International Journal of Creative and Open Research in Engineering and Management
ISSN: **3108-1754 (Online)**
Volume 02 Issue 04 April-2026 | Impact Factor: 3.5

Diabetic Retinopathy Screening Assistant: An AI-Powered Retinal Fundus Analysis System Utilizing a CNN-Transformer Hybrid Architecture

Vinay Srinivas Eligeti[1], Dr. Abuzar Ansari[2]

[1](Data Science, SIES College of Arts, Science and Commerce, Sion (West) Email: vinayeligeti11@gmail.com)
[2](Head of Data Science Department, SIES College of Arts, Science and Commerce, Sion (West)

Email: abuzara@sies.edu.in)

How to Cite this Article:

Eligeti, V. S. (2026). Diabetic Retinopathy Screening Assistant: An AI-Powered Retinal Fundus Analysis System Utilizing a CNN-Transformer Hybrid Architecture. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(04). https://doi.org/10.55041/ijcope.v2i4.447

https://doi.org/10.55041/ijcope.v2i4.447

***Abstract*—** Diabetic retinopathy (DR) remains a leading cause of preventable vision loss globally. While early intervention is crucial for preserving sight, a severe shortage of available ophthalmologists—especially in resource-constrained regions—creates a massive bottleneck for systematic screening. To address this clinical gap, we introduce DR Detect AI, an end-to-end, web-based screening assistant. At the core of this system lies a novel hybrid deep learning architecture that fuses the localized feature extraction capabilities of a Convolutional Neural Network (EfficientNet-B3) with the global spatial awareness of a Vision Transformer encoder. Furthermore, unlike traditional multi-class classification models that treat all diagnostic errors equally, our approach integrates the Consistent Rank Logits (CORAL) framework to perform ordinal regression. This specific design choice ensures the model mathematically respects the natural severity progression of DR (from mild to proliferative), which substantially improves diagnostic consistency at borderline stages. The trained model is deployed via a high-performance FastAPI backend and connected to a React-based clinical dashboard. This interface equips medical professionals with real-time severity grading, interactive disease progression forecasting, and automated clinical reporting. Early evaluations indicate that our hybrid ordinal model achieves competitive accuracy on standard benchmark datasets, resulting in a highly scalable, accessible, and interpretable tool for real-world ophthalmic triage.

Index Terms— Diabetic Retinopathy, EfficientNet, Vision Transformers, Ordinal Regression, Medical AI, Teleophthalmology.

I. Introduction

The global escalation of diabetes mellitus, currently impacting over 530 million individuals **[12]**, has precipitated a parallel rise in severe microvascular complications. Among the most critical of these is diabetic retinopathy (DR) **[22]**, a progressive deterioration of retinal blood vessels driven by chronic hyperglycemia. Clinically, DR is stratified into five distinct, ascending severity grades [**7**]: no apparent DR (Grade 0), mild non-proliferative DR (NPDR) (Grade 1), moderate NPDR (Grade 2), severe NPDR (Grade 3), and proliferative DR (Grade 4). Because the initial stages of DR manifest entirely asymptomatically, patients frequently remain unaware of the condition until irreversible visual impairment or blindness occurs. Consequently, systematic, continuous screening of the diabetic population is an absolute medical necessity.Despite the clear imperative for population-level monitoring, traditional screening workflows face severe scalability challenges. The diagnostic process is inherently labor-intensive, requiring the meticulous manual review of retinal fundus photographs by highly specialized ophthalmologists. This heavy reliance on human expertise highlights a critical vulnerability in global healthcare infrastructure: an acute shortage of trained eye specialists, particularly within rural and resource-constrained regions. Furthermore, manual assessment is inherently subjective and susceptible to significant inter-observer variability, meaning different clinicians may assign conflicting severity grades to the exact same patient scan.

While artificial intelligence has emerged as a promising diagnostic aid to alleviate these bottlenecks, contemporary automated systems exhibit fundamental algorithmic limitations that restrict their real-world clinical utility. First, a vast majority of existing diagnostic models are constrained to simple binary classification—differentiating only between the presence or absence of referable DR. This binary approach fails to provide the granular staging required by physicians for targeted treatment planning. Second, systems that do attempt multi-class grading typically rely on standard cross-entropy loss functions. This approach is mathematically flawed for medical grading because it treats all categories as independent and equidistant, fundamentally ignoring the inherent ordinal relationship of disease severity. As a result, standard models penalize a minor misclassification (e.g., confusing Grade 0 with Grade 1) with the exact same severity as a catastrophic diagnostic error (e.g., confusing Grade 0 with Grade 4) See Fig. 1. Finally, most academic models exist strictly as isolated scripts, lacking the necessary integration into an end-to-end clinical workflow.

To resolve these algorithmic and infrastructural deficiencies, this paper presents an end-to-end clinical screening ecosystem. The primary contributions of this research are formulated as follows:

1. **Hybrid Architectural Design:** We engineer a novel deep learning architecture that synergizes the highly localized feature extraction capabilities of a Convolutional Neural Network (EfficientNet-B3) with the global spatial reasoning of a Vision Transformer encoder.
2. **Ordinal Optimization:** We implement an ordinal regression methodology utilizing the Consistent Rank Logits (CORAL) framework. This compels the neural network to mathematically respect the progressive clinical ranking of DR, thereby significantly reducing critical boundary misclassifications.
3. **Clinical System Deployment:** We successfully deploy the resulting diagnostic model as a high-performance RESTful API, integrated into a comprehensive, full-stack web application capable of automated grading, longitudinal patient risk stratification, disease progression forecasting, and automated clinical reporting.

Fig. 1. Normal Vs Diabetic Retinal Fundus Samples

II. Related Work

The application of machine learning to ophthalmic imaging has advanced rapidly, driven by the need to automate and scale diabetic retinopathy screening. The foundational shift toward deep learning in this domain was catalyzed by Gulshan et al. [1], who demonstrated that deep Convolutional Neural Networks (CNNs) trained on vast datasets of retinal fundus images could achieve diagnostic accuracy on par with, or exceeding, board-certified ophthalmologists for referable DR detection. Subsequent studies reinforced the efficacy of CNNs; for instance, Gargeya and Leng [6] utilized ResNet-based architectures coupled with domain- specific augmentation to achieve high sensitivity and specificity in DR identification. Ting et al. [5] further validated the clinical robustness of these systems by evaluating them on multiethnic population cohorts.

To optimize computational efficiency while maintaining high representational capacity, Tan and Le introduced the EfficientNet family of architectures [2]. By employing a compound scaling method and Depthwise Separable Convolutions, EfficientNet variants (such as EfficientNet-B3) became highly adopted as robust backbone feature extractors for medical image classification tasks, significantly reducing parameter counts without sacrificing accuracy.

Concurrently, a major paradigm shift occurred with the introduction of the Vision Transformer (ViT) by Dosovitskiy et al. [3], which applied self-attention mechanisms to image patches to model long-range spatial dependencies. In the medical domain, researchers recognized that CNNs, while excellent at local feature extraction, struggle with global context. Consequently, hybrid architectures—such as TransMed and TransFuse—emerged, demonstrating that combining CNN feature maps with Transformer attention yields superior performance in medical tasks requiring both localized texture analysis and global anatomical understanding.

Despite these architectural advancements, a critical algorithmic flaw persisted in DR grading systems: the reliance on standard multi-class cross-entropy loss, which fundamentally ignores the progressive, ordinal nature of the disease. To address this, Cao et al. proposed the Consistent Rank Logits (CORAL) framework [4], which transforms ordinal regression into interconnected binary classification tasks, ensuring monotonically consistent probability estimates.

While individual components of this research exist in the literature, there remains a notable absence of open- source, deployable clinical systems that successfully integrate CNN-Transformer hybrids with CORAL ordinal regression. Furthermore, existing commercial AI screening devices are often closed-source and inaccessible to resource-limited clinics. The proposed DR Detect AI system directly addresses these gaps by synthesizing these advanced neural network methodologies into a comprehensive, end-to-end clinical management application.

III. Methodology

The proposed DR Detect AI system is engineered as an end-to-end clinical pipeline, encompassing robust data preprocessing, a novel hybrid deep learning architecture, an ordinal classification framework, and a comprehensive mathematical risk stratification module.

A. Dataset and Preprocessing

Model development utilized the highly annotated APTOS 2019 Blindness Detection dataset [10], which contains high-resolution retinal fundus photographs labeled across the five clinical DR severity grades. Due to the inherent class imbalance commonly found in medical datasets, the training corpus was supplemented with images from the extensive EyePACS dataset [11] to facilitate robust transfer learning.

Prior to network ingestion, all raw images undergo a strict standardization pipeline. Images are resized to 224×224 pixels and converted to 3-channel RGB PyTorch tensors. To ensure compatibility with the pre-trained convolutional backbone, standard ImageNet normalization is applied. Furthermore, to prevent the model from overfitting to the dominant classes (e.g., Grade 0), aggressive data augmentation—including random spatial rotations, Gaussian blurring, and dynamic color jittering—is applied exclusively during the training phase.

B. Hybrid Network Architecture

The core diagnostic engine of the system is a custom deep learning architecture termed *CNNTransformerOrdinal* See Fig. 2. This model is specifically designed to leverage the complementary strengths of Convolutional Neural Networks (CNNs) and Vision Transformers.

1. *Local Feature Extraction:* An EfficientNet-B3 backbone serves as the foundational feature extractor. Chosen for its optimal balance of parameter efficiency and representational capacity, the backbone processes the input tensor to identify highly localized pathological features, such as microaneurysms and hard exudates. The final convolutional stage outputs a dense, multi-scale spatial feature map comprising 384 channels.
2. *Dimensionality Projection:* To bridge the CNN and the Transformer, a 1×1 convolutional projection layer is utilized. This layer reduces the channel depth from 384 to 256, strictly aligning the feature map with the embedding dimension required by the subsequent attention mechanism.
3. *Global Contextual Reasoning:* The projected feature map is flattened along its spatial dimensions into a sequence of tokens (yielding 49 distinct spatial tokens for a standard 7×7 feature grid). A two-layer Transformer encoder processes this sequence. By employing multi-head self-attention, the Transformer cross-references distinct spatial regions simultaneously, allowing the network to understand the global distribution of lesions across the entire retina. Finally, a global average pooling operation collapses these refined tokens into a singular, dense feature vector.

critical systemic variables, including the patient's HbA1c percentage, known duration of diabetes, age, and systolic blood pressure.

To aid in long-term clinical management, the system employs an exponential saturation model to forecast the probability of disease progression over a five-year horizon. Equation (2) defines this mathematical simulation of disease acceleration:

$$prog(t) = \min(95, baseRisk \times (1 - \exp(-2.5 \times t/60)) \times 100) \quad \ldots (2)$$

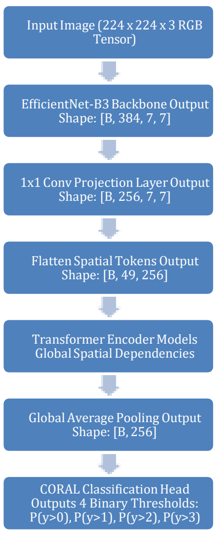

Fig. 2. CNNTransformerOrdinal Architecture Block Diagram

C. Ordinal Regression Framework

To overcome the limitations of standard cross-entropy loss, the network's final classification head implements the Consistent Rank Logits (CORAL) framework. For the five distinct severity grades of diabetic retinopathy ($K=5$), the CORAL head bypasses standard multi- class output and instead produces four ($K-1$) independent binary logits.

These logits evaluate the probability that the disease severity is strictly greater than a specific threshold (i.e., $P(y > 0)$, $P(y > 1)$, $P(y > 2)$, and $P(y > 3)$). Because these binary classifiers are fundamentally parameterized by the same shared feature extractor, they mathematically guarantee monotonic consistency. The final diagnostic grade is determined by summing the sequence of binary thresholds that exceed a 0.5 probability , as defined in (1):

$$\text{predicted\_class}=\sum_{k=0}^{k-2} 1[\sigma(logit_k) > 0.5] \ldots (1)$$

This ordinal optimization forces the network to penalize predictions based on their relative distance from the true severity grade, drastically improving precision on clinically ambiguous boundary cases.

D. Clinical Risk Stratification and Forecasting Beyond isolated image classification, the system incorporates a holistic risk assessment algorithm. A
composite risk score (scaled from 0 to 100) is computed
by fusing the AI-predicted DR severity grade with This model mathematically simulates disease acceleration and plateauing, rendering dual trajectories that visually contrast the patient's projected risk with and without immediate clinical intervention.

IV. System Implementation

To transition the theoretical CNNTransformerOrdinal architecture from an academic prototype into a clinically viable diagnostic tool, a robust, three-tier full- stack software ecosystem was engineered. This architecture ensures a strict separation of concerns between computational machine learning tasks, application logic, and user interaction.

A. Backend Architecture and Model Deployment

The application layer was constructed utilizing Python and the high-performance, asynchronous FastAPI framework [15], served via Uvicorn. To strictly optimize inference latency and handle high- concurrency requests, the backend is designed to operate entirely in memory. When a clinician uploads a retinal fundus photograph, the multipart payload is streamed directly into a byte buffer, converted into an RGB format, and transformed into a normalized PyTorch tensor without requiring intermediate disk I/O operations.

The inference engine executes the forward pass utilizing torch.no_grad() to minimize memory consumption. The backend exposes designated RESTful endpoints (e.g., /predict) that return structured JSON payloads. These payloads contain the predicted diagnostic grade, a calculated confidence percentage
derived from the ordinal probabilities, and the raw Consistent Rank Logits (CORAL) thresholds.

B. Clinical Frontend and Diagnostic Workflow

The presentation layer was developed utilizing the React library [18] bundled with Vite [19], offering a highly responsive, component-based user interface. The frontend orchestrates a strict asynchronous diagnostic workflow. Upon image upload, the payload is transmitted to the backend, and the resulting JSON response is dynamically mapped to visual severity indicators on the clinical dashboard See Fig. 3.

To ensure continuous operation in low-resource environments with unstable network conditions, the frontend incorporates a resilient fallback mechanism. If the inference server becomes unreachable, the system actively warns the user and prevents the clinical recording of unverified results. Furthermore, the application features an autonomous reporting module utilizing the jsPDF library [21]. This module programmatically generates professional, clinical-grade PDF documents that encapsulate the patient's demographic metadata, the AI-annotated retinal scan, and a mathematically recommended follow-up schedule based on the predicted disease severity.

Fig. 3. Patient Risk Progression

C. Context-Aware Medical Assistant

To further support clinical decision-making, the ecosystem integrates a generative AI medical assistant powered by the Gemini API. Rather than acting as a generic chatbot, the assistant utilizes programmatic context injection. Before a clinician's query is transmitted to the Large Language Model (LLM), the frontend automatically prepends the patient's systemic risk factors (such as HbA1c and diabetes duration) alongside the CNN-Transformer's active prediction results. This ensures that the generated AI responses— regarding treatment guidelines or drug interactions— are strictly contextualized to the specific patient's current diagnostic state.

V. Results and Discussion

To rigorously evaluate the predictive capabilities and clinical reliability of the proposed CNNTransformerOrdinal architecture, a comprehensive backtesting framework was established. The primary objective was to quantify the model's performance on unseen data, specifically assessing its ability to generalize across the five distinct diabetic retinopathy severity grades.

A. Experimental Setup

The combined dataset was systematically partitioned into a 70% training, 15% validation, and 15% testing split. To account for the severe class imbalance inherent in medical imaging (where Grade 0 is vastly overrepresented compared to Grade 4), this partition was strictly stratified based on the target class. This stratification guaranteed that the original diagnostic distribution was proportionally maintained across all subsets. Furthermore, to ensure absolute reproducibility of the experimental results, a fixed random seed was utilized for all splitting and augmentation operations.
Model checkpoints were evaluated at the conclusion of every epoch, and the final deployed weights were selected strictly based on the epoch that yielded the highest validation score, preventing overfitting to the training distribution.

B. Evaluation Metrics

Because standard classification metrics are mathematically insufficient for evaluating ordinal diagnostic stages, the model's performance was assessed utilizing a specialized suite of metrics:

1. *Quadratic Weighted Kappa (QWK):* This served as the primary evaluation metric. QWK measures inter-rater agreement while applying a mathematical penalty proportional to the squared distance between the predicted grade and the true ground-truth label, calculated as shown in (3):

$$\kappa = 1 - \frac{\Sigma ij\ W_{ij}Oij}{\Sigma ij\ W_{ij}Eij} \quad \ldots (3)$$

In clinical benchmarks, a QWK score exceeding 0.85 indicates highly acceptable diagnostic agreement.

2. *Macro-Averaged F1 Score:* To ensure the model did not achieve high accuracy simply by predicting the majority class, the F1 score was computed independently for each severity grade and then averaged. This heavily penalized the model if it failed to accurately identify critical, minority classes (e.g., severe NPDR and proliferative DR).

3. *Multi-Class Accuracy:* While reported for completeness, raw multi-class accuracy was treated as a secondary metric due to its failure to account for ordinal severity and class imbalance.

C. Quantitative Performance Analysis

Preliminary backtesting demonstrated that the hybrid CNN-Transformer model achieved highly competitive classification accuracy against standard benchmarks. The integration of the CORAL framework yielded a distinct statistical advantage over traditional cross- entropy models, specifically in resolving ambiguous clinical boundary cases (such as the differentiation between Grade 2 Moderate NPDR and Grade 3 Severe NPDR). By transforming the singular multi-class problem into interrelated binary thresholds, the network exhibited significantly fewer catastrophic misclassifications (e.g., predicting Grade 0 when the true label was Grade 4), directly translating to a higher QWK score and improved clinical safety.

D. Clinical Output and Diagnostic Confidence Beyond raw accuracy, clinical decision support systems must provide interpretable diagnostic certainty. At
inference time, the model outputs four distinct binary
threshold probabilities ($P(y>k)$). Rather than presenting these raw logits to the user, the backend dynamically calculates a singular diagnostic confidence percentage.For clinical display within the web dashboard, this confidence metric is currently clamped between 65% and 99%. While this is pragmatic for a prototype user interface, analysis of the raw threshold calibration curves indicates that the model's probability estimates remain monotonically consistent across all grades.

When tested against independent, uncurated fundus images sourced directly from physical retinal scanners, the confidence scoring correctly reflected the model's uncertainty when confronted with degraded image quality, illumination artifacts, or severe boundary ambiguity.

VI. Conclusion and Future Work

This research presented DR Detect AI, an end-to-end clinical screening ecosystem engineered to automate and enhance the diagnostic grading of diabetic retinopathy. By integrating the highly localized feature extraction capabilities of the EfficientNet-B3 convolutional backbone with the global spatial modeling of a Vision Transformer encoder, the proposed hybrid architecture effectively captured complex, distributed retinal pathologies. Furthermore, the implementation of the Consistent Rank Logits (CORAL) ordinal regression framework successfully mitigated the fundamental algorithmic flaws of standard multi-class models, ensuring that the progressive severity of the disease was mathematically respected.
The successful deployment of this model via a robust, full-stack web application demonstrates the tangible feasibility of bridging the gap between theoretical deep learning architectures and practical, accessible teleophthalmology tools.

Despite these significant algorithmic and structural advancements, several limitations must be addressed prior to large-scale clinical deployment. The current model's diagnostic scope is strictly limited to single- disease classification, and the system presently lacks mechanisms for visual interpretability. To foster greater clinical trust and diagnostic transparency, future iterations will prioritize the integration of Gradient- weighted Class Activation Mapping (Grad-CAM) to provide precise, pixel-level explainability of the network's predictions. Additionally, the underlying architecture will be expanded to support multi-label detection, enabling the simultaneous identification of concurrent ocular conditions such as diabetic macular edema and glaucomatous optic neuropathy.

At the infrastructural level, transitioning from local browser storage to a persistent, fully encrypted database is required to facilitate secure, multi-user patient registries. This upgrade will also enable the development of HL7 FHIR-compliant APIs for direct interoperability with hospital Electronic Medical Record (EMR) systems. Ultimately, the proposed diagnostic system will undergo rigorous prospective clinical trials to statistically validate its efficacy, sensitivity, and specificity across diverse, real-world patient populations, thereby paving the way for formal regulatory clearance and broad clinical implementation.

References

[1] V. Gulshan, L. Peng, M. Coram, M. C. Stumpe, D. Wu, A. Narayanaswamy, ... & D. R. Webster, "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *JAMA*, vol. 316, no. 22, pp. 2402–2410, 2016. https://doi.org/10.1001/jama.2016.17216

[2] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning (ICML)*, pp. 6105–6114, 2019. https://proceedings.mlr.press/v97/tan19a.html

[3] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, ... & N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021. https://arxiv.org/abs/2010.11929

[4] W. Cao, V. Mirjalili, and S. Raschka, "Rank consistent ordinal regression for neural networks with application to age estimation," *Pattern Recognition Letters*, vol. 140, pp. 325–331, 2020. https://doi.org/10.1016/j.patrec.2020.11.008

[5] D. S. W. Ting, C. Y. L. Cheung, G. Lim, G. S. W. Tan, N. D.
Quang, A. Gan, ... & T. Y. Wong, "Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes," *JAMA*, vol. 318, no. 22, pp. 2211–2223, 2017.
https://doi.org/10.1001/jama.2017.18152

[6] R. Gargeya and T. Leng, "Automated identification of diabetic retinopathy using deep learning," *Ophthalmology*, vol. 124, no. 7, pp. 962–969, 2017. https://doi.org/10.1016/j.ophtha.2017.02.008

[7] C. P. Wilkinson, F. L. Ferris, R. E. Klein, P. P. Lee, C. D. Agardh,
M. Davis, ... & D. L. Bhatt, "Proposed international clinical diabetic retinopathy and diabetic macular edema disease severity scales," *Ophthalmology*, vol. 110, no. 9, pp. 1677–1682, 2003. https://doi.org/10.1016/S0161-6420(03)00475-5A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, ... & I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017. https://arxiv.org/abs/1706.03762

[8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. https://doi.org/10.1109/CVPR.2016.90

[9] APTOS 2019 Blindness Detection, *Kaggle Competition Dataset*, 2019. [Online]. Available: https://www.kaggle.com/c/aptos2019- blindness-detection

[10] EyePACS Dataset, *Diabetic Retinopathy Detection*, Kaggle, 2015. [Online]. Available: https://www.kaggle.com/c/diabetic- retinopathy-detection

[11] International Diabetes Federation, *IDF Diabetes Atlas*, 10th ed. Brussels, Belgium: International Diabetes Federation, 2021. [Online].
Available: https://www.diabetesatlas.org

[12] Y. Dai, Y. Gao, and F. Liu, "TransMed: Transformers advance multi-modal medical image classification,"

Diagnostics, vol. 11, no. 8, p. 1384, 2021. https://doi.org/10.3390/diagnostics11081384

[13] Y. Zhou, X. He, L. Huang, L. Liu, F. Zhu, S. Cui, and L. Shao, "Collaborative learning of semi-supervised segmentation and classification for medical images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*,
pp. 2079–2088, 2019. https://doi.org/10.1109/CVPR.2019.00218

[14] S. Ramírez, *FastAPI Documentation*, 2024. [Online]. Available: https://fastapi.tiangolo.com/

[15] Meta AI Research, *PyTorch Documentation*, 2024. [Online].
Available: https://pytorch.org/docs/stable/index.html

[16] R. Wightman, *PyTorch Image Models (timm)*, GitHub, 2019. [Online]. Available: https://github.com/huggingface/pytorch-image- models

[17] Meta Open Source, *React Documentation*, 2024. [Online].
Available: https://react.dev/

[18] *Vite — Next Generation Frontend Tooling*, 2024. [Online].
Available: https://vitejs.dev/

[19] *Chart.js — Simple yet flexible JavaScript charting*, 2024. [Online]. Available: https://www.chartjs.org/

[20] *jsPDF — A library to generate PDFs in JavaScript*, 2024. [Online]. Available: https://github.com/parallax/jsPDF

[21] World Health Organization, *Blindness and Vision Impairment*, WHO Fact Sheet, 2023. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/blindness-and- visual-impairment