IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

A HOLISTIC SURVEY ON RETINAL DISEASE DETECTION WITH DEEP LEARNING

A Comprehensive Review of Deep Learning Techniques for Ocular and Systemic Disease Detection Using Retinal Imaging

[1]P Srilakshmi, [2]P Rajesh, [3]Nagesh Makanapur, [4]Mohan Kumar H S, [5]Manoj M H G

[1]Assistant Professor, [2]Student, [3] Student, [4] Student, [5] Student

[1] Computer Science And Engineering

[1] HKBK College Of Engineering Bengaluru, India

Abstract: This paper focuses on advances in deep learning for early detection of eye disease using images of the retinal fundus. It analyzes retinal images using convolutional neural net- works (CNN's) and other algorithms to classify ocular diseases, such as diabetic retinopathy, glaucoma, hypertensive retinopathy, macular edema, and cataracts, as well as systemic conditions such as cardiovascular diseases. We explore various method- ologies, including data preprocessing, model architectures, and performance metrics, along with challenges like datasets diversity and model generalizability. Proposed future directions include improved preprocessing and datasets, vision transformers, and explainable AI to improve access for researchers, scientists, and practitioners in medical imaging and artificial intelligence. These advances provide comprehensive methodologies, best practices, and information for retinal image classification, data preprocessing, model selection, performance evaluation, and clinical applications. Future deep learning developments will support model deployment in real-world healthcare, integrating hybrid models such as MGSCNN (Multi Glowworm Swarm Convolutional Neural Networks) and pre-trained data to aid early detection of retinal diseases.

***Index Terms* - Deep Learning, Retinal Imaging, Disease Detection, Convolutional Neural Networks, Medical Imaging.**

I. INTRODUCTION

Retinal imaging, including fundus photography and optical coherence tomography (OCT), provides a non-invasive view of the retina, the only part of the human body where blood vessels can be directly observed with high-resolution cameras. This makes it a vital tool for detecting ocular diseases, such as diabetic retinopathy (DR), glaucoma, and age-related macular degeneration (AMD), and systemic conditions such as cardiovascular disease (CVD). Deep learning, particularly through convolutional neural networks (CNNs), has revolutionized automated retinal image analysis, achieving high accuracy and reducing the reliance on manual expert diagnosis.

The accessibility of the retina enables the early detection of diseases through vascular or tissue changes. For example, microaneurysms and hemorrhages indicate DR, optic nerve damage signals glaucoma, and vessel tortuosity reflects CVD risk. Automating these diagnoses with deep learning improves efficiency and accessibility, especially in resource-limited settings with few experienced ophthalmologists. Integrating deep learning with retinal imaging promises to transform medical diagnostics, enabling timely interventions and better patient outcomes. This survey examines studies using deep learning for the detection of retinal disease, analyzing their methodologies, results, advantages, limitations, and future directions to provide a comprehensive overview for researchers and clinicians.

This retinal image highlights key anatomical features and signs of diabetic retinopathy, including the optic disk, major and minor blood vessels, and microaneurysms. Pathological signs, such as hemorrhages, soft exudates, and hard exudates, are also marked. These indicators help diagnose and classify the severity of retinal disease [1].

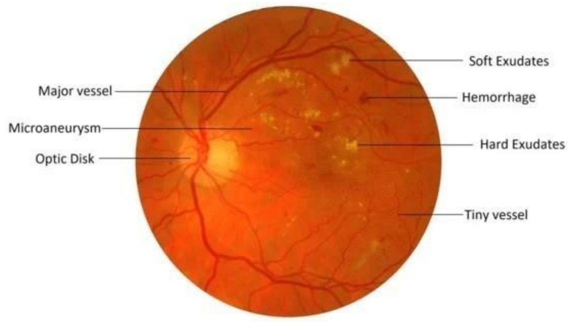

Figure 1: Features of Retinal Fundus Image [1]

The paper is organized as follows: Section II provides background on retinal imaging, Section III reviews the literature, Section IV discusses methodologies, Section V presents a discussion, Section VI explores innovative dataset training strategies, Section VII addresses drawbacks and challenges, Section VIII proposes methodologies and future directions, Section IX concludes the survey, and Section X lists references.

II. BACKGROUND

Retinal imaging techniques capture detailed views of the retina, optic disc, and macula, allowing the diagnosis of various conditions. Fundus photography produces two-dimensional images, revealing vascular and surface abnormalities, while OCT provides cross-sectional views of the retinal layers, highlighting structural changes. These images are critical for detecting symptoms [1, 3, 4]:

- **Diabetic Retinopathy (DR):** A diabetes complication marked by microaneurysms, hemorrhages, and exudates in the retina, caused by damaged blood vessels due to chronic high blood sugar.
- **Glaucoma:** A chronic condition that damages the optic nerve, often due to high intraocular pressure, causing irreversible vision loss and an enlarged cup-to-disc ratio, requiring regular screenings.
- **Cataracts:** A condition characterized by lens opacity that clouds the eye's natural lens, leading to blurred or reduced vision.
- **Age-Related Macular Degeneration (AMD):** Marked by drusen deposits beneath the retina and choroidal neovascularization, disrupting central vision and causing significant visual impairment.
- **Cardiovascular Diseases (CVD):** Reflected in retinal vasculature changes, such as vessel narrowing or tortuosity, serving as biomarkers for systemic health.

Early detection is essential to prevent vision loss or systemic complications, making automated deep learning systems increasingly valuable. The ability to analyze large volumes of retinal images quickly and accurately positions deep learning as a transformative technology in medical diagnostics [3].

III. LITERATURE REVIEW

These studies are categorized by their focus: retinal diseases and systemic diseases. Table 1 summarizes key aspects, including diseases addressed, models used, datasets, and performance metrics.

A. Detection of Ocular Diseases

The studies primarily focus on ocular diseases, using deep learning to classify conditions based on retinal image features. Figure 2 illustrates the accuracy of different deep learning models for classifying retinal diseases, highlighting their performance metrics. Most models achieve high accuracy, ranging from 90% to 99%, indicating robust predictive capabilities. However, the custom 16-layer convolutional neural network (CNN) records a slightly lower accuracy of 89.6%, suggesting it is less effective than other models. These variations may stem from differences in model architecture, training data, or optimization techniques. Such insights are crucial for selecting optimal models for automated retinal disease diagnosis, enhancing clinical decision-making.

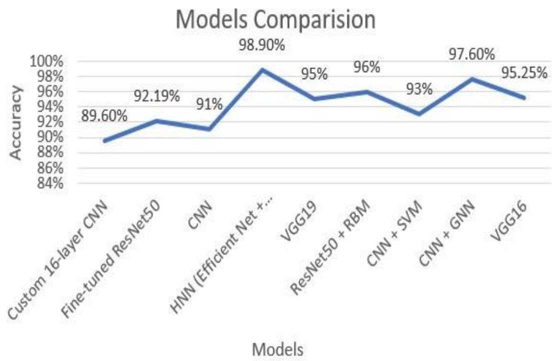

Figure 2: Comparison of Model Accuracies for Retinal Disease Detection [1]

1) Diabetic Retinopathy (DR)

Diabetic retinopathy, a diabetes complication, is a common focus due to its prevalence and severe consequences if untreated.

- **Automated Retinal Disease Classification Using Fine-Tuned ResNet50 [2]:** This study classifies five retinal diseases (AMD, DR, diabetic neuropathy (DN), macular hole (MH), and optic disc cupping (ODC)) using a fine-tuned ResNet50 model. It achieved 92.19% accuracy, with precision ranging from 0.81 to 0.97 and recall from 0.85 to 1.00. The dataset included annotated fundus images, pre-processed with resizing to 224 × 224 pixels and normalization.
- **Early Detection and Classification of Eye Diseases Using Deep Learning Techniques [4]:** Compared DenseNet and ResNet architectures for classifying DR, cataracts, and glaucoma, with DenseNet achieving 95% accuracy on a Kaggle dataset of 4,217 retinal images. This highlights the model's potential for early diagnosis, crucial for preventing vision impairment.
- **Empirical Evaluation of Retinal Disease Identification [5]:** Developed a hybrid neural network (HNN) combining Efficient Net and SVM, achieving 98.9% accuracy for DR, cataracts, and glaucoma on a dataset of 4,000 images.
- **A Combined ResNet50 and Restricted Boltzmann Machine [7]:** Proposed a hybrid model for multi-label classification, including DR, achieving an F1-score of 96% on datasets like CTEH and Messidor.
- **Hybrid CNN and Support Vector Machine Approach [8]:** Employed a CNN-SVM model for multiple retinal diseases, including DR, with over 93% accuracy on 2,000 images.
- **Graph-Aware Multimodal Deep Learning [9]:** Introduced DR-Diag, integrating CNNs and graph neural networks (GNNs) for DR classification, achieving 97.6% to 98% accuracy on APTOS2019 and Messidor-2 datasets.

2) Glaucoma

Glaucoma, characterized by optic nerve damage, is a critical focus due to its asymptomatic nature in early stages.

- **Early Eye Disease Detection Using Machine Learning [3]:** Developed a CNN-based system for detecting glaucoma, DR, and cataracts, achieving 91% accuracy on a Kaggle dataset of 500 images. A graphical user interface (GUI) facilitated real-time diagnosis [3, 4, 5, 8].
- Studies [4, 5, 7, 8] included glaucoma in their classification tasks, using DenseNet, HNN, ResNet50+RBM, and CNN-SVM models, respectively.

3) Cataracts

Cataracts, marked by lens opacity, affect vision clarity and are prevalent in aging populations.

- **AI-Driven Eye Disease Detection: A VGG19-Based Approach [6]:** Focused on cataract detection using VGG19, achieving 95% accuracy on a Kaggle dataset of 2,112 images. The model used transfer learning with fine-tuned pretrained weights.
- Studies [3, 4, 5] addressed cataracts alongside other diseases, reporting high accuracies.

4) **Other Ocular Diseases**

Several studies tackle additional retinal conditions.

- Study [2] classified AMD, DN, MH, and ODC, demonstrating ResNet50's versatility.
- Study [7] included macular edema, macular degeneration, and retinal vascular occlusion in its multilabel framework.
- Automated Retinal Disease Detection: VGG16-Based Approach for OCT Images [10]: Classified choroidal neovascularization (CNV), diabetic macular edema (DME), and drusen from OCT images using VGG16, achieving 95.25% accuracy on a dataset of 84,495 images.

Table1: Summary of Reviewed Studies

Ref.	Diseases	Models	Datasets	Accuracy
[1]	Cardiovascular Diseases	Custom 16-layer CNN	Messidor, e-Ophtha, etc. (8,752 images)	89.6%
[2]	AMD, DR, DN, MH, ODC	Fine-tuned ResNet50	Annotated fundus images	92.19%
[3]	DR, Glaucoma, Cataracts, etc.	CNN	Kaggle (500 images)	91%
[4]	DR, Cataract, Glaucoma	HNN (EfficientNet + SVM)	Eye Diseases Classification (4,000 images)	98.9%
[5]	Cataracts	VGG19	Kaggle (2,112 images)	95%
[6]	Multilabel Eye Diseases	ResNet50 + RBM	CTEH, Messidor, EYEPACS (3,785 images)	F1: 96%
[7]	Retinal Diseases	CNN + SVM	2,000 images	93%
[8]	Diabetic Retinopathy	CNN + GNN	APTOS2019, Messidor-2	97.6%–98%
[9]	CNV, DME, Drusen	VGG16	Kaggle OCT (84,495 images)	95.25%

B. Detection of Ocular Diseases

The studies primarily focus on ocular diseases, using deep learning to classify conditions based on retinal image features. Figure 2 illustrates the accuracy of different deep learning models for classifying retinal diseases, highlighting their performance metrics. Most models achieve high accuracy, ranging from 90% to 99%, indicating robust predictive capabilities. However, the custom 16-layer convolutional neural network (CNN) records a slightly lower accuracy of 89.6%, suggesting it is less effective than other models. These variations may stem from differences in model architecture, training data, or optimization techniques. Such insights are crucial for selecting optimal models for automated retinal disease diagnosis, enhancing clinical decision-making.

VI. METHODOLOGIES

The studies employ diverse deep learning methodologies, categorized into data preprocessing, model architectures, training strategies, and evaluation metrics.

A. Data Preprocessing

Preprocessing ensures images are suitable for model input, addressing variations in resolution, lighting, and noise:

- **Resizing:** Images are resized to 224 × 224 pixels to match model input requirements [1, 2, 4, 6, 10].
- **Normalization:** Pixel values are scaled to [0,1] or standardized to zero mean and unit variance [1, 2, 4, 10].
- **Augmentation:** Techniques like random rotation, flipping, zooming, and shifting enhance dataset diversity and prevent overfitting [1, 2, 4, 5, 8, 10]. Study [1] applied random rotations, shifts, shearing, zooming, and flipping to increase robustness.
- **Contrast Enhancement:** Some studies, like [3], applied histogram equalization to improve image contrast, aiding feature detection.

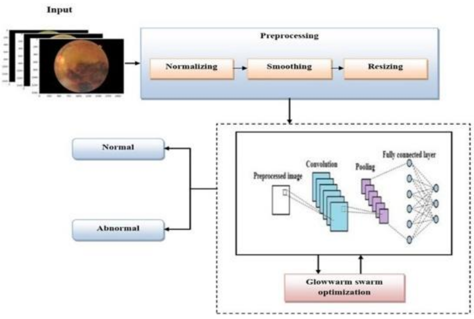

Figure 3: CNN Architecture for Retinal Disease Detection [1]

B. Training Strategies

Training strategies focus on optimizing model parameters to minimize loss and enhance performance, with optimizers playing a critical role.

Optimizers: Adam (Adaptive Moment Estimation), introduced by Kingma and Ba in 2014, combines the advantages of stochastic gradient descent (SGD) extensions AdaGrad and RMSProp. Adam is widely used in training deep neural networks due to its efficiency, low memory requirements, and ability to handle sparse gradients and noisy data.

C. Evaluation Metrics

Performance is assessed using:

- **Accuracy:** Overall correctness of predictions, reported in all studies (e.g., 98.9% in [5]).
- **Precision, Recall, F1-Score:** Measure class-specific performance, critical for imbalanced datasets [1, 3, 5, 7, 10].
- **AUC-ROC:** Evaluates model discrimination, reported in [1] (AUC: 97%).

V. DISCUSSION

The reviewed methodologies highlight the diversity of machine learning approaches in healthcare. Deep learning [1] excels in medical image analysis, achieving high accuracy in tasks like pneumonia detection. However, its computational costs, requiring specialized hardware, limit accessibility for smaller institutions. Additionally, its lack of interpretability poses a barrier to clinical adoption, as clinicians need transparent decision-making processes. Ensemble methods [2] offer robustness for predictive tasks, such as hospital read-mission forecasting, by combining multiple algorithms to handle imbalanced data. However, their computational complexity and sensitivity to noisy inputs, such as inconsistent EHR data, reduce efficiency. Federated learning [3] addresses privacy concerns critical for compliance with regulations like GDPR, enabling collaborative model training without sharing sensitive data. However, it introduces challenges in communication overhead and security risks, such as model poisoning by malicious participants.

A. Comparative Analysis

A detailed comparison reveals key trade-offs:

- **Performance vs. Efficiency:** Deep learning [1] achieves superior accuracy (e.g., 92% for image classification) but requires significant computational resources, with training times often exceeding 24 hours on high-end GPUs. Ensemble methods [2] balance performance (AUC of 0.85) and robustness but need hours for training and are unsuitable for real-time applications. Federated learning [3] prioritizes privacy, resulting in moderate performance (F1-score of 0.78) due to constraints like decentralized data and privacy-preserving mechanisms.
- **Interpretability:** All three approaches lack sufficient explainability. Deep learning models [1] are black-box systems, offering no insight into feature importance. Ensemble methods [2] provide some feature ranking but remain complex for non-experts. Federated learning [3] complicates interpretability due to distributed model updates. Techniques like SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model-agnostic Explanations) could address this gap by providing post-hoc explanations.
- **Scalability:** Federated learning [3] is inherently scalable across institutions, supporting collaborative healthcare systems. However, it requires robust network infrastructure for frequent model updates. Deep

learning [1] and ensemble methods [2] are better suited to centralized systems with access to large datasets and computational resources, limiting scalability in resource-constrained settings.

- **Ethical Considerations:** Deep learning risks amplifying biases in training data, such as under representation of minority groups in medical images, leading to inequitable outcomes. Ensemble methods [2] may propagate biases through feature selection, excluding rare conditions. Federated learning [3] mitigates data-sharing ethical concerns but raises fairness issues if participating institutions have unequal computational capabilities.

B. Research Gaps and Opportunities

The analysis identifies critical research gaps:

- **Lightweight Models:** Developing efficient machine learning models that operate on low-resource hardware is essential for democratizing access, especially in underserved regions. Techniques like model pruning or quantization could reduce the computational burden of deep learning [1] and ensemble methods[2].
- **Explainable AI:** Enhancing interpretability is crucial for clinical trust. Attention mechanisms, rule-based explanations, or hybrid models combining machine learning with decision trees could address the black-box nature of deep learning [1] and ensemble methods.

C. Broader Implications

The adoption of machine learning in healthcare has profound implications:

Patient Care: Accurate models improve diagnosis and treatment. Standardizing evaluation metrics (e.g., fairness-adjusted accuracy, interpretability scores) and benchmarks could accelerate progress. Collaborative initiatives, such as open-source datasets and model repositories, would make significant progress.

VI. INNOVATIVE DATASET TRAINING STRATEGIES FOR RETI- NAL DISEASE DETECTION

Advancements in deep learning for retinal disease detection rely on robust dataset training strategies to overcome limitations in data availability and model generalizability. While preprocessing and transfer learning are standard practices [1, 2, 4], emerging approaches like active learning and domain adaptation offer new ways to improve dataset efficiency and model performance. These strategies address key issues in dataset diversity and annotation costs, providing a foundation for tackling the challenges discussed in Section VII.

A. Active Learning for Optimized Annotation

Active learning (AL) streamlines dataset training by selectively annotating the most informative retinal images, reducing the need for large labeled datasets. In AL, a model identifies images with high predictive uncertainty (e.g., ambiguous cases of diabetic retinopathy) for expert review, optimizing the annotation process. This is critical for retinal imaging, where expert labeling is resource-intensive [3]. For instance, a study using a CNN-based system on a Kaggle dataset of 500 images [3] could benefit from AL to prioritize labeling complex cases, such as early-stage glaucoma, improving model accuracy with fewer annotations. AL is particularly effective for rare conditions like macular holes, ensuring underrepresented samples are included in training [2].

By focusing on high-impact images, AL addresses the dataset diversity concerns raised in [1], where models trained on specific populations showed limited generalizability. It also mitigates overfitting risks noted in [2], as selective annotation introduces varied data distributions. AL could improve clinical dataset creation, enabling efficient training in settings with limited access to ophthalmologists.

B. Domain Adaptation for Robust Generalization

Domain adaptation (DA) enhances model performance across diverse retinal datasets, which often differ in imaging modalities or patient demographics [1]. DA aligns feature distributions between a source dataset (e.g., Messidor) and a target dataset (e.g., APTOS2019) to ensure models generalize to new clinical environments. For example, the multimodal approach in [9], which integrated CNNs and graph neural networks, could leverage DA to adapt features across fundus and OCT images. A DA-enhanced model could improve the 95.25% accuracy reported for VGG16 on OCT images [10] by aligning features from varied imaging devices. DA directly tackles the demographic bias issue highlighted in [1], where dataset skew reduced model applicability. By adapting to diverse datasets, DA ensures robust performance across patient groups, enhancing clinical reliability. It also reduces overfitting, as observed in [2], by training models on broader data distributions. Combining DA with federated learning [3] could further enable privacy-preserving model training across institutions, a promising direction for future research.

C. Implications for Clinical Adoption

Active learning and domain adaptation offer significant potential for retinal disease detection. AL reduces annotation costs, making dataset creation feasible in resource-limited settings, while DA ensures models perform consistently across diverse clinical contexts. These strategies address critical limitations, such as dataset diversity and computational efficiency, setting the stage for the challenges discussed in Section VII. By adopting these innovative approaches, researchers can develop more scalable and clinically viable deep learning models, advancing automated diagnostics for ocular and systemic diseases.

D. Example of a Sample Trained Dataset

The MESSIDOR dataset comprises 1,200 color fundus images captured using a Topcon TRC NW6 non-mydriatic retinography with a 45-degree field of view. These images are categorized into four diabetic retinopathy severity levels:

- Grade 0: No DR
- Grade 1: Mild
- Grade 2: Moderate
- Grade 3: Severe

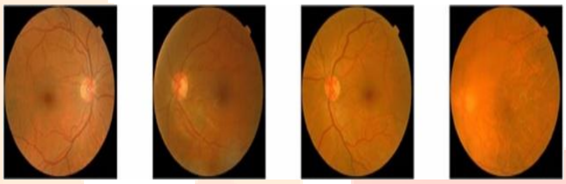

Figure 4: Sample Image from MESSIDOR Dataset for Diabetic Retinopathy Detection

VII. DRAWBACKS AND CHALLENGES

Despite promising results, the studies face several limitations that hinder clinical adoption:

1. **Dataset Diversity:** Many datasets lack demographic diversity (e.g., age, ethnicity, gender), limiting model generalizability. Study [1] noted that its dataset primarily included images from specific populations, potentially skewing results. Similarly, [3] used a small dataset (500 images), insufficient for robust training.
2. **Overfitting:** Deep models, especially complex ones like ResNet50, show overfitting after extended training. Study [2] reported training accuracy reaching 100% while validation accuracy plateaued at 92.19%, indicating overfitting beyond the 10th epoch.
3. **Clinical Validation:** Most studies lack real-world clinical testing, raising concerns about practical utility. Studies [2, 3] emphasized the need for validation in clinical settings to ensure reliability.
4. **Interpretability:** The black-box nature of deep learning models poses challenges for clinical trust. Study [3] highlighted the need for explainable AI to provide insights into model decisions, such as feature importance.
5. **Computational Resources:** Training large models like VGG19 [6] or DenseNet [4] requires significant computational power, which may be impractical for resource-constrained environments.

These challenges underscore the need for improved datasets, model robustness, and clinical integration to bridge the gap between research and practice.

VIII. PROPOSED METHODOLOGIES AND FUTURE DIRECTIONS

To address the identified challenges, future research should focus on the following areas:

1. **Larger and Diverse Datasets:** Collecting datasets with broader demographic representation (e.g., diverse ethnicities, ages, and disease severities) will enhance model generalizability. Collaborative efforts to create global retinal image repositories could address the limitations noted in [1, 3].

2. **Advanced Architectures:** Exploring vision transformers, which excel at capturing global image context, could outperform traditional CNNs. Transformers have shown promise in medical imaging [12], offering potential for improved feature extraction in retinal analysis.
3. **Explainable AI:** Developing interpretable models, such as attention-based networks or gradient-weighted class activation mapping (Grad-CAM), will enhance clinical trust. Study [3] emphasized explainability to make model decisions transparent to clinicians.
4. **Multimodal Integration:** Combining retinal images with patient metadata (e.g., age, medical history, blood pressure) could improve diagnostic accuracy. Multimodal deep learning, integrating image and tabular data, is a promising direction for comprehensive diagnostics.
5. **Federated Learning:** To address privacy concerns and dataset limitations, federated learning could enable model training across multiple institutions without sharing sensitive patient data. This approach could create robust models while maintaining compliance with regulations like HIPAA.
6. **Clinical Trials:** Conducting large-scale clinical trials to validate models in real-world settings will ensure practical utility. Partnerships with hospitals and clinics could facilitate this, addressing the validation gap noted in [2].

These directions aim to enhance the robustness, interpretability, and clinical applicability of deep learning systems for retinal disease detection.

IX. CONCLUSION

Deep learning has significantly advanced the automated detection of diseases using retinal images, achieving high accuracies for ocular conditions like diabetic retinopathy, glaucoma, and cataracts, as well as systemic conditions like cardiovascular diseases. The reviewed studies demonstrate the effectiveness of CNNs, pre-trained models like ResNet50 and VGG19, and hybrid architectures in analyzing fundus and OCT im- ages. However, challenges such as limited dataset diversity, overfitting, lack of clinical validation, and model interpretability persist, hindering widespread adoption.

Future research should prioritize larger, more diverse datasets, advanced architectures like vision transformers, and explainable AI to address these limitations. Multimodal integration and federated learning offer additional opportunities to enhance diagnostic accuracy and scalability. By overcoming these challenges, deep learning can transform medical diagnostics, enabling early detection, improving patient outcomes, and reducing the burden on healthcare systems. This survey provides a foundation for researchers and clinicians to build upon, advancing the field toward practical, impactful solutions.

X. REFERENCES

References

[1] M. K. A. Hawkins, "Classification of Cardiovascular Diseases through Retinal Fundus Images using Deep Learning," 2025.

[2] R. Rani, "Automated Retinal Disease Classification Using Fine-Tuned ResNet50: A Deep Learning Approach," 2025.

[3] A. Mishra, N. Choubey, P. Jumde, A. Samiya, A. Garg, and A. Kale, "Early Eye Disease Detection Using Machine Learning," 2025.

[4] L. Rahunathan and D. Sivabalaselvamani, "Early Detection and Classification of Eye Diseases Using Deep Learning Techniques," 2025.

[5] N. Meenakshisundaram and S. G, "Empirical Evaluation of Retinal Disease Identification Based on Deep Learning Approaches," 2025.

[6] A. Sharma and S. Gupta, "AI-Driven Eye Disease Detection: A VGG19-Based Approach to Retinal Image Classification," 2025.

[7] N. Fadhillah and I. Nurtanio, "A Combined ResNet50 and Restricted Boltzmann Machine for Multilabel Eye Diseases Classification," 2025.

[8] Vishal and S. Mehta, "Hybrid Convolutional Neural Network and Support Vector Machine Approaches for Enhanced Retinal Disease Detection," 2024.

[9] A. Zedara, "Graph-Aware Multimodal Deep Learning for Retinal Disease Detection," 2025.

[10] P. Kaushik and S. Choudhary, "Automated Retinal Disease Detection: VGG16-Based Deep Learning Approach for OCT Images," 2025.

[11] S. Muchuchuti and S. Viriri, "Retinal Disease Detection Using Deep Learning Techniques: A Comprehensive Review," https://www.mdpi.com/2313433X/9/4/84, 2023.

[12]A. K. Kakumani and K. Tanuja, “Vision Transformers for Retinal Disease Classification using Optical Coherence Tomography Images,” 2025.

[13]N. Krishna and A. D. V. A, “RETINOSCAN: Retinal Image-based Diabetic Retinopathy Detection,” 2025.

[14]H. Gavel and C. Azad, “Retinal Disease Detection for Low Resource Devices Through Pruned and Quantized CNNs,” 2025.

Quantized CNNs,” 2025.