

Optimizing Clinical Decision-Making through the Integration of MobileNetV2 and Explainable AI in a Diabetic Retinopathy Screening System

Muhammad Faiz Abror
Faculty of Economic and Business
Universitas Negeri Surabaya
Surabaya, Indonesia

muhammadfaiz.23347@mhs.unesa.ac.id

Nanang Hoesen Hidroes Abbrori
Faculty of Economic and Business
Universitas Negeri Surabaya
Surabaya, Indonesia

nangryo@gmail.com

Abstract— Diabetic Retinopathy is a severe diabetic complication and a leading cause of global blindness. While recent studies predominantly rely on computationally heavy deep learning models that evaluate images in isolation without considering patient clinical histories, their implementation in primary healthcare is frequently hindered by massive hardware demands and the inherent opacity of the algorithms. Addressing this research gap, this study aims to develop an efficient and interpretable hybrid Clinical Decision Support System (CDSS) for Diabetic Retinopathy detection. The proposed methodology integrates a lightweight Convolutional Neural Network architecture, MobileNetV2, to extract pathological features from retinal fundus images. To mitigate the limitations of purely image-centric models, the network's probability outputs are dynamically calibrated utilizing a Rule-Based algorithm that processes real-world patient medical parameters, specifically age and the duration of the diabetic history. Furthermore, the system incorporates Grad-CAM-based Explainable Artificial Intelligence (XAI) to generate real-time spatial mappings of retinal damage. The evaluation results demonstrate robust diagnostic performance, achieving an Accuracy of 95.60%, a Precision of 95.64%, a highly optimal Sensitivity (Recall) of 97.40%, an F1-Score of 95.60%, and an Area Under Curve (AUC) score of 98.82%. The hybrid mechanism successfully improves diagnostic decisions by contextually adjusting risk weights, while XAI spatial heatmap projections accurately localize retinal lesions to offer valid anatomical justification for clinical practitioners. In conclusion, the comprehensive integration of lightweight computational efficiency, clinical parametric logic, and analytical transparency produces a robust, hardware-friendly screening platform highly applicable for accelerating patient triage in primary healthcare facilities.

Keywords—Clinical Decision Support System, Explainable AI, Grad-CAM, MobileNetV2, Retinopathy Diabetic

I. INTRODUCTION

Diabetic Retinopathy (DR) represents one of the most vision-threatening microvascular complications of diabetes mellitus and stands as the leading cause of irreversible blindness among the global working-age population. According to global epidemiological reports [1], the prevalence of diabetes continues to surge massively, with critical statistics indicating that nearly half of all affected individuals remain undiagnosed. This condition frequently leads to delayed clinical interventions, as early-stage DR symptoms often go unnoticed by patients until the disease progresses to an advanced stage. To mitigate the risk of widespread blindness, the demand for broad-scale, early screening frameworks has become increasingly urgent. Over the past

decade, the implementation of Deep Learning (DL) technologies, particularly Convolutional Neural Networks (CNNs), has been established as the new gold standard in medical artificial intelligence. This paradigm shift is driven by their supreme capability to automatically extract and classify high-dimensional visual features [2], [3].

A review of the current literature reveals intensive exploration by previous researchers aiming to optimize automated classification performance by designing increasingly deep and complex CNN models [4], [5]. The application of transfer learning across various popular architectures has proven capable of delivering exceptionally high diagnostic accuracy when evaluated on standard retinal image databases such as Messidor [6], [7], [8]. Applied experiments at both domestic and international levels further confirm the comparative reliability of convolutional models in detecting retinal damage [9], [10], including the development of transformer-based hybrid variations [11]. Nevertheless, these massive models consume an enormous number of parameters, thereby requiring high-end and costly computational hardware specifications [12]. Consequently, these conventional models lack adaptability and pose significant practical challenges for implementation in healthcare facilities with limited technological infrastructure [13]. These barriers have prompted a shift in research focus toward the utilization of lightweight architectures, which are specifically tailored for computational efficiency in applied edge devices [14], [15].

Beyond the constraints of hardware efficiency, a systemic flaw in the majority of current automated detection systems lies in their purely image-centric approach. Most models evaluate retinal images in isolation, failing to integrate contextual variables derived from real-world patient medical records [16], [17], [18]. Based on fundamental statistical learning theories, there is an absolute dilemma—or trade-off—between a model's complex flexibility and its analytical interpretability [19]. The highly non-linear nature of operational DL models causes them to act as "black-boxes," rendering them incapable of elucidating the logical rationale behind their diagnostic decisions. This absence of transparency triggers ethical resistance, decision biases, and a crisis of trust among medical professionals, ultimately hindering the adoption of artificial intelligence into actual clinical workflows [20]. Therefore, the integration of Explainable Artificial Intelligence (XAI) instruments has become a

crucial prerequisite to provide objective visual explanations for clinical practitioners [21], [22].

While recent studies have achieved high diagnostic accuracy using complex architectures such as ResNet or Vision Transformers, they present significant research gaps: they demand massive computational resources unsuited for primary healthcare facilities, and they predominantly operate as purely image-centric black-boxes, ignoring crucial clinical contextual variables. To address this complex gap between computational efficiency, medical record context, and the demand for clinical transparency, this study aims to design a hybrid, web-based Clinical Decision Support System (CDSS) for Diabetic Retinopathy screening. The proposed framework adopts a lightweight CNN architecture, MobileNetV2, which is dynamically optimized using the Adam algorithm [23] and rigorously evaluated using Receiver Operating Characteristic (ROC) and Area Under the Curve (AUC) metrics [24]. The primary novelty of this research lies in its integrated hybridization scheme: the visual classification output from MobileNetV2 is complemented by an XAI module based on Gradient-weighted Class Activation Mapping (Grad-CAM) to project visual mappings of retinal pathological lesions. This output is subsequently calibrated using a Rule-Based system that processes patients' clinical parameters, specifically age and the duration of their diabetes history. The scientific contribution of this study is the presentation of a computationally friendly diagnostic assistant platform designed to facilitate rapid patient triage while offering high accountability that can be verified and trusted by frontline medical personnel.

II. METHODS

A. Research Workflow

This study proposes a hybrid Clinical Decision Support System (CDSS) framework designed to automate and transparently validate the diagnosis of Diabetic Retinopathy. Structurally, the research pipeline begins with the acquisition and sorting of medical image datasets into binary classes, followed by image preprocessing and augmentation stages. In the core computational phase, visual images are analyzed using the MobileNetV2 transfer learning architecture to extract pathological features. The output probabilities from this neural network model are subsequently integrated in parallel with the patient's clinical parametric data via a Rule-Based algorithm. The outputs from both processes are combined to generate the final diagnostic decision. As a final step to ensure transparency, the system generates an Explainable AI visualization that maps the areas of retinal damage. Fig. 1 visualizes the entire operational sequence of the proposed hybrid diagnostic system.




Fig. 1. Flowchart of the Hybrid CDSS Research Workflow.

B. Dataset Collection

The primary data source for visual computation was acquired from the APTOS 2019 Blindness Detection public database, accessed directly via the Kaggle application programming interface. The dataset comprises a total of 3,662 fundus images. To optimize early detection sensitivity in the proposed architecture, these image data were sorted and converted into a binary classification system: specifically, the Normal class for healthy eyes and the Abnormal class for eyes indicating Diabetic Retinopathy lesions. The integrated binary dataset was then partitioned using a data splitting function with a ratio of 80% (2,929 images) for the training set to update the model weights, and 20% (733 images) for the validation set to monitor the model's generalization performance on unseen data.

C. Image Preprocessing and Augmentation

The preprocessing stage aims to standardize the image matrix format and enrich data variation before execution by the artificial neural network. All raw images from the dataset were uniformly resized to a resolution of 224×224 pixels to meet the standard input dimensions of the MobileNetV2 architecture. Subsequently, normalization techniques were applied to scale the color pixel intensities from a range of 0 to 255 into a decimal scale of 0.00 to 1.00, which is crucial for stabilizing gradient convergence. To prevent the model from excessive data memorization (overfitting), dynamic augmentation techniques were implemented directly through the image data generator function. This augmentation includes random spatial rotation of up to 20 degrees and horizontal flipping of the images. The extraction process was configured to stream

data into the computational model with a batch size of 32 images per iteration.




Fig. 2. Comparison of Raw and Normalized/Augmented Images.

D. Hybrid Architecture (Integration of MobileNetV2 and Rule-Based Logic)

The primary diagnostic engine was built using the pre-trained MobileNetV2 convolutional architecture, with initial weights derived from training on the massive ImageNet dataset. To accelerate the computational process, the foundational convolutional layers of this model were frozen (trainable = False) to prevent weight updates, allowing the model to focus exclusively on extracting high-level ocular anatomical features. At the architecture's terminal end, a Global Average Pooling 2D layer was added, concluding with a Dense layer utilizing a sigmoid activation function to produce binary probability outputs. This visual computational model was compiled using the Adam optimizer with a learning rate of 0.001 and a binary crossentropy loss function. The single probability value generated by the convolutional model was then fed into the Rule-Based algorithm. If the visual probability approaches a specific threshold, the algorithm evaluates medical record parameters (such as a lengthy history of diabetes) to increase the risk weight or trigger additional clinical intervention alerts. Fig. 3 illustrates the fusion mechanism between the visual network architecture and these medical record parameters.




Fig. 3. Diagram of the Hybrid MobileNetV2 and Rule-Based Logic Architecture.

E. Model Evaluation Metrics and Training

The classification performance of the computational algorithm was continuously measured during each training cycle (epoch) using model performance monitoring instruments (Callbacks). The primary statistical parameters evaluated included Accuracy, Loss, and Sensitivity (Recall). The selection of the Recall metric as

primary indicator holds a strong medical justification; in disease screening systems, the failure to detect a sick patient (False Negative) poses a significantly higher risk than incorrectly predicting a healthy patient (False Positive). Therefore, the best model saving mechanism (ModelCheckpoint) was instructed to exclusively monitor improvements in the Recall score on the validation data (val_recall). Furthermore, an early stopping function (EarlyStopping) was applied to automatically halt the training process if there was no decrease in validation loss for three consecutive epochs, ensuring the model maintained the most efficient convergence weights.

F. Implementation of Explainable AI (Grad-CAM)

To dismantle the black-box nature of artificial neural network algorithms, this system was equipped with post-computational analytical functionality using an Explainable AI (XAI) approach. Gradient-weighted Class Activation Mapping (Grad-CAM) was extracted from the deepest convolutional block to isolate the spatial areas on the retina that contributed most significantly to the diagnostic classification. The system's output interface rendered these analysis results in parallel across three comparative panels. The first panel displayed the patient's original ocular anatomical image. The second panel presented the AI focus map (heatmap), where red color intensity indicated the critical lesion areas detected by the network. The third panel combined the two preceding images into a single overlaid visualization (overlay), thereby providing immediate anatomical justification for ophthalmologists regarding the basis of the model's decision-making. The arrangement of these three diagnostic visual output panels is comprehensively presented in Fig. 4.




Fig. 4. Output of Three XAI Panels: Original Eye Photo, AI Focus Map (Heatmap), and Combined Detection Visualization.

The selection of the Adam optimizer with a learning rate of 0.001 was strategically chosen due to its robust efficiency in handling sparse gradients and its ability to dynamically adapt the learning rate during convergence. Furthermore, an early stopping function (EarlyStopping) was applied to automatically halt the training process if there was no decrease in validation loss for three consecutive epochs, ensuring the model maintained the most efficient convergence weights and strictly preventing overfitting.

III. RESULTS AND DISCUSSION

A. Evaluation of MobileNetV2 Visual Computational Performance

The training process for the MobileNetV2 convolutional neural network model was executed over 10 epochs utilizing the Adam optimizer. Based on the training history logs, the model demonstrated a highly stable convergence rate from the initial epochs. Evaluation on the

validation dataset confirmed that this lightweight architecture adapted exceptionally well to the APTOS 2019 dataset without exhibiting any significant indications of overfitting. This stability is a direct consequence of the base convolutional layer freezing strategy combined with the application of dynamic image augmentation techniques. The latter compels the model to recognize pathological features across various orientations, rather than merely memorizing dominant image pixels.

An in-depth analysis of the performance metrics indicates that selecting Recall (Sensitivity) as the primary indicator in the model's Callback was a highly appropriate strategic decision. In the context of Diabetic Retinopathy screening, the highest risk lies in the system's failure to identify patients who genuinely suffer from retinal abnormalities (False Negative). The visual computational model successfully maintained an optimal validation Recall score concurrently with a decreasing Loss graph. This proves that the MobileNetV2 architecture does not simply guess the majority class but actively constructs feature representations that are highly sensitive to microvascular anomalies in the retina. The training and validation performance dynamics of this system are comprehensively visualized in Fig. 5.




Fig. 5. Training Curve Graphs: Accuracy, Loss, and Recall.

The visual computational model successfully maintained an optimal validation Recall score concurrently with a decreasing Loss graph. To provide a comprehensive assessment of the model's classification capabilities, the detailed evaluation metrics derived from the validation dataset are systematically presented in Table I.

TABLE I PERFORMANCE EVALUATION METRICS OF MOBILENETV2

Evaluation Metric	Score (%)
Accuracy	95.60%
Precision	95.64%
Recall (Sensitivity)	97.40%
F1-Score	95.60%
AUC (Area Under Curve)	98.82%

The evaluation metrics demonstrate the robust diagnostic performance of the proposed lightweight architecture. An Accuracy of 95.60% indicates the model's overall correctness in classifying retinal images. The Precision score of 95.64% reflects high reliability when predicting abnormal cases, which is crucial for reducing false alarms in primary care settings.

Most importantly for a medical screening context, the system achieved a Recall (Sensitivity) of 97.40%. This exceptionally high sensitivity ensures that the model successfully identifies nearly all actual positive cases,

strictly minimizing the risk of misdiagnosing patients with Diabetic Retinopathy as healthy (False Negatives). The F1-Score of 95.60% confirms a well-balanced diagnostic capability, while the AUC (Area Under Curve) score of 98.82% highlights the model's outstanding discriminative power in distinguishing between healthy and pathological eyes.

B. Analysis of the Hybrid CDSS Mechanism

The fundamental advantage of the proposed Clinical Decision Support System (CDSS) lies in its hybrid calibration mechanism, which bridges artificial neural networks with deterministic logic. The majority of current diagnostic systems operate with an image-centric bias, wherein final decisions are drawn purely based on visual captures [Saeidian et al., 2024]. This singular approach frequently yields questionable diagnoses when the uploaded fundus images suffer from poor lighting quality or low resolution.

Through the integration of a Rule-Based Engine, this system successfully overcomes these weaknesses by facilitating contextual evaluation. The mechanism operates through a set of predefined logical thresholds (IF-THEN conditions). For instance, if the visual model MobileNetV2 yields a borderline probability output (e.g., a score of 0.45) due to a lack of captured visual artifacts, the system will not recklessly classify it as a healthy eye. A secondary algorithm proceeds to evaluate the patient's medical record parameters: IF the system detects a prolonged history of diabetes (e.g., > 10 years) OR the patient falls within a highly vulnerable age demographic, THEN the rule-based engine dynamically elevates the final diagnostic risk weight to trigger an urgent clinical intervention alert. This mechanism effectively simulates the cognitive workflow of an ophthalmologist, who consistently correlates visual findings with the patient's pathological history.

C. Diagnostic Transparency via Explainable AI (Grad-CAM)

The ultimate evaluation of a CDSS's viability in the medical sector is measured by its level of transparency. The "black-box" nature of conventional Deep Learning algorithms often creates a chasm of distrust between machines and healthcare professionals. The implementation of Grad-CAM in this study successfully dismantles this barrier.

During real-time interactive testing, the system demonstrated its capability to separate classes with precise visual justification. For example, in a simulated medical record evaluation for a patient named Abror, the system calculated a visually derived risk score at an exceptionally low value of 0.3%, concluding with a "Normal" diagnosis. Analysis of the Grad-CAM heatmap for this case revealed an absence of gradient activation concentration across the retinal area, confirming that the machine genuinely did not detect any exudates or microaneurysms. Conversely, in images with an "Abnormal" diagnosis, the Grad-CAM panel aggressively highlighted pathological lesions with intense red hues; when overlaid onto the original image,

these highlights accurately coincided with retinal hemorrhage locations. This three-panel output not only proves that the model learns the correct features according to ophthalmological principles but also provides a crucial visual forensic foundation for physicians to validate final decisions before taking medical action (see Fig. 4).

D. Clinical Implications

The synthesis of all findings in this study asserts that employing excessively heavy models (such as ResNet or Vision Transformers) is not always a necessity in designing precision medical tools. By marrying the lightweight MobileNetV2 model, real-world parameters, and visual transparency, this research yields highly efficient computational specifications. This efficiency opens substantial opportunities for this web-CDSS platform to be implemented across various primary healthcare facilities or community health centers in resource-limited regions. Ultimately, it accelerates the triage process, reduces the costs of mass screening, and significantly mitigates the threat of permanent blindness in diabetic patients.

Compared to previous studies that predominantly relied on massive and computationally heavy deep learning architectures such as ResNet or Vision Transformers, the proposed lightweight MobileNetV2 model demonstrates highly competitive diagnostic sensitivity while demanding significantly lower hardware resources. Furthermore, while conventional contemporary models operate as isolated, purely image-centric black-boxes, this study's hybridization with a Rule-Based engine and XAI (Grad-CAM) provides a far more holistic diagnostic approach. The Grad-CAM spatial mappings effectively offer valid anatomical justification for clinical practitioners, directly addressing the transparency and accountability gaps frequently cited as major barriers to AI implementation in earlier research.

IV. CONCLUSIONS

This study successfully developed and validated an efficient, hybrid Clinical Decision Support System (CDSS) for Diabetic Retinopathy screening. By integrating the lightweight MobileNetV2 architecture with a Rule-Based clinical algorithm and Explainable AI (Grad-CAM), the proposed framework effectively resolves the primary limitations of contemporary medical AI: massive hardware demands, image-centric biases, and algorithmic opacity. The computational model achieved a highly stable convergence and an optimal Sensitivity (Recall) rate, while dynamically calibrating diagnostic risks utilizing real-world patient parameters (age and diabetes duration). Furthermore, the integration of XAI successfully provided real-time, transparent anatomical justifications through spatial heatmaps.

For future research and development, it is highly recommended to test this hybrid framework using real-world clinical datasets obtained directly from Hospital Information Systems (HIS). Additionally, evaluating the model against a more diverse, multi-ethnic retinal database will further validate its generalization capabilities across various primary healthcare facilities.

ACKNOWLEDGMENT

The authors express their sincere appreciation to the Digital Business Study Program at Universitas Negeri Surabaya for the institutional support and academic facilities provided during the course of this research. Special recognition is directed to the Asia Pacific Tele-Ophthalmology Society (APTOS) 2019 and Kaggle for providing the public retinal fundus image database, which served as the primary material for training the artificial intelligence models in this study. Furthermore, Abror expresses his deepest gratitude to dr. Sekar Ayu Sitoresmi, Sp.M., M.Ked.Klin., for her valuable clinical insights that significantly refined the writing and analytical approach of this manuscript and to Marta Maura Olivia Ilyas for the technical assistance in designing the architectural diagrams presented in this study.

REFERENCES

- [1] S. Webber, *International Diabetes Federation*, vol. 102, no. 2, 2021. doi: 10.1016/j.diabres.2013.10.013.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," pp. 60-1-60-15, 2012, doi: 10.1201/9781420010749.ch60.
- [3] V. Gulshan *et al.*, "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *JAMA - Journal of the American Medical Association*, vol. 316, no. 22, pp. 2402-2410, 2016, doi: 10.1001/jama.2016.17216.
- [4] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pp. 1-14, 2015.
- [5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *ChemInform*, vol. 37, no. 50, pp. 1-9, 2006, doi: 10.1002/chin.200650130.
- [6] E. Decencière *et al.*, "Feedback on a publicly distributed image database: The Messidor database," *Image Analysis and Stereology*, vol. 33, no. 3, pp. 231-234, 2014, doi: 10.5566/ias.1155.
- [7] I. Kandel and M. Castelli, "Transfer learning with convolutional neural networks for diabetic retinopathy image classification. A review," *Applied Sciences (Switzerland)*, vol. 10, no. 6, 2020, doi: 10.3390/app10062021.
- [8] H. Pratt, F. Coenen, D. M. Broadbent, S. P. Harding, and Y. Zheng, "Convolutional neural networks for diabetic retinopathy," *Ingenius*, vol. 00, no. 33, pp. 91-101, 2025, doi: 10.17163/ings.n33.2025.08.
- [9] D. A. Suryaningrum, A. Rasyid, and A. Wiratama, "Klasifikasi retinopati diabetik pada citra retina menggunakan resnet50 dengan perbandingan vgg16 dan mobilenetv2," vol. 11, no. 2, pp. 1627-1637, 2026.
- [10] K. Waruwu and S. Arlis, "IMPLEMENTASI MODEL CONVOLUTIONAL NEURAL NETWORK BERBASIS ARSITEKTUR RESNET50 UNTUK KLASIFIKASI PADA," vol. 3, no. 1, pp. 1-7, 2026.
- [11] J. Zhou, C. Hu, J. Chen, J. Liu, and X. Li, "Artificial Intelligence in Screening and Grading Diabetic Eye Diseases: A Systematic Review From Algorithms to Clinic," *Diabetes Obes. Metab.*, pp. 4286-4302, 2026, doi: 10.1111/dom.70621.
- [12] G. Rajarajeshwari and G. C. Selvi, "Application of artificial intelligence for Classification, Segmentation, Early Detection, Early Diagnosis, and Grading of Diabetic Retinopathy From Fundus Retinal Images : A Comprehensive Review," 1988. doi: 10.1145/43857.43869.
- [13] A. M. Fahmy, "Assessing the Efficacy of AI Models in Medical Imaging Analysis for the Early Identification of Diabetic Foot Ulcers and Gangrene: A Review," *Premier Journal of Science*, vol. 12, pp. 1-14, 2025, doi: 10.70389/PJS.100093.

- [14] A. G. Howard *et al.*, “MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications,” 2017, [Online]. Available: <http://arxiv.org/abs/1704.04861>
- [15] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, “MobileNetV2: Inverted Residuals and Linear Bottlenecks,” *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018, doi: 10.1109/CVPR.2018.00474.
- [16] V. Lakshminarayanan, H. Kheradfallah, A. Sarkar, and J. J. Balaji, “Automated detection and diagnosis of diabetic retinopathy: A comprehensive survey,” *J. Imaging*, vol. 7, no. 9, 2021, doi: 10.3390/jimaging7090165.
- [17] J. Saeidian *et al.*, “Segmentation of choroidal area in optical coherence tomography images using a transfer learning-based conventional neural network: a focus on diabetic retinopathy and a literature review,” *BMC Med. Imaging*, vol. 24, no. 1, 2024, doi: 10.1186/s12880-024-01459-2.
- [18] E. Maliagkani *et al.*, “Artificial Intelligence Algorithms for Epiretinal Membrane Detection, Segmentation and Postoperative BCVA Prediction: A Systematic Review and Meta-Analysis,” *Applied Sciences (Switzerland)*, vol. 15, no. 22, pp. 1–29, 2025, doi: 10.3390/app152212280.
- [19] G. James, T. Hastie, R. Tibshirani, and D. Witten, “An Introduction to Statistical Learning,” *Springer Texts*, vol. 102, p. 618, 2023.
- [20] M. Rezky, K. R. Syarif, and Y. M., “PENERAPAN KECERDASAN BUATAN (AI) DALAM PELAYANAN KEPERAWATAN DAN KESEHATAN: TREN, TANTANGAN, DAN IMPLIKASI ETIS,” *Jurnal Media Keperawatan: Politeknik Kesehatan Makassar*, vol. 12, no. 1, pp. 149–160, 2021.
- [21] E. Tjoa and C. Guan, “A Survey on Explainable Artificial Intelligence (XAI): Toward Medical XAI,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 11, pp. 4793–4813, 2021, doi: 10.1109/TNNLS.2020.3027314.
- [22] R. Hamza Ali and A. Katham Mtasher, “Personalized Medicine and Explainable AI for Diabetic Retinopathy and Tongue Segmentation: A Systematic Review of Current Approaches and Future Directions,” *Global Scientific Journal of Computer Science*, vol. 8, no. 2, pp. 2023–2029, 2023, doi: 10.5281/gsjcs.2023.8174203.
- [23] D. P. Kingma and J. L. Ba, “Adam: A method for stochastic optimization,” *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pp. 1–15, 2015.
- [24] T. Fawcett, “An introduction to ROC analysis,” *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, 2006, doi: 10.1016/j.patrec.2005.10.010.