



Research Paper

GAN-BASED RETINAL IMAGES SYNTHESIS AND EYE DISEASE DIAGNOSIS

V.Hethrika, P. Jithendra Kumar, T. Chaitanya Sri,
S. Harsha Vardhan Raju

Department of CSE (Data Science)
Raghu Engineering College, Dakamarri
Vizianagaram, Andhra Pradesh, India

D.HIMA BINDU Asst.Professor
Department of Computer Science & Engineering
Raghu Engineering College, Dakamarri
Vizianagaram, Andhra Pradesh, India
hethrikavaddadi@gmail.com

Abstract- This paper presents a retinal fundus image analysis framework that combines vessel segmentation, GAN-based minority-class augmentation, disease classification, and explainable inference for automated ocular screening. The study focuses on four classes: normal, diabetic retinopathy, glaucoma, and age-related macular degeneration. ODIR-5K was used for classification, while a retinal vessel segmentation dataset was used to train PUNet-33 and a lightweight CycleGAN-style generator. The generator learned a vessel-mask to fundus-image mapping and was used to synthesize additional glaucoma and AMD images, increasing those classes from 162 and 164 images to 300 each in the training set. The strongest practical classifier in the project was a weighted EfficientNet-B3 ensemble, while the GAN augmentation experiment retrained an Inception-v3 classifier for comparison. Fixed-split testing reached 73.65% accuracy for the weighted ensemble, whereas patient-aware grouped cross-validation provided a more reliable pooled accuracy of 67.84%. The GAN-augmented retraining run reached 70.61% test accuracy, showing that synthetic balancing improved class availability but did not yet outperform the original baseline. The project also integrates GradCAM, SHAP-style attribution, and a layman-readable explanation to support interpretable deployment.

Keywords- GAN augmentation; retinal fundus image; diabetic retinopathy; glaucoma; AMD; vessel segmentation; explainable AI; ocular disease screening

I. INTRODUCTION

Retinal fundus imaging is a non-invasive and widely used technique for screening chronic eye diseases such as diabetic retinopathy, glaucoma, and age-related macular degeneration. Manual review by ophthalmologists remains the clinical standard, but large-scale screening programs require automated support tools that can process images rapidly and consistently [1].

A major difficulty in fundus classification is class imbalance. In real datasets, normal and diabetic retinopathy images are often much more common than glaucoma and AMD samples. This imbalance reduces minority-class learning, especially when disease patterns are subtle or limited to small retinal regions. Synthetic data generation with generative adversarial networks has therefore become a useful research direction [4].

The main contribution of this B.Tech project is a complete GAN-assisted retinal analysis pipeline. It includes dataset preparation, fundus cropping and CLAHE enhancement, PUNet-33 vessel segmentation, mask-to-fundus synthetic image generation, classifier retraining, grouped cross-validation, and explainable inference. The work studies not only whether GAN augmentation can balance the dataset, but also whether it improves practical classification results.

II. LITERATURE SURVEY

Recent retinal disease screening studies combine transfer learning, lesion-sensitive preprocessing, and explainable deep learning. The literature can be grouped into three main directions: direct fundus-image classification, vessel-aware retinal analysis, and data augmentation with generative models.

A. Retinal Disease Classification

ODIR-style fundus screening has been widely explored using CNN backbones such as Inception and EfficientNet. These models learn global retinal structure together with local lesion cues, but performance strongly depends on data quality, class balance, and proper patient-level evaluation [1], [5], [6].

B. Vessel-Aware Retinal Analysis

Retinal vessels carry clinically useful information for glaucoma and diabetic retinopathy analysis. U-Net based segmentation has become the standard approach for vessel extraction because it preserves fine vascular detail while remaining computationally efficient for biomedical images [3], [9].

C. GAN-Based Augmentation

GANs are increasingly used to address minority-class scarcity in medical imaging. Cycle-consistent generators can translate structure from one domain to another without requiring large paired datasets. However, synthetic realism and downstream generalization remain key challenges, especially in retinal imaging where small artifacts can mislead a classifier [4].


TABLE II. Comparison of Retinal Screening and Augmentation Approaches

Category	Core Method	Strength	Limitation
Fundus CNN screening	Transfer learning on ODIR	Good overall baseline	Sensitive to class imbalance
Vessel segmentation	U-Net / PUNet family	Captures fine vascular detail	Not a full disease classifier
Explainable inference	GradCAM and SHAP	Improves model interpretability	Depends on base model quality
GAN augmentation	CycleGAN-style synthesis	Adds minority-class samples	Synthetic realism may vary
Fixed-split ensemble	Weighted EfficientNet-B3	Best practical test result	Less reliable than grouped CV
Proposed study	Segmentation + GAN + classifier	Complete research pipeline	Generalization still limited

III. SYSTEM DESIGN AND ARCHITECTURE

The proposed system follows five connected stages: dataset preparation, vessel segmentation, GAN-based minority synthesis, multi-class disease classification, and explainable inference. Fundus images are first cropped to the retinal field and enhanced with CLAHE. PUNet-33 extracts vascular structure, the generator synthesizes additional glaucoma and AMD training images, and the final classifier predicts one of four disease classes.

GAN-Augmented Retinal Disease Analysis I



ODIR-5K images are cropped and enhanced, vessel masks guide GAN-based augmenta

Fig. 1. Overall workflow of the proposed GAN-assisted retinal disease analysis pipeline.

A. Segmentation-GAN-Classification Flow

Segmentation and augmentation are used as support modules rather than isolated tasks. Vessel masks help the generator learn retinal structure, and the augmented data are then used to improve minority-class availability for classifier retraining. The last stage combines class prediction with visual explanation maps so that the output is suitable for both research analysis and demonstration use.

IV. FEATURE GUIDANCE AND EXPLAINABILITY

Retinal diagnosis depends on both global and local visual patterns. Diabetic retinopathy may appear as microaneurysms or small hemorrhages, glaucoma can show vessel compression around the optic disc, and AMD affects the central macular region. The project therefore uses both vessel-based structure learning and end-to-end image classification.

To improve interpretability, the inference pipeline generates GradCAM heatmaps, SHAP-style attribution maps, class confidence scores, and a layman-readable explanation. This makes the system easier to present in a research paper, demo interface, and clinical discussion setting [7], [8].

V. DATASET AND PREPROCESSING

Two datasets were used in this study. ODIR-5K provided the classification images for the four selected classes: normal, diabetic retinopathy, glaucoma, and AMD [1]. A combined retinal vessel segmentation dataset provided paired fundus images and vessel masks for segmentation training and GAN learning [2].

A. Preprocessing Pipeline

Each fundus image was cropped to the circular retinal field, resized, and enhanced using CLAHE to improve local contrast. The classification pipeline used consistent validation transforms, while segmentation used image-mask aligned preprocessing. These steps reduced black-border bias and made lesion regions easier to learn.

B. Minority Class Balancing

The original training split contained 1471 normal, 963 diabetic retinopathy, 162 glaucoma, and 164 AMD images. To reduce the class gap, the generator was used to produce 138 synthetic glaucoma images and 136 synthetic AMD images, increasing both minority classes to 300 images each.

VI. PROPOSED GAN-AUGMENTED RETINAL FRAMEWORK

The proposed framework joins three learning blocks. PUNet-33 learns vessel segmentation from retinal image-mask pairs, the CycleGAN-style generator maps vessel structure to realistic fundus appearance, and a retinal disease classifier is trained on the augmented classification set. The generator is intentionally lightweight so that the full experiment remains feasible on a single research workstation.

Mask-to-Fundus GAN Sa
Real mask and generated minority-class ima




Fig. 2. Real vessel-guided fundus samples and GAN-generated minority-class images used for augmentation.

A. GAN Generator and Discriminator

A lightweight U-Net style encoder-decoder was used as the generator, and a three-layer PatchGAN was used as the discriminator. The training objective combined adversarial loss, cycle consistency loss, and identity loss. This setup was chosen to keep the research implementation simple while still allowing visually meaningful synthesis.

B. Classifier Training

The strongest practical classifier in the project was a weighted EfficientNet-B3 ensemble under the fixed-split protocol. To test the direct effect of GAN augmentation inside the academic experiment, a fresh Inception-v3 classifier was retrained using the new augmented glaucoma and AMD images and then compared with the earlier fixed-split baseline.

C. Training Protocol

The GAN was trained with batch size 16, learning rate $2e-4$, and early stopping patience 15. The best generator checkpoint was saved after the validation loss reached its minimum. Classification training used the same project preprocessing logic, class-aware loss weighting, and standard versus TTA inference selection on validation data.

VII. IMPLEMENTATION

A. Development Environment

The complete project was implemented in Python using PyTorch, torchvision, PIL, NumPy, and albumentations. Training and inference were executed on an NVIDIA RTX PRO 6000 Blackwell GPU with 102 GB VRAM. Mixed

precision, cudnn benchmark mode, and torch.compile were used where stable.

B. Inference and Visualization

The final inference script accepts a single fundus image, applies preprocessing, predicts the disease class, segments vessels, and produces explanation maps. The output is saved as a panel-style figure and can also be served through a Streamlit interface for interactive demonstration.

C. Deployment Considerations

For deployment, the non-GAN weighted ensemble remained the most reliable model because it outperformed the GAN-retrained classifier on the tested split. Nevertheless, the GAN module is valuable from a research perspective because it demonstrates a full synthetic data workflow and highlights the practical limits of naive augmentation.

VIII. RESULTS AND DISCUSSION

Three result views were considered in this project: a fixed-split ensemble result, patient-aware grouped cross-validation, and the GAN-augmented retraining experiment. The fixed-split ensemble achieved the strongest practical score, but grouped cross-validation provided the most trustworthy estimate of generalization.

A. Grouped Evaluation

Five-fold patient-aware grouped cross-validation was used to reduce leakage between images from the same patient. The pooled result was 67.84% accuracy, 60.58% macro F1, and 59.16% balanced accuracy. The main error pattern remained confusion between normal and diabetic retinopathy, which indicates that fine lesion separation is still the dominant problem.




Fig. 3. Patient-aware five-fold cross-validation accuracy of the retinal classifier.




Fig. 4. Pooled confusion matrix from the patient-aware cross-validation experiment.

B. Quantitative Performance

Table I summarizes the major experiments. The weighted fixed-split ensemble achieved 73.65% test accuracy and 70.70% macro F1. The grouped cross-validation result was lower but more reliable. The GAN experiment successfully balanced glaucoma and AMD counts, but the retrained classifier did not outperform the earlier non-GAN baseline.

TABLE I. Performance Summary of the Proposed Retinal Disease Analysis Pipeline

Metric	Value
Original training images	2760
Training images after GAN	3034
Glaucoma images before / after GAN	162 / 300
AMD images before / after GAN	164 / 300
Weighted ensemble fixed-split accuracy	73.65%
Weighted ensemble fixed-split macro F1	70.70%
Patient-aware CV accuracy (mean +- std)	67.84% +- 1.79%
Patient-aware CV macro F1 (mean +- std)	60.47% +- 1.29%
Patient-aware CV balanced accuracy	59.16% +- 1.36%
GAN-augmented Inception test accuracy	70.61% (-3.04% vs baseline)

C. Effect of GAN Augmentation

GAN training itself was stable: the generator validation loss reduced from 16.13 to 3.36, and 274 synthetic minority-class images were produced successfully. However, the downstream Inception-v3 retraining result was 70.61% test accuracy, which is 3.04 percentage points below the earlier 73.65% reference. This suggests that balancing class counts alone is not enough; synthetic realism and label consistency also matter.

D. Comparative Analysis

The project therefore offers two useful outcomes. First, it establishes a complete GAN-assisted retinal disease pipeline suitable for academic demonstration. Second, it shows through controlled experiments that synthetic augmentation must be validated carefully instead of being assumed beneficial. The fixed-split ensemble remained the best deployable model, while grouped cross-validation remained the best benchmark for honest reporting.




Fig. 5 (left). Class distribution before and after GAN augmentation. Fig. 6 (right). Accuracy, macro-F1, and balanced-accuracy comparison across major experiments.

E. Interpretation

In practical terms, the GAN module solved the shortage of glaucoma and AMD training examples, but it did not fully solve the clinical difficulty of distinguishing mild disease from normal retinal appearance. This result is still valuable because it identifies where future work should focus: higher-fidelity synthesis, lesion-aware crops, and stronger minority-class validation.

F. Limitations and Scope

The main limitations are the small number of GAN training pairs, domain differences between datasets, and the gap between fixed-split and patient-aware evaluation. The current study should therefore be viewed as a solid research prototype rather than a final clinical model. Even so, it provides a complete framework that can be extended with stronger generators, higher-resolution training, or expert quality control.

IX. CONCLUSION

This work presented a GAN-augmented retinal fundus image analysis system for ocular disease screening. The project combined vessel segmentation, synthetic minority-class generation, disease classification, grouped cross-validation, and explainable inference within a single pipeline. The GAN stage successfully increased glaucoma and AMD class availability and demonstrated a practical research path for addressing data imbalance in retinal imaging.

At the same time, the experimental results showed that synthetic data must be evaluated carefully, because the GAN-retrained classifier did not surpass the strongest non-GAN baseline. Therefore, the project contributes both an end-to-end implementation and an evidence-based discussion of the benefits and limitations of GAN augmentation in retinal disease classification.

X. ACKNOWLEDGMENT

The authors thank the Department of Computer Science and Engineering, Raghu Institute of Technology, and the project guide for supporting this B.Tech research work on GAN-assisted retinal disease classification and explainable medical image analysis.

REFERENCES

- [1] Andrew MVD, "Ocular Disease Recognition (ODIR5K)," Kaggle Dataset. [Online]. Available: <https://www.kaggle.com/datasets/andrewmvd/ocular-disease-recognition-odir5k>
- [2] Pradosh123, "Retinal Vessel Segmentation Combined," Kaggle Dataset. [Online]. Available: <https://www.kaggle.com/datasets/pradosh123/retinal-vessel-segmentation-combined>
- [3] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in Proc. MICCAI, pp. 234-241, 2015.
- [4] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks," in Proc. ICCV, pp. 2223-2232, 2017.
- [5] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. ICML, pp. 6105-6114, 2019.
- [6] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," in Proc. CVPR, pp. 2818-2826, 2016.
- [7] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in Proc. ICCV, pp. 618-626, 2017.
- [8] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in Proc. NeurIPS, pp. 4765-4774, 2017.
- [9] A. Budai, R. Bock, A. Maier, J. Hornegger, and G. Michelson, "Robust Vessel Segmentation in Fundus Images," *Int. J. Biomed. Imaging*, vol. 2013, pp. 1-11, 2013.
- [10] H. Pratt, F. Coenen, D. M. Broadbent, S. P. Harding, and Y. Zheng, "Convolutional Neural Networks for Diabetic Retinopathy," *Procedia Comput. Sci.*, vol. 90, pp. 200-205, 2016.