IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511125

Retina AI: An Explainable Deep Learning-Based Clinical Decision Support System for Diabetic Retinopathy Assessment

Trupti Gangakhedkar[1], Bhavana M Patel[2], Abhijna Hiremath[3], Manasa L[4], Gowrika H S[5]

Assistant Professor, East West Institute of Technology, Bengaluru-91[1]

UG Scholar, East West Institute of Technology, Bengaluru-91[2-5]

Abstract: Diabetic Retinopathy (DR) is a diabetes-related retinal disorder that can progress through different stages and may lead to vision impairment when not appropriately assessed and managed. Retinal fundus photography provides an important source of information for evaluating retinal abnormalities, but manual assessment can require substantial specialist time. RetinaAI is a doctor-oriented web-based clinical decision support system developed to assist ophthalmologists in assessing diabetic retinopathy from retinal fundus images. The proposed system accepts patient information and a retinal image, performs image preprocessing, and applies a trained Convolutional Neural Network (CNN) to classify the image into five classes: No DR, Mild, Moderate, Severe, and Proliferative DR. The system displays the primary assessment, disease stage, severity, confidence information, and clinical-support text. To improve interpretability, Grad-CAM is used to generate a visual explanation of the image regions that contributed to the classification. A structured PDF report is also generated for documentation. The current implementation integrates the web interface, image processing, model inference, explainability and reporting into one workflow. RetinaAI is designed as a decision-support tool for professional review; its output is not intended to replace ophthalmological examination or clinical judgment.

Index Terms: Diabetic Retinopathy, RetinaAI, Fundus Image, Deep Learning, Convolutional Neural Network, CNN, Grad-CAM, Explainable AI, Clinical Decision Support, Medical Image Analysis

I. INTRODUCTION

Diabetic Retinopathy is a major retinal complication associated with diabetes. The condition is characterized by progressive retinal changes that can be difficult to identify at an early stage without appropriate retinal examination. Fundus photography provides a non-invasive visual representation of the retina and is therefore useful for screening and assessment workflows. As the number of retinal images increases in clinical environments, computer-assisted image analysis can provide useful preliminary information to support specialist review. Traditional retinal assessment depends strongly on specialist expertise and careful examination of retinal structures. Although expert assessment. remains essential, a software system can help organize the information presented to the clinician, provide a consistent image-processing workflow, and reduce repetitive tasks.

A useful clinical-support application should therefore do more than display a classification label. It should provide the predicted stage, confidence information, visual evidence, and a structured record that can be reviewed by the ophthalmologist. The primary assessment. The application additionally generates a Grad-CAM visualization and a downloadable PDF report.

II. MOTIVATION

The motivation for Retinal is the need for a unified assessment workflow rather than a stand-alone image classifier. In a practical setting, a doctor may need to enter patient information, inspect a fundus image, interpret an assessment, review confidence information, examine visual evidence, and document the result. Separating these tasks across different tools can make the workflow less convenient.

The project therefore combines these functions into a single browser-based application. The design is intentionally doctor-facing: the system presents information for clinical review rather than presenting itself as an autonomous patient diagnosis system. The architecture also separates the user interface, backend request handling, image preprocessing, model inference, explainability and report generation so that individual components can be improved independently.

III. PROBLEM STATEMENT

The project addresses the following problem: how to develop a doctor-oriented application that can accept a retinal fundus

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511125

image and patient information, perform automated five-stage diabetic-retinopathy assessment, provide confidence information and visual explanation, and generate a structured report within one workflow
The proposed system aims to use deep learning models such as VGG16 and InceptionV3 to classify images into five stages—No DR, Mild, Moderate, Severe, and Proliferative DR—while using Grad-CAM to provide visual explanations of the prediction. The proposed solution is intended to support early screening and assist healthcare professionals, rather than replace clinical diagnosis.

IV.OBJECTIVES

- Develop a doctor-oriented web application for retinal fundus-image assessment.
- Classify retinal images into No DR, Mild, Moderate, Severe, and Proliferative DR.
- Standardize image input through resizing to 224 × 224 pixels and pixel normalization.
- Display the primary assessment, disease stage, severity, and confidence information.
- Provide confidence scores for all five supported classes.
- Generate Grad-CAM visualizations to support interpretation of the model output.
- Provide structured clinical-support information associated with the predicted stage.
- Generate a downloadable PDF report containing patient and assessment information.
- Maintain a workflow in which the ophthalmologist remains responsible for final clinical interpretation.

V.RELATED WORK

Deep learning has become an important approach for medical image analysis because neural networks can learn hierarchical visual representations from image data. CNN-based systems are particularly suited to image classification because convolutional layers learn local spatial patterns and deeper layers combine these patterns into higher-level representations.

Gulshan et al. demonstrated the use of deep learning for detection of diabetic retinopathy from retinal fundus photographs, showing the potential of automated retinal image analysis. Ting et al. further investigated deep learning for diabetic retinopathy and related retinal diseases using large-scale retinal datasets. These studies provide the broader technical motivation for using learned image representations for retinal assessment.
A limitation of a classification-only interface is interpretability. A predicted class does not directly show which regions of an image influenced the result. Selvaraju et al. introduced Grad-CAM as a method for generating class-discriminative localization maps from convolutional networks. Retina AI uses this explainability concept to provide a visual reference alongside the classification result.
The project therefore combines three ideas: image classification, visual explanation and structured clinical reporting. The contribution of the prototype is not to replace specialist assessment, but to integrate these functions into a single doctor-oriented workflow.

VI.RESEARCH GAP

Existing approaches may focus on automated classification or on explainability as a separate research component. A practical assessment interface also needs patient context, confidence information, clear presentation and documentation. RetinaAI addresses this workflow-level gap by integrating image upload, patient information, model inference, five-stage output, confidence display, Grad-CAM visualization and PDF reporting in one web application.
The system also distinguishes between a model prediction and a clinical decision. The prediction is treated as supporting information, while the ophthalmologist is responsible for correlating the result with the retinal image and other clinical findings.

VII. PROPOSED SYSTEM

The proposed system consists of six major functional stages: input collection, image preprocessing, model inference, result interpretation, explainability, and reporting.
In the input stage, the ophthalmologist enters patient information such as Patient ID, patient name, age, gender and diabetes duration, and uploads a retinal fundus image. The Flask backend receives the form data and image file.

During preprocessing, the uploaded image is loaded and resized to 224 × 224 pixels. The image is converted into a numerical array, a batch dimension is added, and pixel values are divided by 255.0 so that the values are represented in the range 0–1.

The trained CNN model then performs inference. The implementation supports five classes: NoDR, Mild, Moderate, Severe and Proliferated. The model produces a score for each class. The class with the maximum score is selected as the primary result. The displayed result is converted to user-friendly labels such as No DR and Proliferative DR.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511125

The result layer maps each class to a disease stage, severity description, explanatory text and clinical-support message. The application also records the prediction date and time.
For explainability, Grad-CAM is generated using the model's convolutional feature representation. The heatmap is combined with the original retinal image so that the clinician can inspect the regions.
Finally, the reporting module uses Report Lab to create a PDF containing patient information, prediction, disease stage, severity, confidence, description, clinical recommendation, clinical decision.

VIII. IMAGE PREPROCESSING

Image preprocessing is essential because the trained model expects a fixed input shape. In the current implementation, the uploaded image is resized to 224 × 224 pixels. The image is then converted into an array and expanded to represent a single image batch.

If p denotes an original pixel intensity in the range 0–255, normalization is performed as:

p' = p / 255 Therefore, p' lies in the range 0–1. This scaling provides a consistent numerical representation for model inference. The preprocessing pipeline can be represented as: Uploaded Image Resize(224× 224) → Array Conversion → Batch Dimension → Normalization → Model Input.

IX.CNN MODEL ARCHITECTURE

The trained model used by the current implementation is a sequential image-classification model. The model inspection performed during development showed the following major layers:

1. Conv2D
2. MaxPooling2D
3. Conv2D
4. MaxPooling2D
5. Conv2D
6. MaxPooling2D
7. Flatten
8. Dense
9. Dropout
10. Dense layer

The convolutional layers learn visual patterns from the retinal image. Max-pooling reduces the spatial representation while retaining important features. The Flatten layer converts the learned feature maps into a vector. Dense layers combine these features for classification, while dropout is used as a regularization mechanism during training. The final dense layer produces the class scores used for the five-stage assessment. The model is stored in Kera's format and loaded by the Flask application using TensorFlow/Kera's.

X.FIVE-STAGE CLASSIFICATION

The classification output contains five predefined classes:

a. No DR — no diabetic-retinopathy class.
b. Mild — mild diabetic-retinopathy class.
c. Moderate — moderate diabetic-retinopathy class.
d. Severe — severe diabetic-retinopathy class.
e. Proliferative DR — advanced proliferative diabetic-retinopathy class.

The application maps the internal class names to user-facing labels and associates each class with a stage number from Stage 0 to Stage 4. The stage and severity displayed by the interface are derived from the selected model class.

XI.CONFIDENCE ESTIMATION

For an input image, the model returns a score for each supported class. The current implementation selects the class with the largest score using an argmax operation. If the output vector is p = [p1, p2, ..., p5], the selected class is:

$$\hat{y} = \arg\max_i p_i \quad (1)$$

and the displayed confidence is:

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511125

$C = \max_i p_i \times 100\%$ (2)

The application also displays a confidence table containing the score of each of the five classes. Confidence should be interpreted as model output rather than as a guarantee of clinical correctness. A high model score does not remove the need for professional review, and a low score for a severe class does not by itself prove that severe disease is absent.

XII.GRAD-CAM EXPLAINABILITY

Grad-CAM is included to make the classification result more interpretable. The technique uses gradients associated with a target class and a convolutional feature layer to estimate which spatial regions contributed strongly to that class score. The resulting Should correlate the visualization with the original fundus image. However, the proposed system is intended as a decision-support and preliminary screening tool rather than a replacement for professional diagnosis. Further validation using larger datasets and real-world clinical data is required before deployment in clinical environments.

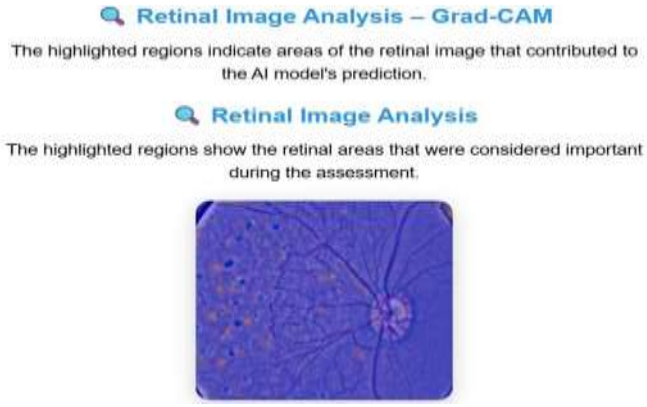


Fig .1. Grad-Cam Image

The design is intentionally focused on professional use. It avoids presenting the output as a final diagnosis and instead frames the result as an assessment that should be reviewed by the ophthalmologist.
Should correlate the visualization with the original fundus image. However, the proposed system is intended as a decision-support and preliminary screening tool rather than a replacement for professional diagnosis. Further validation using larger datasets and real-world clinical data is required before deployment in clinical environments.
The system has the potential to support early screening, particularly in areas where access to specialized eye-care professionals may be limited. This can help users understand where the model is focusing when making its prediction.

XIV. REPORT GENERATION

The report-generation module creates a PDF document using Report Lab. The report includes:

- Patient ID
- Patient name
- Age
- Gender
- Diabetes duration
- Prediction
- Disease stage
- Severity
- Confidence
- Description
- Prediction date and time
- Model information

The purpose of the report is to provide a structured record that can be saved or shared as part of the project demonstration. A production clinical deployment would require appropriate authentication, access control, secure storage, audit trails and privacy safeguards.

The system classifies the retinal image into one of five categories: No Diabetic Retinopathy, Mild, Moderate, Severe, or Proliferative Diabetic Retinopathy. Along with the prediction, an AI-generated explanation is provided to make the result easier to understand.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511125

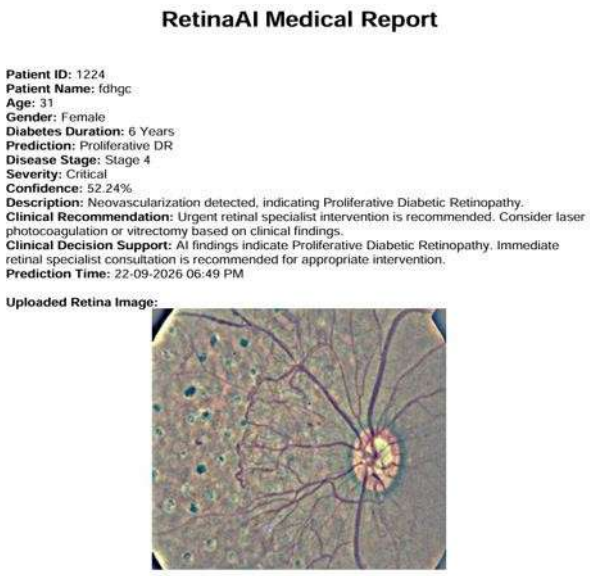


Fig .2. Report Generation

XV. SYSTEM ARCHITECTURE

The architecture follows a sequential processing flow. The ophthalmologist provides patient information and a retinal fundus image Flask application receives the request and invokes the preprocessing pipeline. The preprocessed image is passed to the trained CNN classifier. The resulting class scores are used to produce the primary assessment and confidence information.

The same model is used by the Grad-CAM module to create a visual explanation. The application then presents the clinical-support output and generates a PDF report.

The architecture separates input handling, preprocessing, inference, explainability and reporting. This modular organization makes it easier to test individual components and replacing or improving a component without redesigning the complete interface.

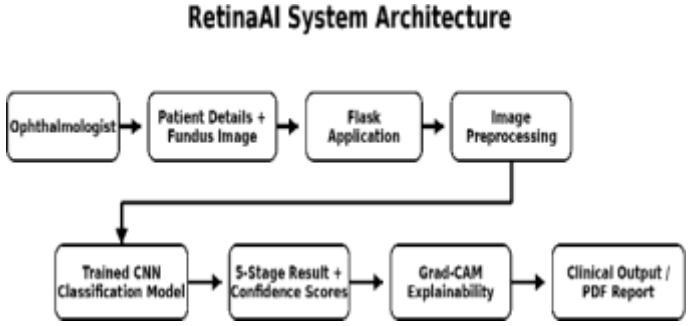


Fig .2. System architecture of the doctor-oriented Retinal workflow

XVI. ALGORITHM

Algorithm: Retinal Assessment

Input: Retinal fundus image I and patient information D.

1. Receive D and I from the ophthalmologist.
2. Validate that an image has been uploaded.
3. Store the uploaded image for processing.
4. Resize I to 224 × 224 pixels.
5. Convert the image to a numerical array.
6. Add the batch dimension.
7. Normalize pixel values using I' = I/255.
8. Pass I' to the trained CNN model.
9. Obtain scores for the five DR classes.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511125

10. Select the class with the maximum score.
11. Map the selected class to its display label, stage and severity.
12. Generate confidence information for all classes.
13. Generate Grad-CAM using the selected convolutional layer.
14. Display the assessment and supporting information.
15. Generate the PDF report.
16. Allow the ophthalmologist to review the output.

Output: Five-stage assessment, confidence information, visual explanation and structured report.

XVII. IMPLEMENTATION DETAILS

The application is implemented as a web-based system using Python and Flask for the backend and HTML/CSS for the user interface. TensorFlow/Kera's is used to load the trained model and perform inference. The current classification model contains convolutional, max-pooling, flatten, dense, and dropout layers. Uploaded images are resized to 224 × 224 pixels and normalized to the range 0–1 before prediction.

The system supports five diabetic-retinopathy classes and displays the predicted class, stage, severity, confidence, clinical-support text, and prediction timestamp. A Grad-CAM utility produces a heatmap using the model's convolutional layer. Report generation is implemented using Report Lab.

XVIII. TESTING AND VALIDATION

The prototype was tested at multiple levels. Model loading was tested independently to confirm that the saved Kera's model could be loaded successfully. Model-layer inspection was used to identify the convolutional layers required for explainability. Image prediction. tested produced. Grad-CAM generation was tested separately and successfully produced a saved visualization.
The Flask interface was tested for images uploaded, patient information submission and display of prediction results. PDF report generation was also tested by creating a report containing the assessment information.
The current project evidence supports functional demonstration of these components. Quantitative model-performances metrics should be calculated from the actual held-out evaluation datasets.

XIX. FUNCTIONAL AND NON-FUNCTIONAL REQUIREMENTS.

Functional requirements:

- FR1. The system shall accept patient details.
- FR2. The system shall accept a retinal fundus image.
- FR3. The system shall preprocess the image to the expected input format.
- FR4. The system shall classify the image into five supported classes.
- FR5. The system shall display the primary assessment and confidence.
- FR6. The system shall display confidence scores for all five classes.
- FR7. The system shall generate a Grad-CAM visualization.
- FR8. The system shall display clinical-support information.
- FR9. The system shall generate a downloadable PDF report.

Non-functional requirements:

- NFR1. Usability — the interface should be understandable to the intended professional user.
- NFR2. Reliability — valid inputs should follow a consistent processing workflow.
- NFR3. Maintainability — backend, interface, model and explainability components should remain modular.
- NFR4. Performance — inference and report generation should complete within a practical response time.

XX. RESULTS AND DISCUSSION

The developed prototype demonstrates an end-to-end workflow from retinal image upload to assessment and report generation. The trained model loads successfully in the application. The prediction module produces a five-class result and confidence values. The interface displays the assessment, disease stage, severity and supporting information. The Grad-CAM module successfully generated a visualization during testing. The report-generation module produced. For scientific reporting, the following metrics should be computed from the project's held-out test set: accuracy, precision, recall, F1-score, sensitivity and specificity.
the current project information supplied for this paper does not include the complete test-set counts; numerical performance values are intentionally not fabricated.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.1511125

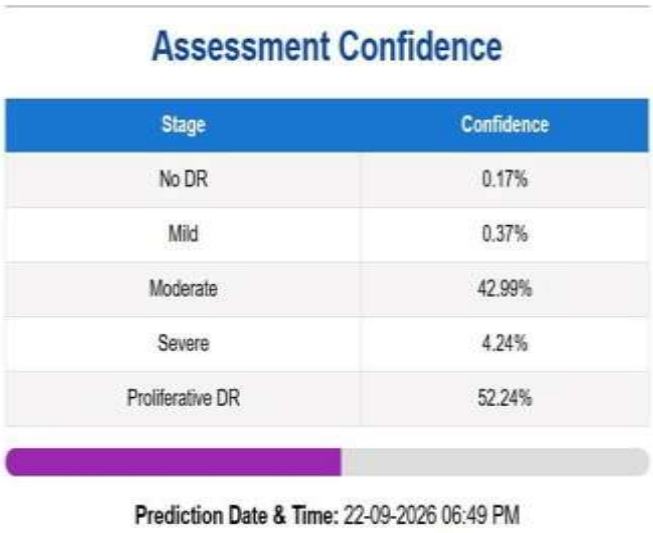


Fig .2. Diabetic Retinopathy Classification based on Confidence Scores

XXI. LIMITATIONS

The current prototype has several limitations. First, a model prediction is not equivalent to a clinical diagnosis and must be interpreted by an ophthalmologist. Second, confidence scores are model probabilities/scores and are not automatically calibrated clinical probabilities.

Third, formal generalization performance depends on the diversity and quality of the training and test datasets. Fourth, the current prototype stores uploaded images and generates reports locally; production deployment would require stronger security, access control and patient-data management.

Fifth, Grad-CAM is an explanatory visualization and should not be treated as a precise lesion detector. Findings.

Future versions can also support multi-disease retinal assessment and integration with hospital information systems after appropriate clinical validation.

XXII.CONCLUSION

Retinal is a doctor-oriented clinical decision support prototype for diabetic retinopathy assessment from retinal fundus images. The system combines a trained CNN classifier with image preprocessing, five-stage assessment, confidence display, Grad-CAM explainability, clinical-support information and PDF reporting. The modular web architecture allows the different components to operate as one workflow while keeping the ophthalmologist responsible for the final clinical decision.

The project demonstrates the feasibility of integrating automated retinal image analysis and explainability into a practical application.

XXIII.FUTURE WORK

Future work can extend Retinal in several directions. A larger and more diverse retinal dataset can be used for improved generalization. Formal validation can be performed using an independent held-out test set and clinically meaningful metrics.

Confidence calibration and uncertainty estimation can be investigated so that the displayed confidence better reflects model reliability.

An automated image-quality check can be added before classification to identify blurred, poorly illuminated or unsuitable images.

REFERENCES

[1]. G. Gulshan et al., “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.

[2]. D. S. W. Ting et al., “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.

[3]. R. R. Selvaraju et al., “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization,” in Proc. IEEE ICCV, 2017, pp. 618–626.

[4]. M. Abadi et al., “TensorFlow: A System for Large-Scale Machine Learning,” in Proc. 12th USENIX OSDI, 2016,

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.1511125

pp. 265–283.
[5]. F. Chollet, Deep Learning with Python, 2nd ed. Manning Publications, 2021.
[6]. W. McKinney, Python for Data Analysis, 3rd ed. O'Reilly Media, 2022.
[7]. Li, F., et al. (2019). Automatic Detection of Diabetic Retinopathy in Retinal Fundus Photographs Based on Deep Particularly relevant because the study uses Inception-v3 for automatic DR detection
[8]. Li, F., et al. (2022). Deep learning-based automated detection for diabetic retinopathy and diabetic macular oedema in retinal fundus photographs. Eye, 36, 1433–1441.
[9]. Oh, K., et al. (2021). Early detection of diabetic retinopathy based on deep learning and ultra-wide-field fundus images. Scientific Reports, 11, 1897.
[10]. Krause/related authors — systematic evidence: Deep learning algorithms for detection of diabetic retinopathy in retinal fundus photographs: A systematic review and meta-analysis. (2020).

 This work is licensed under a Creative Commons Attribution 4.0 International License