAM saliency visualization. The system reached 94.7% accuracy and a 0.97 AUC-ROC, outperforming several recent baselines. The Grad-CAM overlays and quadrant-level activation analysis offer explanations that make clinical sense, with high-risk cases consistently showing superior-temporal activation that matches known vascular biomarkers of CVD. Overall, the framework shows real promise as a non-invasive, low-cost, explainable screening tool that could fit into ophthalmology workflows and community health settings. Future directions include multi-modal fusion with intraocular pressure and OCT data, prospective clinical validation, and extending the approach to multi-class CVD severity grading.

References

[1]. P. Bizopoulos and D. Koutsouris, "Deep learning in cardiology," IEEE Rev. Biomed. Eng., vol. 12, pp. 270–293, 2019.

[2]. M. K. A. Hridoy, S. R. Barman, and R. Murugan, "Classification of cardiovascular diseases through retinal fundus images using deep learning," J. Med. Imaging Health Inf., vol. 14, no. 3, pp. 410–418, 2024.

[3]. V. Willis, B. Zhou, and Q. Liu, "Detection and prediction of cardiovas-cular disease using fundus images with deep learning," in Proc. IEEE Int. Conf. Bioinform. Biomed. (BIBM), 2024, pp. 1120–1125.

[4]. Muthuraja M. et al., "A vessel-aware deep learning framework for cardiovascular risk prediction using retinal images," IEEE Trans. Med. Imaging, vol. 44, no. 1, pp. 102–114, 2025.

[5]. S. P. Hon and M. Chavan, "Retinal imaging based deep learning for cardiovascular disease risk prediction: A systematic review," Arch. Comput. Methods Eng., vol. 32, no. 2, pp. 545–560, 2025.

[6]. U. H. Himi, S. Akter, and A. Islam, "Eyes on the heart: Detecting cardiovascular disease using retinal images with explainability," in Proc. Quality Phenotypes in AI and Nuances (QPAIN) Conf., 2025, pp. 74–83.

[7]. W. Zhang, D. Shi, and M. He, "Improving consistency in cardiovascular disease risk assessment: Cross-camera adaptation for retinal images," IEEE J. Biomed. Health Inform., vol. 28, no. 11, pp. 6210–6221, 2024.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4788-4795
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0629

[8]. C. Kotteeswaran, Ankitha, and Harish, "Cardiovascular disease predic-tion from retinal images using deep learning," Int. J. Comput. Vis., vol. 133, no. 4, pp. 889–901, 2025.

[9]. Y. P. Prakash, B. Siri, and J. Jabez, "Identifying the abnormalities in retinal images towards the prediction of cardiovascular disease using deep learning," IEEE Trans. Cybern., vol. 55, no. 2, pp. 1340 1351,2025.

[10]. S. Seniaray et al., "Cardiovascular disease prediction through retinal image analysis using deep learning," Biomed. Signal Process. Control, vol. 99, p. 106850, 2025.

[11]. A. S. Nandini, N. Srinu, and M. Senthil, "Hybrid deep learning ar- chitectures for cardiovascular disease prediction using retinal imaging and explainable AI," in Proc. IEEE Int. Conf. Innov. Comput. Intell. (ICINVENTS),2025, pp. 210–216.

[12]. Pryadarshne D., S. R. Madireddy, and M. P. Vaishnnave, "Diagnosis of cardio vascular diseases using retinal fundus scans via deep learning," IEEE Access, vol. 12, pp. 145200–145212, 2024.

[13]. P. S. Priya and L. V. CH, "Automated cardiovascular risk assessment using retinal imaging and deep learning techniques," in Proc. IEEE Int. Conf. Comput. Intell. Commun. Netw. (CICN), 2025, pp. 45–50.

[14]. B. P. A.V. and G. Srinitya, "Cardiovascular disease prediction with deep learning," IEEE Trans. Neural Netw. Learn. Syst., vol. 36, no. 3, pp. 3105–3116,2025.