IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 11, November 2026

DOI: 10.17148/IJIREEICE.2026.141106

GLAUCOMA DETECTION USING CNN AND SVM CLASSIFICATION: A ROBUST AND MULTIMODAL FRAMEWORK

Sharada Shyam Sundar[1], P. Tritiksha Balaaji[2], Neelam Sanjeev Kumar[3]

Student, Department of CSE (Emerging Technologies), SRM Institute of Science and Technology, Vadapalani Campus, Chennai, India[1,2]

Assistant Professor (S.G), Department of CSE (Emerging Technologies), SRM Institute of Science and Technology, Vadapalani Campus, Chennai, India[3]

Abstract: Glaucoma damages the optic nerve slowly and silently, and vision that has been lost cannot be recovered. Screening with colour fundus photographs is therefore attractive, but deep learning models that perform well on curated test sets often struggle with images from other cameras, other clinics or eyes that carry a second disease. This paper describes and analyses a hybrid pipeline in which a convolutional neural network (CNN) is used only to extract features and a Support Vector Machine (SVM) makes the final decision. The design combines three kinds of evidence: segmentation-based biomarkers such as the vertical cup-to-disc ratio (vCDR), holistic CNN embeddings of the optic nerve head, and optional clinical variables including intraocular pressure and age. The distance of a case from the separating hyperplane, together with a Platt-scaled probability, lets the system answer "ungradable" instead of forcing a diagnosis. We review the evidence reported by the AIROGS challenge, G-RISK, PKRT-Net, a lightweight fuzzy segmentation network and GlaBoost, define an evaluation protocol built on AUC, sensitivity at 95% specificity, Dice and the Brier score, and discuss limitations. The study is analytical: no new clinical experiments are reported, and every performance figure quoted is taken from the cited publications.

Keywords: glaucoma screening, fundus photography, convolutional neural network, support vector machine, optic cup segmentation, out-of-distribution detection, robustness, multimodal learning.

I. INTRODUCTION

Glaucoma is not one disease but a family of optic neuropathies in which retinal ganglion cells die off gradually and the optic nerve head changes shape as a result. The damage is permanent. It is already a leading cause of irreversible blindness, and the number of people affected is expected to keep growing as populations age [1]. Raised intraocular pressure (IOP) is the best-known risk factor, although many patients have normal pressure [2]. In the common open-angle form there is no pain, and the early field loss is peripheral, so people tend to notice a problem only when central vision is already under threat. The Ocular Hypertension Treatment Study and the Early Manifest Glaucoma Trial showed that lowering IOP delays onset and progression [3], [4], which is why finding the disease early matters so much.

In the clinic, diagnosis rests on colour fundus photographs, optical coherence tomography (OCT) and visual-field tests. From the photograph an examiner estimates the vertical cup-to-disc ratio (vCDR), looks for thinning of the neuroretinal rim and checks whether the inferior–superior–nasal–temporal (ISNT) pattern has been broken. This grading is slow, differs from one grader to the next and cannot be scaled to whole populations. Deep learning has therefore attracted a great deal of interest, and convolutional neural networks (CNNs) [5] now reach very high accuracy on curated datasets. The trouble starts afterwards. When the same networks meet photographs from a different camera, poorly dilated pupils, imaging artefacts or eyes with diabetic retinopathy or other disease, performance drops and, worse, the prediction often stays confident [6], [7].

This paper examines one practical answer to that problem: a CNN that only extracts features, followed by a Support Vector Machine (SVM) that takes the decision [8], [9]. The SVM takes the place of the usual Softmax layer. Because it fits a maximum-margin boundary that depends on a small set of support vectors, it is a natural place to attach a reject option for images the system should not judge. The principal contributions are: (i) a three-part feature design that combines segmentation-derived biomarkers, holistic CNN embeddings and optional clinical variables; (ii) an SVM decision stage with Platt-scaled probabilities and an explicit "ungradable" output; (iii) a screening-oriented evaluation protocol based on AUC, sensitivity at 95% specificity, Dice and the Brier score, with patient-level and cross-cohort

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

DOI: 10.17148/IJIREEICE.2026.141106

splits; and (iv) a critical reading of recent evidence from AIROGS, G-RISK, PKRT-Net, a lightweight fuzzy segmentation network and GlaBoost. The paper is a design and analysis study. It reports no new clinical experiments, and every performance figure quoted comes from the cited publications.

II. RELATED WORK

Automated glaucoma detection has followed two routes. One segments the optic disc (OD) and optic cup (OC) and computes the vCDR, which is easy to explain to a clinician. The other classifies the whole photograph end to end. Recent work blends the two, with robustness, efficiency and multimodal input as the common themes. The AIROGS challenge [6] supplied roughly 113,000 photographs from about 60,000 people at 500 screening sites and, unusually, rewarded methods that could recognise ungradable or unexpected images; the top-scoring team detected ungradable images with an AUC of 0.99, and the headline glaucoma metric was sensitivity at a fixed 95% specificity. G-RISK [7] attacked generalisation from another side. A regression CNN trained on data from a single tertiary centre was tested on thirteen external sources from six countries, reaching AUC values between 0.854 and 0.988 on eleven public datasets and sensitivities of 87.3% and 90.3% at 95% specificity in two population cohorts.

Segmentation is harder than it looks, because the cup edge is a gradual transition rather than a sharp line and blood vessels run across it. PKRT-Net [10] splits the fundus image into a vessel-feature branch and a general local-feature branch, then uses a relation transformer to fuse clinical prior knowledge with the local features, and reports state-of-the-art results on Drishti-GS, RIM-ONE r3 and REFUGE. For devices with little computing power, Song et al. [11] proposed a compact network in which fuzzy learning absorbs the uncertainty of the cup boundary and a multi-layer perceptron keeps the architecture simple and fast to train. Dice scores above 95% for the disc and above 85% for the cup, as reported for a recent ViT–CNN network [12], indicate the accuracy needed for a trustworthy vCDR. Finally, GlaBoost [13] fuses fundus embeddings from a pretrained CNN, structured ophthalmic indicators and expert text descriptions in a gradient-boosting classifier and reports 98.71% validation accuracy on a real-world annotated dataset. Its fusion head is XGBoost rather than an SVM, but the idea of one classifier working on heterogeneous features carries over directly.

Placing an SVM on CNN features is not new. Tang [9] replaced the Softmax layer with a linear SVM and reported gains on several image benchmarks, and the soft-margin SVM itself goes back to Cortes and Vapnik [8]. Platt scaling [14] converts SVM scores into probabilities, whereas Hendrycks and Gimpel [15] showed that Softmax confidence is only a weak baseline for spotting out-of-distribution (OOD) inputs. Hybrid pipelines that pair learned representations with a classical classifier have also been used in other domains [16].

Four gaps motivate this work. Benchmark scores are usually measured on clean images, so they say little about how a system treats poor ones. Models are often trained at one centre and checked only on that centre's held-out split. Softmax confidence is rarely questioned even though it is the number a clinician would actually see. And segmentation, image-level classification and clinical data tend to be studied separately, although a real screening decision draws on all three.

III. DATASET DESCRIPTION AND PREPROCESSING

The pipeline is meant to be trained and checked on public fundus datasets (Table I and Fig. 3). AIROGS [6] offers about 113,000 photographs labelled as referable glaucoma, no referable glaucoma or ungradable, so it supports both detection and rejection. REFUGE [17] (1,200 images), Drishti-GS1 [18] (101), ORIGA [19] (650) and RIM-ONE [20] (159 in release 3) carry expert OD/OC masks and glaucoma labels and are used for segmentation and external validation.

TABLE I. PUBLIC BENCHMARK DATASETS FOR THE PROPOSED FRAMEWORK

Dataset	Modality	Volume	Annotations	Intended use here
AIROGS	Fundus	≈113,000 images	Referable / non-referable / ungradable	Robustness and ungradable-image detection; sensitivity at 95% specificity
REFUGE	Fundus	1,200 images	OD/OC masks, glaucoma label	Joint OD/OC segmentation and external validation
Drishti-GS1	Fundus	101 images	OD/OC masks, glaucoma label	Cup-boundary segmentation on a small dataset
ORIGA	Fundus	650 images	OD/OC masks, glaucoma label	vCDR extraction and cross-cohort testing

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

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 11, November 2026

DOI: 10.17148/IJIREEICE.2026.141106

Dataset	Modality	Volume	Annotations	Intended use here
RIM-ONE r3	Fundus	159 images	OD/OC masks, glaucoma label	Rim-tissue and segmentation precision

Preprocessing comprised: (1) discarding corrupt files and duplicate images; (2) mapping every dataset's labels to a common two-class scheme (referable versus non-referable glaucoma) while keeping "ungradable" as a separate label where it exists; (3) cropping a region of interest around the optic nerve head (ONH) with a small localiser network and resizing it to a fixed size such as 224 × 224 pixels; (4) normalising illumination, for example by applying Contrast Limited Adaptive Histogram Equalisation (CLAHE) [21] to the green channel, where the contrast between cup and disc is highest; (5) scaling pixel values with training-set statistics; (6) standardising clinical variables (age, IOP, central corneal thickness) and flagging missing values; and (7) augmenting training images with rotations, flips, small translations and brightness, contrast and hue jitter, so that the network sees some of the variation it will meet in new clinics.

Splits are made at patient level, so that the two eyes of one person or repeat visits never fall on both sides of a split. Anything learned from data, such as normalisation constants or SVM hyper-parameters, is fitted on training data only. For the cross-cohort test, entire datasets are held out: the model is fitted on some sources and evaluated on others it has never seen. No feature may use the label or any information from the test cohort.

IV. FEATURE ENGINEERING

The classifier receives one vector per eye, built from three groups of features. The first group, *b*, holds explicit biomarkers computed from the segmentation masks: the vCDR, the horizontal cup-to-disc ratio, disc and cup area, and neuroretinal rim area. The second group, *e*, is a latent embedding taken from the last pooling layer of a pre-trained backbone such as ResNet-50 [22] or DenseNet-121 [23]; it can capture rim defects, nerve-fibre-layer changes and peripapillary atrophy that no mask describes. The third group, *c*, contains optional clinical variables: age, IOP and central corneal thickness. When these are not available, a second SVM trained on imaging features alone is used instead.

$$vCDR = h_{cup} / h_{disc}, \quad z = [\, b \,;\, e \,;\, c \,], \quad \tilde{z} = (z - \mu) / s \tag{1}$$

Here h_{cup} and h_{disc} are the vertical diameters of the cup and disc, and μ and *s* are the training-set mean and standard deviation of each feature. Standardisation matters because the RBF kernel is built on Euclidean distance, and a feature with a large numeric range would otherwise dominate all the others. All feature definitions live in one reusable module that is called both when the model is trained and when it is applied to new images, which removes any mismatch between training and use.

V. SYSTEM ARCHITECTURE

The system is organised as a layered pipeline (Fig. 1). An acquisition and quality layer accepts a fundus photograph, rejects clearly unusable files and crops the ONH. A feature layer then runs two CNN branches in parallel: a segmentation network that yields OD and OC masks together with the biomarkers derived from them, and a backbone network that yields the latent embedding. The fusion layer concatenates both with the clinical variables and standardises the result. The decision layer is an RBF-kernel SVM, followed by Platt scaling for a calibrated risk score and a reject rule that can return "ungradable" instead of a label. Each layer has a narrow interface, so the segmentation network, the backbone or the classifier can be replaced independently, and the whole chain could be exposed as a web service to a screening clinic's existing software.

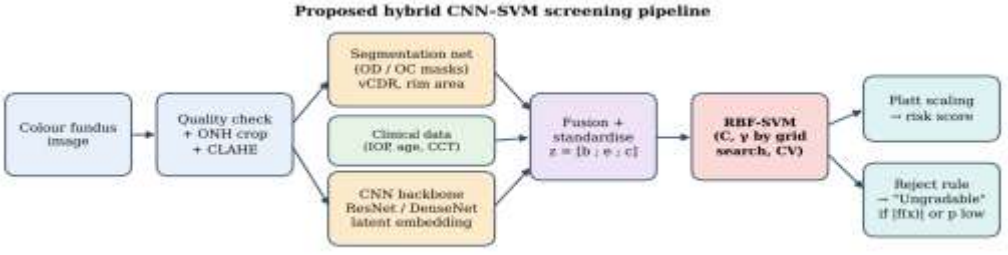


Fig. 1. Architecture of the proposed hybrid CNN–SVM glaucoma screening pipeline

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 11, November 2026

DOI: 10.17148/IJIREEICE.2026.141106

VI. METHODOLOGY

A. Feature extractors. A convolutional layer turns the feature maps of the previous layer into new ones through learned filters, $h_j^{(l)} = \sigma(\Sigma_i W_{ij}^{(l)} * h_i^{(l-1)} + b_j^{(l)})$, with the rectified linear unit $\sigma(x) = \max(0, x)$ as the usual non-linearity. After repeated convolution and pooling the network produces a vector e. In a standard classifier this vector would feed a Softmax layer; here the network is cut before that layer and e is passed on. For segmentation, a fuzzy-learning module and a multi-layer perceptron, following Song et al. [11], keep the model small and express the uncertainty at the cup edge, while a relation-transformer block in the spirit of PKRT-Net [10] brings in clinical prior knowledge. We propose training it with cross-entropy plus a Jaccard-distance term and a larger class weight on the small cup region, so that the network does not settle for predicting background everywhere.

B. SVM classifier. Given standardised training vectors z_i with labels $y_i \in \{-1, +1\}$ (non-referable, referable), the soft-margin SVM [8] solves

$$\min_{w,b,\xi} \; \tfrac{1}{2} \|w\|^2 + C \Sigma_i \xi_i \quad \text{subject to} \quad y_i(w^T \varphi(z_i) + b) \geq 1 - \xi_i, \;\; \xi_i \geq 0 \tag{2}$$

The margin is $2/\|w\|$, so a smaller $\|w\|$ means a wider margin, and the slack variables ξ_i allow a few overlapping examples to be misclassified, which is needed because early glaucoma and healthy eyes overlap. A large C punishes errors heavily and risks overfitting; a small C tolerates more violations and usually generalises better. In the dual form only the samples with non-zero Lagrange multipliers, the support vectors, define the boundary. The kernel trick replaces the inner product $\varphi(z_i)^T\varphi(z_j)$ with a kernel, and we use the Gaussian radial basis function (RBF) kernel

$$K(z_i, z_j) = \exp(-\gamma \|z_i - z_j\|^2), \qquad f(z) = \Sigma_{i \in S} \alpha_i y_i K(z_i, z) + b \tag{3}$$

where S is the set of support vectors and γ sets how far the influence of one support vector reaches: a small γ gives smooth boundaries and a large γ gives tight, overfitted ones. C and γ are chosen by grid search with patient-level k-fold cross-validation. The SVM can be fitted with LIBSVM [24] through scikit-learn [25].

C. Training, calibration and reject option. Training is carried out in two stages. The CNN is first trained with a proxy Softmax (or contrastive) head and then frozen; features are extracted in a forward pass and the SVM is fitted on them. A validation set, kept apart from the training patients, is used to fit Platt scaling [14], $p(y = 1 \mid f) = 1 / (1 + \exp(Af + B))$, and to set the reject thresholds. An image is returned as "ungradable" when it fails the quality check, when $|f(z)|$ falls below a margin threshold τ_f, or when the calibrated probability lies inside an uncertainty band around 0.5. The thresholds are tuned on the validation set so that few gradable images are rejected. For the final evaluation the model is trained on some cohorts and tested on others, and each held-out test set is used once, at the end.

VII. RESULTS AND ANALYSIS

Because this paper proposes and analyses a design, the figures below are those reported by the original studies on their own data. They are not directly comparable with one another, and none of them was produced by the pipeline proposed here. Table II gathers them.

TABLE II. PERFORMANCE REPORTED BY THE KEY STUDIES REVIEWED

Study	Task	Data	Reported result
AIROGS [6]	Screening with ungradable-image detection	≈113,000 photographs, 500 sites	Top team: AUC 0.99 for ungradable images; best teams comparable to 20 clinicians
G-RISK [7]	Risk regression, external validation	13 sources, 149,455 test images	AUC 0.854–0.988 (11 public sets); sensitivity 87.3% and 90.3% at 95% specificity
PKRT-Net [10]	OD/OC segmentation	Drishti-GS, RIM-ONE r3, REFUGE	State-of-the-art segmentation reported on all three
Song et al. [11]	Lightweight OD/OC segmentation	Public fundus datasets	Simpler model with improved segmentation accuracy reported
DST-Net [12]	Semi-supervised OD/OC segmentation	Public fundus datasets	Dice 95.12 (disc) and 85.69 (cup)

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

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

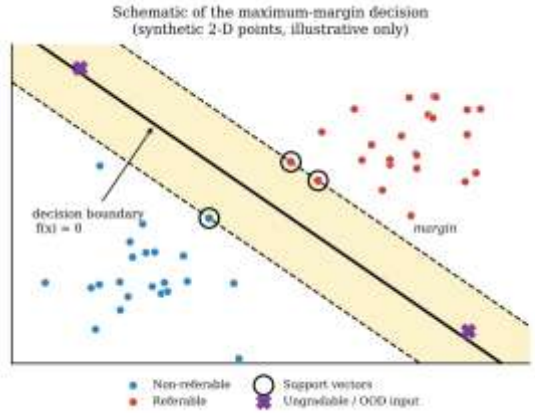
Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 11, November 2026

DOI: 10.17148/IJIREEICE.2026.141106

Study	Task	Data	Reported result
GlaBoost [13]	Multimodal risk prediction	Real-world annotated dataset	98.71% validation accuracy (gradient boosting)

Two points stand out. First, the screening studies measure what a clinic cares about, namely sensitivity at a fixed, high specificity and the ability to refuse bad images, rather than plain accuracy. Second, G-RISK shows that strong cross-cohort behaviour is possible even when training data come from one centre, provided the inputs are standardised carefully. It converts every photograph to a 30° disc-centred view, and the ONH cropping step in our preprocessing serves the same purpose.

Why an SVM stage might help. The argument is geometric rather than empirical. A Softmax layer is trained to make the logits for the right class large, and nothing stops it from producing extreme confidence far from the training data [15]. An SVM looks for the widest margin and depends only on its support vectors, so its decision score changes gradually near the boundary (Fig. 2, left) and a band around that boundary can be declared uncertain. Tang [9] found a linear SVM top layer to beat Softmax on image benchmarks, which is consistent with this view, but none of the glaucoma studies reviewed here tested it directly. We therefore treat the benefit as a hypothesis to be verified, and Table III sets out the differences.



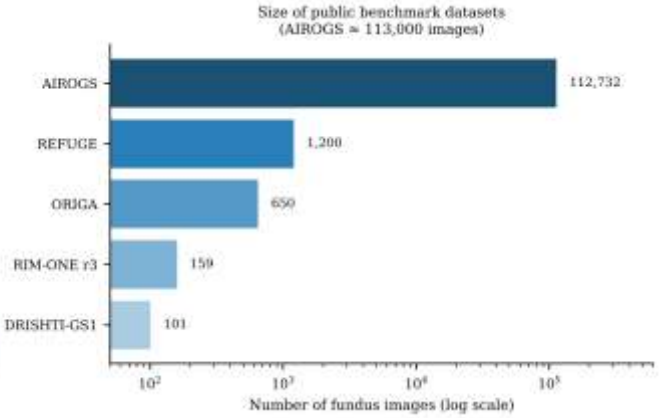


Fig. 2. Schematic of the maximum-margin decision (left) and Fig. 3. Size of the public benchmark datasets (right)

TABLE III. SOFTMAX HEAD VERSUS SVM HEAD ON CNN FEATURES

Aspect	Softmax head	SVM head (proposed)
Training objective	Cross-entropy, empirical risk	Hinge loss with margin penalty, structural risk
Decision depends on	All training samples, through the weights	Support vectors only
Confidence output	Softmax probability, often overconfident on OOD inputs	Margin distance plus Platt-scaled probability
Training procedure	End-to-end back-propagation	Two stages: CNN first, then SVM
Cost at $\approx 10^5$ images	Low	Higher for kernel SVM; linear or approximate solvers needed

Evaluation metrics. Segmentation is scored with the Dice coefficient and the intersection over union (IoU) between the predicted mask P and the expert mask G:

$$DSC = 2\,|P \cap G| \,/\, (\,|P| + |G|\,), \qquad IoU = |P \cap G| \,/\, |P \cup G| \tag{4}$$

Classification is judged by the AUC as a threshold-free summary, by sensitivity at 95% specificity (SE@95SP) as the screening operating point, and by the Brier score for the quality of the probabilities. The reject option is judged by the AUC for separating ungradable from gradable images and by the share of gradable images wrongly rejected. Accuracy alone is not enough, since a false negative means untreated disease and a false positive adds to the load on referral services.

DOI: 10.17148/IJIREEICE.2026.141106

Multimodal use. Clinical variables are appended to the feature vector before the SVM. The RBF kernel treats all standardised coordinates alike, so image and numeric information end up in one decision space, much as GlaBoost [13] combines imaging, structured and textual evidence. Whether an SVM or boosted trees is the better fusion head is an open empirical question.

Planned functional tests. Before any clinical use, the behaviours listed in Table IV should be verified end to end, including the case in which clinical data are missing.

TABLE IV. DESIGN TEST SCENARIOS FOR THE REJECT OPTION AND FUSION

Scenario	Expected behaviour (to be verified)
Clear, well-centred photograph of a healthy eye	Non-referable with a confident score
Clear photograph with large vCDR and rim thinning	Referable with a high score
Blurred or poorly lit photograph	Flagged as “ungradable”; no diagnosis issued
Photograph from a camera not seen in training	Score distribution shifts; flagged for review if confidence is low
Eye with a co-existing disease such as diabetic retinopathy	Flagged rather than labelled as glaucoma
Image without a visible optic nerve head	Rejected at the ROI stage
Clinical data missing	Imaging-only SVM is used and the output is marked as imaging-only
Borderline vCDR	Probability close to 0.5, reported as uncertain

VIII. ADVANTAGES AND LIMITATIONS

Advantages. The margin-based decision tends to resist overfitting when labelled data are scarce, which is the normal situation in medical imaging. Only the support vectors define the boundary, so the classifier stays compact. Every case comes with a distance and a probability, which gives a natural place for a reject option. The two-stage design lets a team reuse a pre-trained backbone and retrain only the SVM when data from a new clinic arrive. Clinical variables can be added without changing the architecture, inference is light enough for an ordinary clinic workstation, and the segmentation branch gives the clinician a readable vCDR next to the score.

Limitations. The main one is that nothing in this paper was tested experimentally: the benefits above are design arguments supported by related studies. The CNN and the SVM cannot be trained jointly by ordinary back-propagation, so the features are not tuned for the SVM, and kernel SVM training scales poorly beyond a few tens of thousands of examples. At AIROGS scale, linear solvers, stochastic gradient descent on the hinge loss or approximate kernels become necessary [24]. Distance from the hyperplane is also not the same as distance from the training data, so an unusual image can still land far on one side of the boundary; the reject rule should be paired with the quality check and, ideally, a feature-space novelty score. Labels are noisy and defined differently across datasets, the public sets are small and demographically narrow, and a fundus photograph cannot show structural change in depth as OCT does. A prospective, multi-centre study is needed before any claim about clinical performance.

IX. CONCLUSION AND FUTURE WORK

This paper set out a hybrid CNN–SVM framework for glaucoma screening and examined it against recent evidence. A CNN supplies segmentation biomarkers and holistic embeddings, optional clinical data are appended, and a margin-maximising SVM with Platt scaling makes the decision and can decline to judge images that are ungradable or unfamiliar. The reviewed studies show that robust screening, cross-cohort generalisation, anatomically informed segmentation, lightweight models and multimodal fusion are each achievable; the proposed design brings them together around one classifier whose confidence is easier to control than that of a Softmax layer. These advantages are argued, not yet measured. Future work includes a prospective multi-centre validation, joint or contrastive training of the CNN and SVM, fusion with OCT and visual-field data, density-based out-of-distribution scores, saliency-based explanations, subgroup fairness analysis, longitudinal progression modelling and deployment on smartphone-based fundus cameras.

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

DOI: 10.17148/IJIREEICE.2026.141106

REFERENCES

[1]. Y.-C. Tham, X. Li, T. Y. Wong, H. A. Quigley, T. Aung, and C.-Y. Cheng, "Global prevalence of glaucoma and projections of glaucoma burden through 2040," *Ophthalmology*, vol. 121, no. 11, pp. 2081–2090, 2014.

[2]. R. N. Weinreb, T. Aung, and F. A. Medeiros, "The pathophysiology and treatment of glaucoma: A review," *JAMA*, vol. 311, no. 18, pp. 1901–1911, 2014.

[3]. M. A. Kass *et al.*, "The Ocular Hypertension Treatment Study: A randomized trial determines that topical ocular hypotensive medication delays or prevents the onset of primary open-angle glaucoma," *Arch. Ophthalmol.*, vol. 120, no. 6, pp. 701–713, 2002.

[4]. A. Heijl *et al.*, "Reduction of intraocular pressure and glaucoma progression: Results from the Early Manifest Glaucoma Trial," *Arch. Ophthalmol.*, vol. 120, no. 10, pp. 1268–1279, 2002.

[5]. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[6]. C. de Vente *et al.*, "AIROGS: Artificial intelligence for robust glaucoma screening challenge," *IEEE Trans. Med. Imaging*, vol. 43, no. 1, pp. 542–557, 2024.

[7]. R. Hemelings *et al.*, "A generalizable deep learning regression model for automated glaucoma screening from fundus images," *npj Digit. Med.*, vol. 6, art. no. 112, 2023.

[8]. C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, 1995.

[9]. Y. Tang, "Deep learning using linear support vector machines," arXiv:1306.0239, 2013.

[10]. S. Lu, H. Zhao, H. Liu, H. Li, and N. Wang, "PKRT-Net: Prior knowledge-based relation transformer network for optic cup and disc segmentation," *Neurocomputing*, vol. 538, art. no. 126183, 2023.

[11]. Y. Song, W. Zhang, and Y. Zhang, "A novel lightweight deep learning approach for simultaneous optic cup and optic disc segmentation in glaucoma detection," *Math. Biosci. Eng.*, vol. 21, no. 4, pp. 5092–5117, 2024.

[12]. Y. Sun, T. Xu, J. Wang, and J. Wang, "DST-Net: Dual self-integrated transformer network for semi-supervised segmentation of optic disc and optic cup in fundus image," *Electron. Res. Arch.*, vol. 33, no. 4, pp. 2216–2245, 2025.

[13]. C. Huang, Z. Han, W. Xie, K. Kooner, T. Lee, J.-K. Wang, and J. Zhang, "GlaBoost: A multimodal structured framework for glaucoma risk stratification," arXiv:2508.03750, 2025.

[14]. J. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," in *Advances in Large Margin Classifiers*, MIT Press, 1999, pp. 61–74.

[15]. D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *Proc. Int. Conf. Learning Representations (ICLR)*, 2017.

[16]. N. S. Kumar, S. V. S, A. M. Enoch, K. T and W. C. R, "Hybrid Machine Learning Framework for Secure Information Verification in Social Media," *2025 IEEE 17th International Conference on Computational Intelligence and Communication Networks (CICN)*, Goa, India, 2025, pp. 907-913, doi: 10.1109/CICN67655.2025.11368151.

[17]. J. I. Orlando *et al.*, "REFUGE Challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs," *Med. Image Anal.*, vol. 59, art. no. 101570, 2020.

[18]. J. Sivaswamy, S. R. Krishnadas, A. Chakravarty, G. D. Joshi, and A. Syed Tabish, "A comprehensive retinal image dataset for the assessment of glaucoma from the optic nerve head analysis," *JSM Biomed. Imaging Data Pap.*, vol. 2, no. 1, p. 1004, 2015.

[19]. Z. Zhang *et al.*, "ORIGA-light: An online retinal fundus image database for glaucoma analysis and research," in *Proc. IEEE Eng. Med. Biol. Soc. (EMBC)*, 2010, pp. 3065–3068.

[20]. F. Fumero, S. Alayón, J. L. Sanchez, J. Sigut, and M. Gonzalez-Hernandez, "RIM-ONE: An open retinal image database for optic nerve evaluation," in *Proc. 24th Int. Symp. Computer-Based Medical Systems (CBMS)*, 2011, pp. 1–6.

[21]. S. M. Pizer *et al.*, "Adaptive histogram equalization and its variations," *Comput. Vis. Graph. Image Process.*, vol. 39, no. 3, pp. 355–368, 1987.

[22]. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[23]. G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4700–4708.

[24]. C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, art. no. 27, 2011.

[25]. F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.

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