PROJECT ABSTRACT

Diabetic Retinopathy Detection using CNN

Diabetic retinopathy is the leading cause of preventable blindness among working-age adults, yet early stages show no symptoms and rural screening access is scarce.

01 Title

Diabetic Retinopathy Detection using CNN

02 Introduction / Background

Diabetic retinopathy is the leading cause of preventable blindness among working-age adults, yet early stages show no symptoms and rural screening access is scarce. Automated grading of retinal fundus photographs with convolutional neural networks has matched specialist-level accuracy in landmark studies, making it one of the most meaningful medical-imaging topics a final-year student can take on.

03 Problem Statement

Manual fundus screening is slow, expensive and unavailable where it is needed most, while generic image classifiers trained on natural images fail on retinal specifics — vessel structure, lesion texture and severe class imbalance. This project builds a complete grading pipeline: fundus-specific preprocessing, transfer learning fine-tuned on the APTOS 2019 dataset, five-grade severity classification and Grad-CAM lesion heatmaps for explainability.

04 Objectives

- Five-grade DR classification (no DR, mild, moderate, severe, proliferative).
- Transfer learning on an EfficientNet/ResNet backbone fine-tuned on APTOS 2019.
- Fundus-specific preprocessing: circular crop, CLAHE enhancement, artifact handling.
- Grad-CAM heatmaps localizing microaneurysms, haemorrhages and exudates.
- Quadratic weighted kappa and per-grade metrics on a held-out test split.
- RetinaScan Streamlit app: screening, batch mode, dashboard.

05 Proposed System / Working

Implemented in Python 3.9+ with TensorFlow/Keras, OpenCV and Streamlit. Fundus images from APTOS 2019 are circular-cropped to remove black borders, resized to the backbone input size and enhanced with CLAHE to sharpen vessels and lesions. An ImageNet-pretrained EfficientNet/ResNet backbone is fine-tuned with its top layers unfrozen, using a class-weighted loss and augmentation

(rotations, flips, zooms, brightness jitter) to fight the dataset's imbalance. Grad-CAM is computed from the final convolutional layer to overlay lesion heatmaps. Evaluation uses quadratic weighted kappa — the standard DR competition metric — plus per-grade precision, recall and a confusion matrix. The RetinaScan Streamlit app serves single-image and batch screening with a grade-distribution dashboard; moderate DR and above triggers an ophthalmologist-referral flag.

The model reaches 78–86% five-class accuracy with quadratic weighted kappa of 0.80–0.88 on the held-out APTOS test split; per-grade analysis shows strongest performance on no-DR and moderate cases and expected weakness on rare proliferative samples. CPU inference takes seconds per image. The live RetinaScan demo grades an uploaded fundus photo, shows the confidence badge and toggles the lesion heatmap — a memorable viva moment.

06 Technologies / Components Used

Python · TensorFlow/Keras · OpenCV · Streamlit

07 Key Features

- Five-grade severity grading with per-grade probability breakdown.
- Transfer learning fine-tuned on the APTOS 2019 dataset.
- Grad-CAM lesion heatmaps for every prediction.
- CLAHE fundus preprocessing and class-imbalance handling.
- RetinaScan app: upload, batch screening, dashboard, referral alerts.
- Held-out test evaluation with kappa and confusion matrix.

08 Applications / Use Cases

- Rural diabetes screening camps
- primary-care triage
- tele-ophthalmology pipelines
- hospital record pre-screening
- research on automated retinal analysis

10 Conclusion

This project delivers a complete, examinable retinal-AI pipeline — from fundus preprocessing to graded, explainable screening output — with honest per-grade metrics and the screening disclaimer a medical-adjacent project must carry.