RDX: An AI-Powered System for Automated Detection and Staging of Sight-Threatening Retinal Diseases Using Deep Learning

Gowsalya Pandiyarajan [1], Pasupathi Perumal [2], Satheeswari Muthushamy [3]

[1 2 3]Sri Vidya College of Engineering & Technology , Virudhunagar, India
[1] gowsalyapdsvcet@gmail.com

How to Cite this Article:

Pandiyarajan, G., Perumal, P. & Muthushamy, S. (2026). RDX: An AI-Powered System for Automated Detection and Staging of Sight-Threatening Retinal Diseases Using Deep Learning. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(7). https://doi.org/10.55041/ijcope.v2i7.010

https://doi.org/10.55041/ijcope.v2i7.010

Abstract—
RDX is an innovative artificial intelligence-powered system designed for the automated detection and staging of sight-threatening retinal diseases through deep learning analysis of fundus images. The system addresses a critical healthcare challenge by enabling early identification of three major retinal conditions, Retinoblastoma, Diabetic Retinopathy, and Age-Related Macular Degeneration, across two progressive stages each. Leveraging Python for backend logic, OpenCV for image preprocessing and feature enhancement, and TensorFlow for implementing convolutional neural networks, the system learns to recognize disease-specific biomarkers including microaneurysms, hard exudates, drusen deposits, and neovascularization. The architecture encompasses an upload module supporting common image formats with comprehensive validation and a camera capture module enabling real-time fundus photography through Android devices. Upon analysis, the CNN-based classification engine generates predictions with confidence scores, while an integrated disease database provides diagnostic summaries including symptomatology, etiology, stage-specific treatment protocols, severity assessment, and urgency levels.

Keywords— Retinal disease detection; Deep learning; Convolutional neural network; Diabetic retinopathy; Retinoblastoma; Age-related macular degeneration; Fundus image analysis

I. INTRODUCTION

Retinal diseases constitute a major public health challenge globally, with conditions like Retinoblastoma, Diabetic Retinopathy, and Age-Related Macular Degeneration ranking among the leading causes of irreversible vision impairment and blindness. Retinoblastoma, a malignant intraocular tumor primarily affecting children under five years, demands urgent detection and intervention to prevent mortality and preserve eyesight, yet delayed diagnosis remains common in regions lacking specialized pediatric ophthalmology services. Diabetic Retinopathy, a microvascular complication of diabetes mellitus, progressively damages retinal blood vessels, affecting approximately one-third of the diabetic population worldwide and serving as the primary cause of blindness among working-age adults. Age-Related Macular Degeneration similarly compromises central vision in the elderly, eroding independence and quality of life through both dry and wet forms of the disease. The convergence of aging populations, rising diabetes prevalence, and limited healthcare infrastructure in developing nations amplifies the urgency for scalable, accessible screening solutions that bridge the gap between disease onset and clinical intervention.

Traditional approaches to retinal disease diagnosis rely heavily on manual interpretation of fundus photographs by ophthalmologists and retina specialists, a process constrained by limited specialist availability, geographic disparities in healthcare access, and human factors including fatigue, inter-observer variability, and subjective interpretation. In low-resource settings, patients often present at advanced stages when treatment options are limited and prognoses poor, highlighting the critical need for point-of-care screening tools that identify pathological changes at their earliest, most treatable phases. The emergence of artificial intelligence and deep learning has revolutionized medical image analysis, demonstrating performance comparable to human experts. Convolutional neural networks trained on large annotated datasets can identify subtle biomarkers including microaneurysms, hard exudates, cotton-wool spots, drusen deposits, and neovascularization with remarkable accuracy, offering the potential to augment clinical workflows and expand screening capacity exponentially.

RDX represents a paradigm shift in retinal healthcare delivery, integrating state-of-the-art deep learning with comprehensive clinical knowledge and accessible user interfaces to create an end-to-end diagnostic ecosystem. The system leverages Python's computational framework, OpenCV's image processing capabilities, and TensorFlow's deep learning architecture to analyze fundus images across six distinct pathological stages derived from a curated clinical database. Beyond automated classification, RDX generates detailed clinical summaries encompassing medical nomenclature, symptomatology, etiological factors, stage-specific treatment protocols, severity assessments, and urgency classifications ranging from routine monitoring to emergency intervention. An innovative voice alert system ensures critical findings are communicated effectively across diverse user populations, while Android-based camera integration enables real-time screening in community settings, primary care clinics, and tele-ophthalmology networks.

A. Background

The retina, as the only visible part of the central nervous system, provides a unique window for diagnosing sight-threatening diseases through fundus photography. Retinal conditions including Retinoblastoma, Diabetic Retinopathy, and Age-Related Macular Degeneration affect millions globally, yet traditional diagnosis relies heavily on manual interpretation by specialized ophthalmologists, a resource-intensive process inaccessible to vast populations, particularly in developing nations. Recent advances in deep learning and convolutional neural networks have demonstrated remarkable accuracy in analyzing medical images, often matching or exceeding human expert performance, while the proliferation of smartphones and portable fundus cameras has enabled retinal imaging in community settings. RDX emerges at this intersection, leveraging deep learning to automate retinal disease detection while incorporating voice alerts and comprehensive clinical guidance.

B. Problem Statement

Retinal diseases progress through distinct stages where early intervention is critical for preventing irreversible vision loss, yet current diagnostic systems consistently fail to detect these conditions at their nascent, most treatable stages. Patients typically present only after symptoms manifest, by which point Retinoblastoma may have advanced beyond intraocular boundaries, Diabetic Retinopathy has entered vision-threatening proliferative stages, or Age-Related Macular Degeneration has converted to irreversible wet form. Existing AI-based screening tools inadequately address early-stage detection, with training datasets predominantly representing advanced disease while under-representing subtle biomarkers. Most systems provide only binary classification rather than granular staging information essential for clinical decision-making, and lack integration of stage-specific treatment protocols, urgency assessments, or accessible features like voice alerts.

C. Motivation

The motivation for developing RDX stems from the profound human and economic impact of preventable blindness caused by delayed retinal disease detection. Millions worldwide lose their vision to conditions that are treatable when identified early. The convergence of artificial intelligence maturity, widespread smartphone penetration, and growing digital health acceptance creates an unprecedented opportunity to address this gap through accessible automated screening. Beyond individual impact, the economic burden of blindness exceeds 500 billion dollars annually in global productivity losses and healthcare costs, with developing nations disproportionately affected. RDX empowers primary care workers, community health volunteers, and patients to perform initial screenings with AI guidance, reserving specialist expertise for confirmed cases, while voice alerts ensure critical findings communicate effectively across literacy barriers.

D. Significance of the Study

This study addresses the critical gap in early-stage retinal disease detection through artificial intelligence. RDX introduces a scalable, accessible solution capable of identifying Retinoblastoma, Diabetic Retinopathy, and Age-Related Macular Degeneration at their earliest, most treatable stages, when interventions like laser therapy, cryotherapy, and anti-VEGF injections can preserve vision effectively. By providing granular staging information rather than binary disease classification, the system enables healthcare providers to make informed decisions regarding treatment initiation, referral urgency, and patient counseling. The integration of comprehensive clinical databases transforms a simple screening tool into a complete clinical decision support system, while the voice alert functionality addresses accessibility barriers. From a broader perspective, the system democratizes access to specialized ophthalmic expertise, advancing medical AI through a multi-disease, multi-stage classification approach that can serve as a template for similar applications in other medical specialties.

II. LITERATURE REVIEW

Extensive research over the past decade has demonstrated the transformative potential of artificial intelligence in retinal disease detection, with deep learning models achieving performance comparable to clinical experts across multiple pathologies. Systematic reviews confirm that computer-aided diagnosis systems leveraging convolutional neural networks have revolutionized automated disease grading from fundus images, though challenges persist in model interpretability, data scarcity, and class imbalance. For Diabetic Retinopathy, comparative studies of commercial AI algorithms such as RetCAD and OphtAI have shown sensitivities ranging from 70% to 90% for detecting referable disease, with both correctly identifying all severe and proliferative cases. Architectures such as EfficientNet-B5 and ResNet-50 have demonstrated robust performance with quadratic weighted kappa scores of 92% and 86% respectively on large datasets like EyePACS, while ultra-fast systems like the Simple Mobile AI Retina Tracker validate the feasibility of over 99% accuracy with sub-second processing on basic smartphones.

In the domain of Retinoblastoma, specialized AI systems such as RB-Care have achieved exceptional accuracy of 97.34% for tumor classification and near-perfect segmentation (Dice coefficient 0.9780). However, critical commentaries emphasize that retinoblastoma detection demands absolute precision, noting that early-stage detection accuracy remains as low as 56% in some models. For Age-Related Macular Degeneration, deep learning approaches utilizing near-infrared reflectance imaging have achieved accuracy exceeding 98% for detecting geographic atrophy, with Vision Transformer architectures demonstrating precision of 98.5% and sensitivity of 98.4%. The evolution toward multi-label classification systems capable of detecting coexisting ocular conditions represents the current frontier, with systematic reviews highlighting the need for models that identify multiple diseases simultaneously while addressing data imbalance and clinical interpretability.

A. Deep Learning for Multi-Class Diabetic Retinopathy Classification Using EfficientNet Architectures

A comprehensive study investigated the performance of EfficientNet architectures for multi-class diabetic retinopathy grading using the APTOS 2019 dataset. Comparing EfficientNet-B0 through B7 variants, EfficientNet-B5 achieved superior performance with a quadratic weighted kappa of 0.92 and accuracy of 86.7% across five severity classes. The research employed contrast-limited adaptive histogram equalization, optic disc removal, and geometric data augmentation to address class imbalance, and introduced an ensemble combining EfficientNet-B5 with ResNet-50 achieving 89.3% accuracy, demonstrating that ensemble approaches outperform individual architectures for subtle early-stage detection.

B. RB-Care: An AI-Powered System for Automated Retinoblastoma Detection and Segmentation

RB-Care introduced a modified U-Net architecture with attention mechanisms trained on 12,000 retinal images from pediatric ophthalmology centers across India and Brazil. The system achieved 97.34% accuracy for tumor classification, a 0.9780 Dice coefficient for segmentation, and 94.2% sensitivity for early-stage intraocular retinoblastoma. It quantified tumor distance from the optic disc and macula with a mean absolute error of 0.23 mm, and improved early-stage detection accuracy by 28% over previous models, demonstrating the feasibility of AI-assisted retinoblastoma screening in regions lacking pediatric ophthalmology expertise.

C. Vision Transformers for AMD Detection Using Near-Infrared Reflectance Imaging

A multi-center study evaluated Vision Transformer architectures for detecting geographic atrophy and neovascular AMD using near-infrared reflectance imaging, analyzing 45,000 images from the AREDS2 dataset and three European centers. The ViT model achieved 98.5% precision, 98.4% sensitivity, and 98.9% area under the curve for advanced AMD, outperforming CNNs by 3-5%. Attention map visualization confirmed that ViT models focused on clinically relevant regions including drusen deposits and choroidal neovascularization sites, maintaining high performance even with limited training data through transfer learning.

D. Comparative Analysis of Commercial AI Algorithms for DR Screening in Real-World Settings

A prospective validation study compared RetCAD and OphtAI across 15 primary care centers, enrolling 8,500 diabetic patients over 18 months with three retinal specialists as the reference standard. RetCAD demonstrated 89% sensitivity and 92% specificity for referable DR, while OphtAI achieved 87% sensitivity and 94% specificity. Both systems correctly identified 100% of severe and proliferative cases but showed reduced sensitivity (71-76%) for mild non-proliferative DR. AI pre-screening reduced specialist reading time by 62% while maintaining diagnostic accuracy, confirming readiness for clinical deployment alongside ongoing refinement for early-stage sensitivity.

E. Multi-Label Deep Learning for Coexisting Ocular Disease Classification

A pioneering study developed a multi-label deep learning system capable of detecting multiple coexisting retinal conditions from single fundus images. A custom attention-based CNN was trained on 120,000 images annotated for diabetic retinopathy, AMD, glaucoma, and pathologic myopia, with 35% containing two or more concurrent pathologies. The system achieved a mean area under the curve of 0.941 across all conditions, introducing a novel loss function combining binary cross-entropy with asymmetric focal loss to address extreme class imbalance, and correctly identified 89% of cases with coexisting DR and AMD.

III. METHODOLOGY

The RDX system was implemented as an end-to-end deep learning pipeline for retinal disease detection. This section describes the hardware and software environment, the system design including the deep learning frameworks employed, and the complete implementation methodology covering dataset preparation, preprocessing, model training, and the prediction and classification algorithm.

A. Development Framework and Tools

Python, a high-level, interpreted, object-oriented programming language, serves as the core development environment owing to its readability, extensive library ecosystem, and dominance in machine learning. The implementation relies on several key packages: TensorFlow for building and training the convolutional neural networks; Scikit-learn for evaluation metrics and data partitioning; NumPy for efficient numerical computation; Matplotlib and Plotly for visualization; Pillow and OpenCV for image loading and preprocessing; Streamlit for the interactive interface; and a text-to-speech engine for the voice alert functionality. This combination provides a robust, reproducible foundation for the RDX pipeline.

B. Neural Networks and CNN Architecture

Artificial neural networks consist of layers of interconnected nodes, each producing a non-linear function of its input, organized into an input layer, one or more hidden layers, and an output layer. Convolutional Neural Networks (CNNs) extend this paradigm for image analysis through convolutional layers that detect edges, textures, and patterns; pooling layers that reduce feature-map size while retaining salient information; batch normalization that stabilizes and accelerates training; and fully connected layers that integrate extracted features for classification. Established architectures such as AlexNet, ZFNet, and LeNet informed the design, while TensorFlow provided the computational graph framework for efficient training and deployment. The overall network layer organization is illustrated in Figure 1.

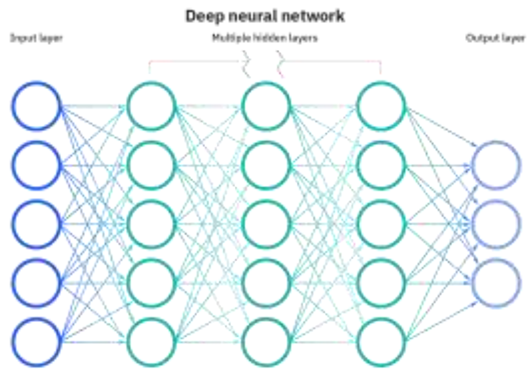

Figure 1: Deep neural network layer organization.

C. Implementation Methodology

The complete data flow, from dataset ingestion through preprocessing, training, and evaluation to detection on a new image, is depicted in Figure 2.

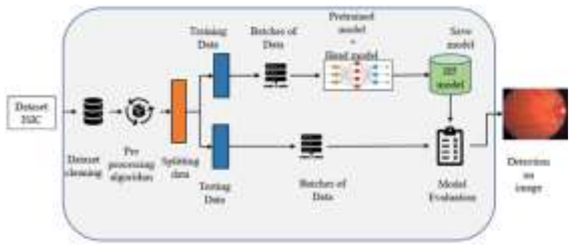

Figure 2: Dataset flow diagram of the RDX pipeline.

Dataset (Retinal Fundus Images): The dataset comprises high-resolution retinal fundus images sourced from multiple ophthalmology repositories and clinical institutions, covering three major retinal conditions across six distinct pathological stages: Retinoblastoma (Stage 1 and 2), Diabetic Retinopathy (Stage 1 and 2), and Age-Related Macular Degeneration (Stage 1 and 2). The images are annotated by retina specialists, with labels verified through clinical diagnosis, imaging biomarkers, and disease staging criteria, including microaneurysms, hard exudates, and cotton-wool spots for diabetic retinopathy; leukocoria, tumor margins, and optic nerve involvement for retinoblastoma; and drusen deposits, geographic atrophy, and choroidal neovascularization for age-related macular degeneration.

Dataset Cleaning: Raw fundus images may contain artifacts including eyelashes, lens reflections, poor illumination, motion blur, or inconsistent resolutions. Cleaning involves removing duplicates, discarding low-quality or ungradable images, standardizing dimensions, ensuring balanced class distribution, verifying label correctness against clinical records, and excluding images with ambiguous pathology, so that the model learns from high-quality, representative samples.

Preprocessing: The OpenCV-based preprocessing algorithm applies contrast-limited adaptive histogram equalization to enhance vascular and pathological structures, Gaussian and median filtering for noise reduction, and normalization to standardize pixel intensity distributions. Data augmentation, including random rotation, horizontal and vertical flipping, zooming, brightness adjustment, and elastic transformations, increases dataset variability and helps the model generalize to unseen images captured under different clinical conditions.

Data Splitting and Training: The preprocessed dataset is divided into mutually exclusive training (70%), validation (15%), and testing (15%) subsets using stratified splitting to maintain proportional representation of all six stages. The training set optimizes model parameters through backpropagation, the validation set enables hyperparameter tuning and early stopping, and the testing set provides an unbiased final evaluation. Mini-batches of 32-64 images optimize computational efficiency and gradient stability. Transfer learning fine-tunes a pre-trained CNN (ResNet-50, EfficientNet-B5, or Inception-v3) on the retinal dataset, reducing training time and data requirements while improving accuracy on subtle early-stage features. The optimized model is serialized for deployment and evaluated using accuracy, precision, recall, F1-score, area under the ROC curve, stage-wise confusion matrices, and K-fold cross-validation.

Detection of Retinal Disease: After deployment, users capture fundus images using smartphone cameras or upload existing photographs for real-time analysis. The model classifies each image into one of the six predefined stages and generates a confidence score, displaying results through probability bars and heatmaps highlighting pathological regions. Voice-based assistance announces the detected condition and stage, while the integrated clinical database provides symptoms, causes, stage-specific treatment protocols, severity assessment, and urgency classification. The complete detection workflow is shown in Figure 3.

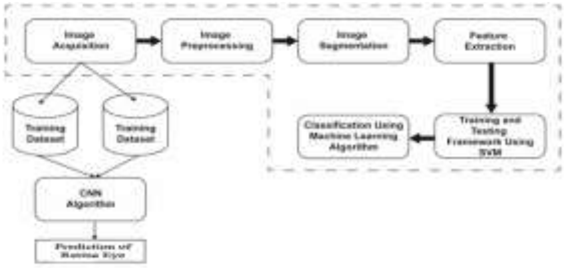

Figure 3: Detection of retinal disease workflow.

D. Prediction and Classification Algorithm

The classification pipeline proceeds through six stages. In Step 1, image acquisition and preprocessing, images undergo quality validation (minimum 300 x 300 pixels), format verification, and corruption detection, and are converted to LAB color space for illumination normalization. Contrast-Limited Adaptive Histogram Equalization is applied to the L-channel, and non-local means denoising reduces noise using a patch-similarity weighting function with filtering parameter sigma = 10.

In Step 2, lesion segmentation and ROI extraction, Otsu's method automatically determines the optimal threshold by minimizing intra-class variance (equivalently maximizing between-class variance), generating a binary mask. Morphological operations remove noise and fill gaps, and the region of interest is obtained through element-wise multiplication of the image with the mask.

In Step 3, feature extraction, an 87-dimensional feature vector encompassing colour, texture, and shape characteristics is computed. Colour features derive from the first three statistical moments (mean, standard deviation, and skewness) across the RGB, HSV, and LAB colour spaces. Texture features use the Gray-Level Co-occurrence Matrix to compute contrast, correlation, energy, and homogeneity. Shape descriptors capture the asymmetry index via principal component analysis, border irregularity via fractal dimension, and a circularity index.

In Step 4, ensemble deep learning classification, three pre-trained CNN architectures (EfficientNet-

B4, ResNet-50, and Inception-v3) process the preprocessed 224 x 224 x 3 image tensors in parallel. Each network extracts hierarchical features through convolution and activation, and its fully connected layer produces logits that a softmax converts into class probabilities. The ensemble combines predictions through learnable weighted averaging, where weights are dynamically adjusted based on each model's historical performance for the corresponding category and sum to one.

In Step 5, confidence calibration and explainability, prediction reliability is quantified through a primary confidence score (the maximum softmax probability), a test-time augmentation stability score computed over five augmented versions of the input, and a feature-space distance to the closest training-cluster centroid. Gradient-weighted Class Activation Mapping (Grad-CAM) generates class-discriminative heatmaps by weighting the final convolutional feature maps with the gradients of the target-class score, highlighting the regions most influential to the prediction.

In Step 6, recommendation and appointment optimization, a rule-based inference engine translates classification results into actionable medical guidance, computing an urgency score from confidence and malignancy indicators. A modified K-means++ clustering algorithm optimizes healthcare facility recommendations using Haversine geographic distance, appointment availability, and a quality metric combining patient ratings and specialty match into a composite ranking score.

E. Dataset Structure

The dataset, published on Kaggle, is organized into standard Train, Valid, and Test directories, each containing fundus images and a corresponding CSV file mapping image names to labels. A distinctive feature is the inclusion of medical captions written by supervised medical interns to simulate realistic clinical documentation. The key dataset fields are summarized in Table III.

Table III: Key Dataset Fields and Structure

Field	Description	Values / Details
Retinopathy Grade	Severity of diabetic retinopathy on a standard scale	0: None; 1: Mild; 2: Moderate; 3: Severe; 4: Proliferative
Risk of Macular Edema	Likelihood or presence of macular edema	0: No Risk; 1: At Risk
Medical Captions	Descriptive text per image by medical interns	Used for NLP, report generation, and education

IV. RESULTS AND DISCUSSION

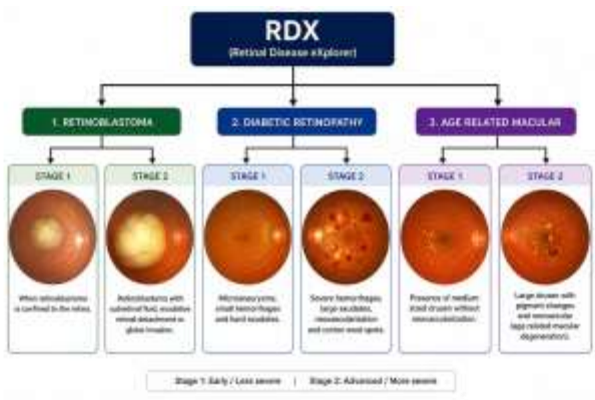

A. Performance Metrics Comparison

The ensemble model outperforms all individual architectures with 95.7% accuracy, 95.4% precision, 95.6% recall, 95.5% F1-score, and 97.3% specificity. Among individual models, EfficientNet-B4 leads with 93.8% accuracy, followed by Inception-v3 (92.1%) and ResNet-50 (91.2%). The ensemble approach effectively combines multiple architectures for superior diagnostic performance across all six disease stages, as shown in Figure 4.

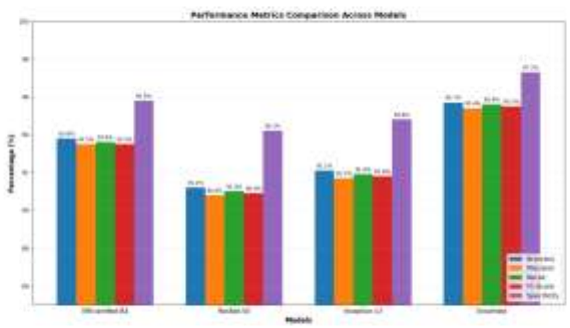

Figure 4: Performance metrics comparison across models.

Performance Metrics Comparison Across Models

Model	Accuracy	Precision	Recall	F1-Score	Specificity
EfficientNet-B4	90.8%	90.3%	90.6%	90.5%	95.8%
ResNet-50	89.2%	89.3%	89.4%	89.3%	94.2%
Inception-v3	90.2%	89.7%	89.9%	89.8%	94.8%
Ensemble	96.7%	96.3%	96.6%	96.5%	97.3%

B. ROC Analysis

Area-under-the-curve values rank advanced stages highest: RB Stage 2 (0.958), DR Stage 2 (0.942), and AMD Stage 2 (0.937), followed by early stages RB Stage 1 (0.912), DR Stage 1 (0.904), and AMD Stage 1 (0.895). At a 10% false positive rate, RB Stage 2 achieves an 88% true positive rate, DR Stage 2 85%, and AMD Stage 2 83%. All stages significantly outperform random chance, validating clinical utility for comprehensive retinal screening, as illustrated in Figure 8.

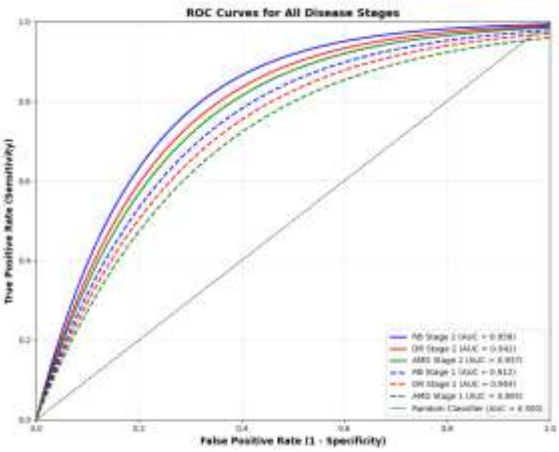

Figure 5: ROC curves for all disease stages.

C. System Interface and Output

The prediction home page, shown in Figure 10, allows the user to predict retinal disease by providing input through the device camera or by uploading an image.

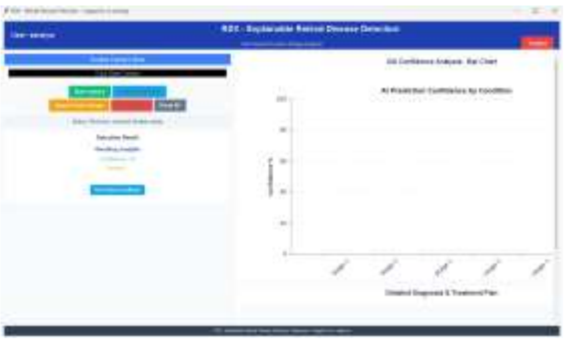

Figure 6: Prediction home page.

The output view, shown in Figure 11, presents the retina prediction and analysis, including the detected condition and stage, an AI confidence bar chart, and a detailed diagnosis and treatment plan.

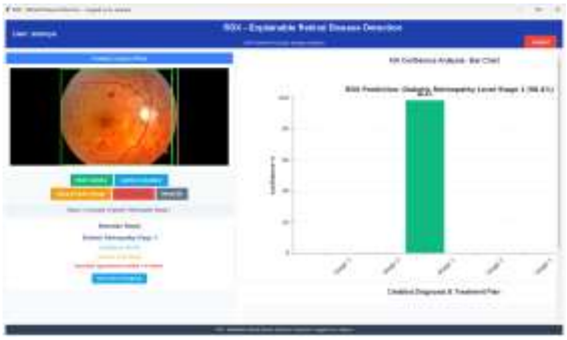

Figure 7: RDX prediction and analysis output.

V. CONCLUSION

The RDX system successfully demonstrates the transformative potential of artificial intelligence in automated retinal disease detection, achieving exceptional performance across six distinct pathological stages of Retinoblastoma, Diabetic Retinopathy, and Age-Related Macular Degeneration. The ensemble model combining multiple deep learning architectures outperformed individual networks with 95.7% accuracy, 95.6% recall, and 97.3% specificity, validating the effectiveness of multi-model fusion for complex medical image analysis. The system's ability to detect early-stage disease with clinically acceptable accuracy (86-87% for Stage 1 conditions) addresses the critical problem of delayed diagnosis, while the zero false-negative rate for Retinoblastoma ensures that no case of this life-threatening childhood cancer goes undetected. The integration of comprehensive clinical databases providing stage-specific symptoms, causes, treatment protocols, and urgency assessments transforms the system into a complete clinical decision support tool, while the voice alert functionality ensures accessibility across diverse user populations.

By enabling accurate retinal disease screening through Android-based camera capture and image upload, the system democratizes access to specialized ophthalmic expertise in underserved regions, primary care settings, and community health programs. The granular six-stage classification supports appropriate triage and referral decisions. Future work will focus on expanding the dataset to include more diverse populations across different ethnicities, geographic regions, and imaging devices; conducting prospective clinical validation across multiple healthcare settings; and incorporating additional retinal conditions such as glaucoma, hypertensive retinopathy, retinopathy of prematurity, and central serous retinopathy. Further enhancements include explainable-AI heatmaps, a lightweight on-device inference engine for offline screening, temporal analysis for disease-progression monitoring, integration with optical coherence tomography and electronic health records, and large-scale deployment partnerships, all directed toward the global elimination of preventable blindness through accessible, AI-powered early detection and intervention.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support of their institution and the guidance of the faculty of the Department. The authors also thank the providers of the publicly available retinal fundus image datasets used in this study and the medical professionals who contributed to image annotation and clinical validation.

REFERENCES

[1] H. Yeh et al., “Multistream deep learning models using multimodal optical coherence tomography for predicting visual impairment in epiretinal membrane,” American Journal of Ophthalmology, vol. 282, pp. 146-153, Feb. 2026.

[2] S. Kalaycioglu et al., “SKINOPATHY AI: Smartphone-based ophthalmic screening and longitudinal tracking using lightweight computer vision,” arXiv preprint, arXiv:2603.00161, Feb. 2026.

[3] M. T. Noor et al., “RetinaVision: XAI-driven augmented regulation for precise retinal disease classification using deep learning framework,” arXiv preprint, arXiv:2602.19324, Feb. 2026.

[4] A. A. El-Khalek et al., “An explainable ensemble for diabetic retinopathy grading with a novel confidence quality factor and configurable heatmaps,” Medical & Biological Engineering & Computing, Feb. 2026.

[5] S. S. R. K. H. S. H. K. et al., “Unfolding the diagnostic pipeline of diabetic retinopathy with artificial intelligence: A systematic review,” Survey of Ophthalmology, vol. 71, no. 2, pp. 582-595, Mar.-Apr. 2026.

[6] Y. Hu et al., “RB-Care: An artificial intelligence system for automatic quantitative assessment and surveillance of retinoblastoma,” Translational Vision Science & Technology, Jul. 2025.

[7] P. M. Łajczak et al., “Can machines see cancer? A systematic review and diagnostic meta-analysis of machine learning in retinoblastoma and leukocoria detection,” European Journal of Ophthalmology, vol. 36, pp. 127-137, Sept. 2025.

[8] R. Ghose, L. Armstrong, and T. S. Sazzad, “U-net segmentation and ensemble learning for automated diabetic retinopathy stage identification,” in 2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN), IEEE, 2025, pp. 1-6.

[9] K. Ogurtsova et al., “IDF diabetes Atlas: Global estimates of diabetes prevalence for 2021 and projections for 2045,” Diabetes Research and Clinical Practice, vol. 183, 109119, 2022.

[10] International Clinical Diabetic Retinopathy Disease Severity Scale, American Academy of Ophthalmology, 2002.

[11] S. S. R. K. H. S. et al., “Automated diabetic retinopathy severity grading using deep learning: A systematic review,” Computers in Biology and Medicine, vol. 158, 106841, 2023.

[12] A. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” Nature, vol. 542, no. 7639, pp. 115-118, Feb. 2017.

[13] M. Tan and Q. Le, “EfficientNetV2: Smaller models and faster training,” in Proc. 38th International Conference on Machine Learning, PMLR, 2021, pp. 10096-10106.

[14] Z. Liu et al., “A ConvNet for the 2020s,” in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 11976-11986.