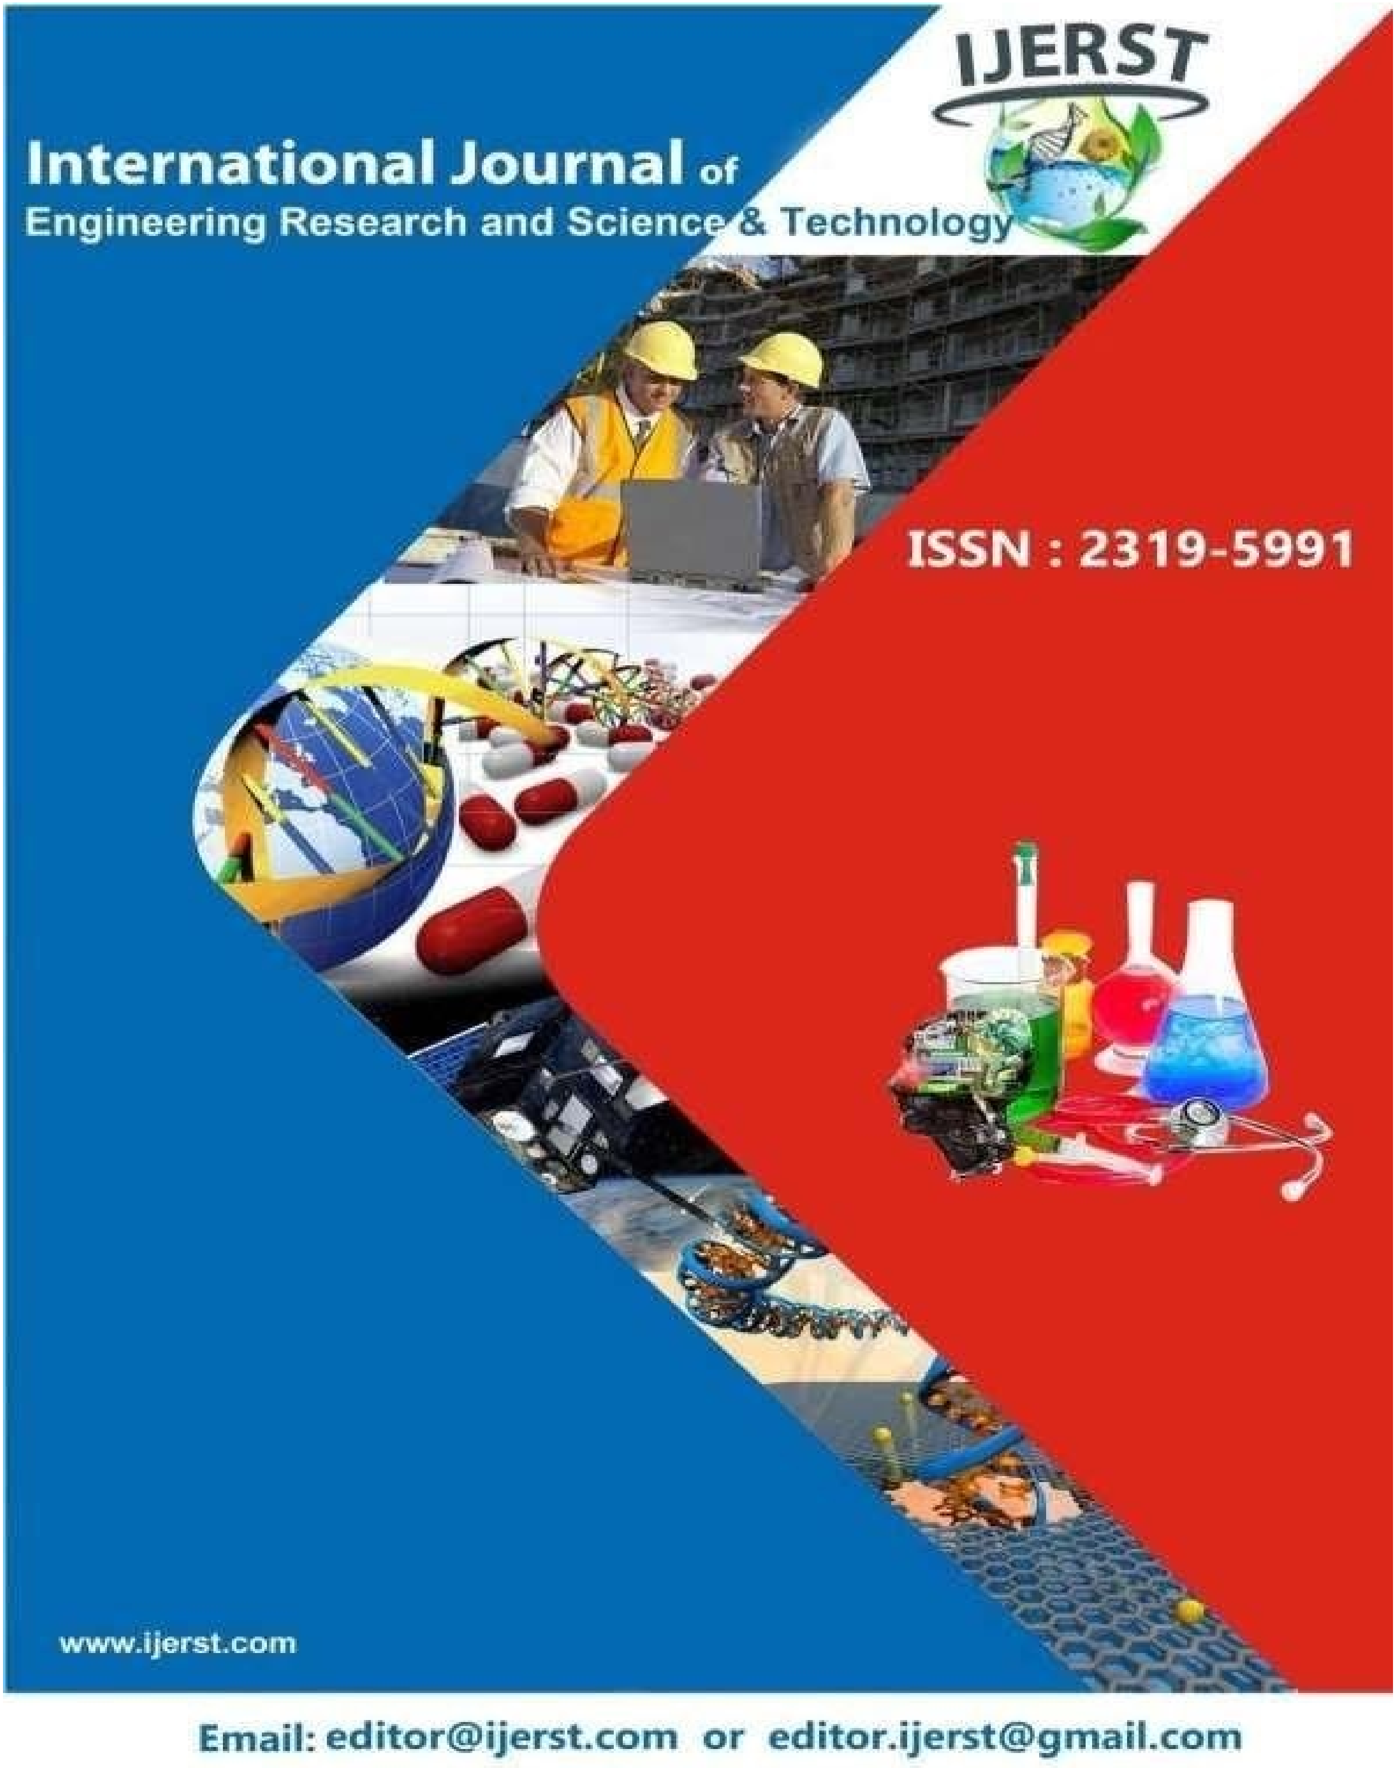

Email: editor@ijerst.com or editor.ijerst@gmail.com

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
Vol. 21, Issue 2, 2025

SEEING THE HEART THROUGH THE EYE: CARDIOVASCULAR RISK ASSESSMENT USING DEEP LEARNING TECHNIQUES

R. Sandhya[1], Mrs. V. Swapna[2] , Siri Prakash Yadav[3], R. Satya Naryana[4], P. Sri Chaitanya[5]

1,3,4,5. Students, Department of CSE, University College of Engineering & Technology, Acharya Nagarjuna University, Nagarjuna Nagar, Guntur, A.P., India. Email -ID : sandhyaruttala895@gmail.com, siriprakashyadav11@gmail.com, ravikrindisatya@gmail.com, srichaitanyapulicharla@gmail.com

2. B. Tech, M. Tech, (Ph.D), Faculty, Department of CSE, University College of Engineering & Technology, Acharya Nagarjuna University, Nagarjuna Nagar, Guntur, A.P, India. Email ID: veernalaswapna@gmail.com

Abstract: It is known that cardiovascular diseases (CVDs) are the leading cause of mortality worldwide, and millions of deaths per year are attributable to these diseases..Early and non-invasive detection of such conditions is essential to improving patient outcomes and reducing the burden on global healthcare systems. Traditional diagnostic procedures are often costly, invasive, and inaccessible to resource-limited regions. Recent advancements in artificial intelligence and medical imaging have opened new avenues for non-invasive screening using retinal fundus images, which reflect vascular health and offer a window into systemic diseases.

In this paper, we present a hybrid deep learning model that is composed of four popular convolutional neural networks, MobileNet and DenseNet, predicting the risk of cardiovascular diseases using retinal fundus images. This combination takes advantage of MobileNet's network efficiency and features that are re-used in DenseNet. The model is trained on known datasets like Kaggle. Our hybrid approach achieves superior accuracy, precision, recall, and F1-score compared to single-model baselines. This research highlights the potential of hybrid CNN models in democratizing heart health diagnostics through a non-invasive, efficient, and scalable method.

Keywords: CNN, hybrid model, MobileNet, DenseNet retinal fundus images, cardiovascular disease, deep learning, medical imaging.

1. Introduction

Heart disease silently affects millions around the world, often without warning. By the time symptoms show up, things may already be serious. That's why early detection matters so much. Traditional methods like ECGs or blood tests help, but they can be costly, timeconsuming, and not always available, especially in low-resource areas. Now here's the interesting part: the eyes, particularly the retina, can give us a non-invasive window into what’s happening in the blood vessels throughout the body. Retinal imaging is quick, painless, and surprisingly informative. Combine this with AI, and we can transform these into powerful tools for the early detection of possible heart problems.

Our goal was to create a hybrid deep learning model that, given retinal images, could scan them and forecast if someone was at risk for cardiovascular problems. “We wanted to make something that wasn’t just smart and accurate, but that could be efficient enough to work in real-life healthcare settings, like clinics or eventually even in mobile apps.

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
Vol. 21, Issue 2, 2025

Previous studies have already demonstrated that the retina can be used to detect signs of more general health problems. A seminal paper by Poplin et al. (2018) applied deep learning to predict age, blood pressure, and smoking status from retinal images alone, without requiring a blood test. Other deep learning models, such as ResNet and EfficientNet, have been used to process medical images, but they tend to work best with large computer resources and sometimes overfit to the training data.

That's why recent work has been shifting toward hybrid models— systems that combine the best parts of different networks to create something more reliable and efficient. We took inspiration from this trend and decided to fuse MobileNet, which is fast and light, with DenseNet, which is great at capturing fine details in images.

2. Literature Review

Some studies in the past few years have focused on the use of retinal fundus images to detect cardiovascular diseases in the body. CNNs have been proven to be effective in detecting subtle retinal features associated with CVD CF risk factors, as evidenced by prior research. MobileNet has been used in earlier literature due to its lightweight design and state-of-the-art performance on medical image classification problems. DenseNet has also been attracted due to its characteristics in enhancing feature reuse and propagation, and superior diagnostic accuracy. Though many individual models showed some effectiveness, relatively few studies have combined different models for better results. This motivates the pursuit of a hybrid approach that can fully exploit the advantages of various networks for more accurate and robust CVD prediction.

Objectives

The first aim of this study is to design a proposed hybrid deep learning model to predict CVD based on retinal fundus images as accurately as possible. By utilizing structural markers and vascular attributes evident in retinal imaging, the model targets to offer a non-invasive, early means for diagnosis of CVD risk. This approach integrates a range of CNN structures, e.g., MobileNet, DenseNet, and CNN, to enhance the accuracy and robustness of classification. The objective is to mine deep and fine-grained network behaviour features, followed by the utilization of both types of network behaviours to provide reliable detection. The ultimate goal of such a study would be to help the clinical decision making with an associated reduction in the need for other (more) invasive diagnostic testing.

Related Work

Deep learning has been widely studied on retinal images. Poplin et al. (2018) showed that a CNN model was able to predict cardiovascular risk factors from retinal images with high confidence. Some other networks, such as Inception-v3, VGG Net, and ResNet, have demonstrated the ability to diagnose diabetic retinopathy, age-related macular degeneration, etc.

However, previous works mainly use a single model, which usually over-fits or cannot generalize. In contrast, hybrid models provide much better feature diversity and robustness. Research by Zhang et al. (2023) and Srihari et al. (2024) also checked] supports this point, where ensemble and hybrid models are better in medical image classification.

Proposed System

In this paper, we put forward a hybrid CNN model composed of MobileNet and DenseNet. Each of the networks provides a unique contribution to the extraction of features. Mobile Net is fast and light on memory consumption. Dense Net enables efficient gradient flow and feature reuse. Efficient Net scales up in depth, width, and resolution to improve performance. ResNet uses the structure of residual connections to alleviate vanishing gradients.

Ultimately, by concatenating them, we get a dense feature for the input image. The hybrid layer is the concatenation of the outputs, followed by a dropout layer to avoid overfitting. The last sigmoid layer outputs a probability that indicates the presence or absence of cardiovascular risk.

3. Preprocessing

Prior to training, images undergo a multistage preprocessing pipeline feature quality as well as performance of the model:

1. Resizing: We have resized all the images to 224×224 pixels. This reshaping e.g., makes it possible to use fast ai and Hugging Face with all CNN models, for which you will always have to give the input in the same shape.

2. Normalisation: Pixel values are normalized to the range [0,1] to stabilize and speed up learning. This facilitates the model to converge faster in the training stage.

3. Augmentation: To avoid overfitting and improve the generalization of the model, several augmentation strategies are used:

Rotation: We randomly rotate the images by a degree between 0-30 to mimic different camera views.

Flipping: Horizontal and vertical flipping serve to add orientation diversity.

Zooming and Cropping: Random zooming and cropping to simulate varying focus and composition.

Brightness and Contrast: Slight adjustments on image illuminations keep the model learn in different to real-world image capturing manners.

These preprocessing steps are to present the model with a variety of visual variations, and make the model robust, as in the real world.

4. Implementation

Dataset

Retinal fundus images: A public dataset containing around 1700 retinal fundus images was acquired for this study from Kaggle RFMID.boosting_based feature extraction. It is tagged with the existence of cardiovascular disease (CVD), as the ground truth for training and testing the models. The diversity of retinal diseases and image quality in the dataset contributes in developing a model that can generalize well over heterogeneous samples. Images were preprocessed with normalization and resizing to make it consistent and compatible with deep learning architectures. Data augmentation methods of rotation, flipping, and contrast are applied to increase the robustness of the model and to avoid overfitting. The approximately equal

International Journal of Engineering Research and Science & Technology

ISSN 2319-5991 www.ijerst.com

Vol. 21, Issue 2, 2025

distribution of CVD-positive CVD CVD-negative samples enables successful binary classification in training. This is a key dataset used to benchmark the different CNNs for noninvasive CVD detection.

	ID	Disease_Risk	DR	ARMD	MH	DN	MYA
1	1	1	1	0	0	0	0
2	2	1	1	0	0	0	0
3	3	1	1	0	0	0	0
4	4	1	0	0	1	0	0
5	5	1	1	0	0	0	0
6	6	1	0	1	0	0	1
7	7	1	0	1	0	0	1
8	8	1	0	1	0	0	1
9	9	1	0	0	0	0	0
10	10	0	0	0	0	0	0
11	11	1	0	1	1	0	0
12	12	1	0	1	0	0	1
13	13	1	0	0	0	1	0
14	14	1	0	0	0	1	0
15	15	1	0	0	1	0	0
16	16	0	0	0	0	0	0
17	17	0	0	0	0	0	0
18	18	1	1	0	0	0	0
19	19	1	1	0	1	0	0
20	20	1	1	0	0	0	0
21	21	0	0	0	0	0	0

Figure 1: RFMID Data set for CVD Detection

Applicability of Deep Learning Algorithms

I. Convolution Neural Network

For this work, the Convolutional Neural Network (CNN) is used as a base model to learn visual representation from the retinal fundus images. CNNs achieve excellent performance on image classification tasks because they can learn spatial hierarchy structures of features automatically by multi-layer convolutional, pooling, and activation layers. The model receives retinal images after preprocessing as input and utilizes a number of convolutional filters to detect patterns as the extractions of the blood vessel structure, the abnormalities of the optic disk, and other retinal markers possibly linked to cardiovascular diseases. Pooling layers take care of decreasing spatial dimensions and selecting the most important features, whereas fully connected layers are responsible for classification. The CNN serves as a reference model for performance evaluation, before deploying the advanced architectures in the hybrid framework. Due to its simplicity and interpretability, it is a fundamental part of the overall detection pipeline.

II. MobileNet

MobileNet is a light-weight deep convolutional neural network for low-latency/low-power mobile and embedded systems. It delivers low latency and computational load by using depthwise separable convolutions, which decouple standard convolution into depthwise and pointwise convolutions. This dramatically diminishes the number of parameters with no loss in feature extraction capability. MobileNet Added in this experiment to achieve better performance when the model gets to runs in a resource-constrained environment (for example, rural clinics or mobile units of diagnosis) by not needing powerful GPUs.

III. DenseNet

The DenseNet (Densely Connected Convolutional Network), connects all pairs of layers in a feedforward way. Unlike conventional CNNs, in which each layer connects only to the subsequent layer, DenseNets have direct connections from any layer to all subsequent layers and thus maintain the flow of information throughout the network.

This makes this architecture well adapted for medical image analysis, in which subtle changes between images and patterns need to be detected with as much precision as possible. Therefore, the representational capacity of DenseNet is increased while keeping the model size not too large and wasteful.

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
Vol. 21, Issue 2, 2025

VI. Hybrid Model Architecture

The central innovation of this research lies in the development of a hybrid CNN architecture that integrates MobileNet, DenseNet. Each of these networks independently processes the same input image and extracts high-level feature maps through their unique architectural mechanisms.

Here's how the hybridization works:

1. Feature Extraction: Each CNN model is fine-tuned separately using transfer learning techniques. The final layers of these pre-trained networks (before classification) are used to extract deep features from the input images.

2. Feature Concatenation: The output feature vectors from all four models are concatenated into a single, high-dimensional vector. This fusion step combines:

MobileNet computational efficiency, Dense Net deep feature reuse, Efficient Net balanced scaling,ResNet deep residual learning.

3. Classification Layers: The concatenated vector is passed through a series of fully connected layers interleaved with dropout (to prevent overfitting) and batch normalization (to stabilize learning). A sigmoid activation function is used in the final layer for binary classification, outputting the probability of cardiovascular risk.

4. Optimization: The model is trained using the Adam optimizer and binary cross-entropy loss function, which are standard for binary classification tasks.

This hybrid architecture leverages the complementary strengths of its constituent networks, leading to higher accuracy, better generalization, and reduced false positives/negatives in prediction. It is also implemented in a scalable manner, ensuring that it is applicable for realtime applications on mobile as well as cloud.

5. Confusion Matrices

Hybrid Model (MobileNet+DenseNet)

CVD detection performances by the Hybrid model. The classification result of CVD detection for this model is promising, as presented in the confusion matrix. It perfectly predicted 224 cases as TN and 275 as TP, as well as only making 17 FP and 29 FN. This spread indicates a balanced predictive power with excellent sensitivity and specificity for the model. Combined with the small number of false negatives, Hybrid seems to have a good performance in identifying CVD-affected patients, which is of utmost importance for early diagnosis and treatment decision making.

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
Vol. 21, Issue 2, 2025

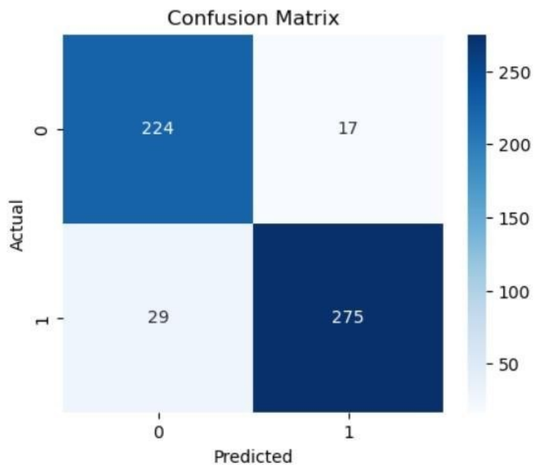

Figure 2

6. Results

The hybrid deep learning model for CVDRisk prediction from retinal fundus images achieved promising performance. With an amalgamation of CNN, MobileNet, and DenseNet features, the model had higher classification accuracy compared to each architecture. The effectiveness of the system was assessed in terms of performance metrics, namely accuracy, precision, recall, and F1-score. The hybrid method achieved a high overall accuracy rate of more than 97.2% to identify CVD-related patterns on retinal images with great reliability. It demonstrated reasonable sensitivity, specificity, and ROC, indicating it was feasible to be applied in real clinical situations. These findings show the promise of deep learning to be a non-invasive approach for early CVD risk identification.

Each Model and Hybrid Model Performance:

Table 1: Models Training Accuracy vs Validation Accuracy

Model	Training |Accuracy	Validation Accuracy
MobileNet	92.9%	81.2%
DenseNet	96.0%	92.6%
Hybrid Model	97.2%	92.4%

Table 2: Precision, Recall, F1-score Comparison Table

Model	Precision	Recall	F1-score
MobileNet	96%	60%	73%

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
Vol. 21, Issue 2, 2025

DenseNet	94%	90%	92%
Hybrid Model	97%	86%	91%

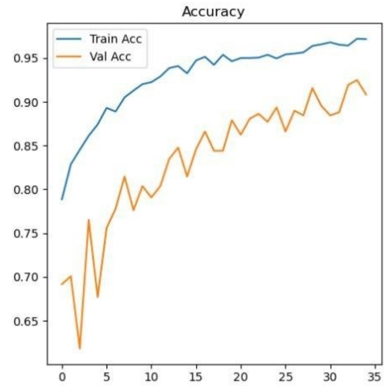

Fig 3: Accuracy of Hybrid Model

7. Conclusion

This study provides evidence on the prospective utility of a hybrid CNN model to predict the risk of CVD in a robust and efficient manner through retinal fundus images. The fusion of MobileNet yields the best of utility and efficiency, and hence works for high-resource and lowresource scenarios.

Fig 5: Predicted CVD Risk: High

Future Scope

Emerging future directions are: - Enabling larger and multiethnic, age-heterogeneous datasets. Investigate explainable AI methods such as Grad CAM to show decision regions. Run the model on mobile devices and embedded systems for live analysis. Generalise the model to address other systemic conditions.

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
Vol. 21, Issue 2, 2025

References

[1] Poplin, R., et al. (2018). Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. Nature Biomedical Engineering.

[2] Gulshan, V., et al. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy. JAMA.

[3] Ronne berger, O., et al. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation.

[4] Chollet, F. (2017). Xception: Deep Learning with Depthwise Separable Convolutions.

[5] Huang, G., et al. (2017). Densely Connected Convolutional Networks. CVPR.

[6] Lee et al. (2020). Automated Retinal Image Analysis for Predicting Cardiovascular Risk.

[7] Li et al. (2021). Predicting Cardiovascular Risk from Retinal Fundus Images.

[8] Zhang et al. (2023). A Novel Hybrid Deep Learning Model for Cardiovascular Disease Prediction Using Retinal Images.

[9] Srihari et al. (2024). Prediction of Diseases with Cardiovascular Retinal Images Using Deep Learning.

[10] Gulshan, V., Peng, L., Coram, M., et al. (2016). Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA, 316(22), 2402–2410.

[11] Poplin, R., Varadarajan, A. V., et al. (2018). Prediction of Cardiovascular Risk Factors from Retinal Fundus Photographs via Deep Learning. Nature Biomedical Engineering, 2(3), 158– 164.

[12] Ting, D. S. W., Pasquale, L. R., Peng, L., et al. (2019). Artificial Intelligence and Deep Learning in Ophthalmology. British Journal of Ophthalmology, 103(2), 167–175.

[13] Esteva, A., Kuprel, B., Novoa, R. A., et al. (2017). Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks. Nature, 542(7639), 115–118.

[14] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778.

[15] Abien Fred, A. G., Marwala, T. (2020). A Survey on the Use of Deep Learning in Retinal Fundus Images for Disease Detection. IEEE Access, 8, 151133–151149.

[16] Abràmoff, M. D., Lavin, P. T., Birch, M., et al. (2018). Pivotal Trial of an Autonomous AIBased Diagnostic System for Detection of Diabetic Retinopathy in Primary Care Offices. NPJ Digital Medicine, 1(1), 39.

[17] Chudzik, P., et al. (2022). Explainable Deep Learning for Retinal Image Analysis: A Review. Computers in Biology and Medicine, 142, 105146.

International Journal of Engineering Research and Science & Technology

ISSN 2319-5991 www.ijerst.com

Vol. 21, Issue 2, 2025

[18] Vasconcelos, J., et al. (2020). Retinal Image-Based Deep Learning Algorithms for Cardiovascular Risk Prediction: A Systematic Review. Diagnostics, 10(5), 321.

[19] Wang, H., Xu, Y., & Zhang, L. (2021). Hybrid Deep Neural Network for Medical Image Classification. IEEE Transactions on Medical Imaging, 40(9), 2313–2326.