




Semarak International Journal of Public Health and Primary Care

Journal homepage:
<https://semarakilmu.my/index.php/sijphpc/index>
ISSN: 3083-8401



DiRetina: Anatomically Guided Hybrid Deep Learning for Glaucoma Detection from Retinal Fundus Images

Mohammad Loqman Hakeem¹, Muhammad Irsyad Abdullah¹, Ahgalya Subbiah^{2,*}

¹ Faculty of Information Sciences and Engineering, Management and Science University, Malaysia

² Software Engineering and Digital Innovation Centre, Management and Science University, Malaysia

ARTICLE INFO

Article history:

Received 26 July 2026

Received in revised form 30 September 2026

Accepted 30 September 2026

Available online 9 October 2026

Keywords:

Glaucoma detection; retinal fundus imaging; deep learning; optic nerve head; hybrid feature fusion; medical artificial intelligence

ABSTRACT

Glaucoma is a major cause of irreversible visual impairment, and early screening remains difficult because disease progression may be asymptomatic while manual fundus assessment depends on specialist availability and observer judgement. This study develops DiRetina, a multi-stage artificial intelligence framework for automated glaucoma screening from color retinal fundus photographs, with the objective of improving robustness under heterogeneous acquisition conditions while retaining interpretable anatomical markers. A curated analytical dataset of 3,608 unique images was formed from seven public retinal repositories and stratified into training, validation, and test subsets using a 70:15:15:15 allocation. The pipeline applies contrast normalization, vessel-guided optic nerve head localization using U-Net++, Chan-Vese active-contour delineation of the optic disc and cup, vertical cup-to-disc ratio extraction, and hybrid feature fusion combining Local Binary Patterns with pretrained convolutional representations. A knowledge-enhanced attention module incorporates frozen retinal foundation-model embeddings to reduce task-irrelevant background influence. Baseline ResNet-50, VGG-16, and Xception models achieved 86.70%-88.46% accuracy, with specificity between 78.62% and 82.63%. The optimized Xception configuration reached 95.07% accuracy and an area under the receiver operating characteristic curve of 0.9880, while the hybrid feature-Support Vector Machine configuration achieved 96.4% accuracy, 95.1% sensitivity, 97.2% specificity, and an area under the curve of 0.981. Reported one-way analysis of variance results indicated significant performance differences among model variants at $p < 0.05$. The findings support a structured screening strategy that combines anatomical localization, handcrafted texture descriptors, and deep representations to improve diagnostic consistency across heterogeneous fundus images. Prospective clinical validation remains necessary before deployment.

* Corresponding author.

E-mail address: ahgalya_subbiah@msu.edu.my

1. Introduction

Glaucoma remains a major cause of irreversible vision loss and presents a substantial screening challenge because structural damage can progress before patients experience obvious symptoms. Colour retinal fundus photography offers a non-invasive and scalable basis for computer-assisted screening, particularly where specialist ophthalmic assessment is limited. Recent reviews show that deep learning has become central to automated glaucoma identification from retinal fundus images, although performance varies across datasets, acquisition devices, and evaluation protocols [1].

Automated screening commonly relies on clinically meaningful changes at the optic nerve head, including optic disc and optic cup morphology, neuroretinal rim thinning, and the vertical cup-to-disc ratio. Convolutional architectures can learn discriminative retinal representations, while hybrid pipelines can incorporate handcrafted descriptors or clinical measurements to strengthen decision support [7-9]. These approaches reduce dependence on subjective manual grading, but their reliability is influenced by image quality, class imbalance, and the extent to which the model focuses on anatomically relevant regions.

A persistent methodological limitation is that purely data-driven models may learn camera-specific or background artefacts rather than disease-related structures. Uneven illumination, vignetting, low contrast, and irregular optic disc-cup boundaries can degrade localisation and segmentation. Vessel-aware localisation, teacher-student segmentation, and retinal-knowledge attention have therefore been explored to constrain model attention toward clinically meaningful regions [2,6,10]. Cross-domain generalisation is an additional concern, especially when retinal appearance is altered by high myopia or when datasets are acquired using different imaging systems [12].

The research gap addressed in this study is the absence of a single screening workflow that consistently integrates anatomical localisation, explicit morphometric extraction, local texture information, and deep retinal representations while accounting for heterogeneous image acquisition and imbalanced class distributions. The significance of this work lies in treating anatomical structure as an explicit model constraint rather than relying only on end-to-end image classification. Accordingly, the objective of this study is to develop and evaluate DiRetina, a multi-stage glaucoma screening framework that combines vessel-guided optic nerve head localisation, active-contour disc and cup delineation, vertical cup-to-disc ratio measurement, hybrid Local Binary Pattern and convolutional features, and retinal knowledge-enhanced attention. The evaluation focuses on classification performance, structural interpretability, and reported cross-domain robustness across a curated multi-repository dataset.

2. Methodology

DiRetina was designed as a sequential screening pipeline in which retinal image enhancement, anatomical localisation, morphometric extraction, feature fusion, and classification are performed as linked stages. The operational configuration of the principal stages is summarised in Table 1.

Table 1

Operational parameters of the DiRetina image

Algorithmic Phase	Parameter / Model	Operational Value / Configuration
Dataset aggregation	Curated analytical set; stratified partition	3,608 unique images; 70:15:15 train/validation/test
Preprocessing and contrast	Target size; LAB contrast enhancement	224 x 224 pixels; clip limit 2.0; 8 x 8 tile grid
Optic nerve head localisation	Vessel segmentation; spatial grid	U-Net++ (depth L = 3); 6 x 6 grid
Disc/cup delineation	Morphology; active contour	Black-hat radius 15 px; Chan-Vese: disc 300 iterations, cup 150
Feature fusion and attention	Deep + texture features; retinal prior	ResNet-50/Xception + LBP; 1000-dimensional RETFound prior
Classification	Dense/Softmax or Support Vector Machine	Dropout 0.3-0.5; radial-basis-function kernel for SVM

2.1 Dataset and Study Design

The source manuscript reports a curated analytical dataset of 3,608 unique colour retinal fundus photographs assembled from seven publicly available repositories: G1020, ORIGA, REFUGE, PAPILA, ACRIMA, Drishti-GS, and RIM-ONE v3 [1,2]. The final analytical set comprises 804 glaucoma-positive images and 2,804 normal images. A class-aware stratified split allocated 70% of the images to training, 15% to validation, and 15% to testing. Class-weighted loss was used during deep-model training to reduce the influence of the majority class and preserve screening sensitivity [2,3]. The primary evaluation uses the fixed train-validation-test split, while the methodology also reports stratified five-fold cross-validation as an additional robustness procedure. This separation is retained because the source manuscript describes both procedures. The exact number of images available before curation and deduplication is not reported and therefore is not reconstructed here.

2.2 Retinal Image Preprocessing and Enhancement

Input images are standardised before anatomical localisation to reduce acquisition-dependent variation. The pipeline applies bilinear resizing, noise suppression using Gaussian or bilateral filtering, and Contrast Limited Adaptive Histogram Equalization to the luminance channel of the LAB colour space. The reported contrast configuration uses a clip limit of 2.0 and an 8 x 8 tile grid [3-5]. Colour-channel separation is then used to support anatomical analysis: the red channel emphasises the optic disc boundary, whereas the green channel provides stronger contrast for retinal vessels and the optic cup interface [6,7].

2.3 Optic Nerve Head Localisation and Delineation

Optic nerve head localisation follows a vessel-guided coarse-to-fine strategy. A U-Net++ network with nested skip connections and depth L = 3 segments retinal vessels. The binary vessel map is divided into a 6 x 6 grid, and the rows and columns with the highest local vessel density define the candidate optic nerve head region [2]. Within the reported 256 x 256 candidate region, a black-hat morphological operation enhances local transitions according to Eq. (1).

$$B_norm = [close(I_inp, S_15) - I_inp]_{norm} \quad (1)$$

In Eq. (1), I_inp is the input candidate-region image, S_15 is a disk structuring element with radius 15 pixels, $close(.)$ denotes morphological closing, and the subscript $norm$ denotes normalisation. Adaptive local Otsu thresholding with a 64×64 sliding window produces initial masks, after which Chan-Vese active contours refine the optic disc and cup boundaries. The reported iteration limits are 300 iterations for the optic disc and 150 for the optic cup, with local contrast enhancement used to stabilise low-contrast boundaries [2,6]. Least-squares ellipse fitting is applied to the refined optic disc and optic cup contours to obtain their vertical diameters. The vertical cup-to-disc ratio is then calculated using Eq. (2), providing an interpretable structural marker of glaucomatous change [2,11].

$$vCDR = D_cup,vert / D_disc,vert \quad (2)$$

In Eq. (2), $D_cup,vert$ is the vertical diameter of the optic cup and $D_disc,vert$ is the vertical diameter of the optic disc.

2.4 Deep Feature Extraction, Fusion, and Classification

The classification stage combines high-level convolutional representations with explicit local texture information. ResNet-50, VGG-16, and Xception are used as baseline or comparison backbones, while Local Binary Patterns capture micro-scale variations in the segmented neuroretinal region [8,9]. The source manuscript reports a fused representation of the form shown in Eq. (3).

$$F_fused = [\alpha g || \beta H_LBP] \quad (3)$$

In Eq. (3), g denotes the deep feature vector, H_LBP is the Local Binary Pattern histogram, α and β are learned scaling coefficients, and $||$ denotes vector concatenation. A Knowledge-Enhanced Convolutional Block Attention Module further introduces a frozen 1000-dimensional retinal representation derived from RETFound to weight backbone features and reduce sensitivity to task-irrelevant background patterns [10]. Final classification is performed either by a dense layer with dropout regularisation between 0.3 and 0.5 and a Softmax output, or by a Support Vector Machine with a radial-basis-function kernel [2,9].

2.5 Evaluation and Statistical Analysis

Performance is evaluated using accuracy, sensitivity, specificity, F1-score, and the area under the receiver operating characteristic curve; the methodology also identifies the Matthews Correlation Coefficient as a class-imbalance-aware measure [3,10]. The source manuscript reports one-way analysis of variance across six repetitions and paired t-tests for between-model comparisons, using a significance threshold of $\alpha = 0.05$. These tests are used to assess whether observed performance differences are unlikely to result from random training variation.

3. Results

3.1 Analytical Sample and Stratified Partition

After curation and deduplication, the reported analytical sample contains 3,608 unique fundus images. As shown in Table 2, 804 images are labelled glaucoma-positive and 2,804 are labelled normal. The stratified split yields 2,530 training images, 538 validation images, and 540 test images.

Table 2

Clinical sample distribution and stratified split counts

Class	Train (70%)	Validation (15%)	Test (15%)	Total
Glaucoma	565	119	120	804
Normal	1,965	419	420	2,804
Total	2,530	538	540	3,608

3.2 Classification Performance

Baseline performance indicates that the unaugmented and unweighted convolutional models maintain high sensitivity but lower specificity. Xception achieved 88.46% accuracy, 93.73% sensitivity, 82.63% specificity, and an area under the curve of 0.9402. ResNet-50 achieved 86.70% accuracy with 78.62% specificity, while VGG-16 achieved 87.80% accuracy with 82.22% specificity. The lower specificity across the baseline models is consistent with the class-imbalance concern identified in the study design.

The optimised configurations improved overall performance. The augmented Xception configuration reached 95.07% accuracy and an area under the curve of 0.9880. The hybrid feature-Support Vector Machine configuration achieved 96.4% accuracy, 95.1% sensitivity, 97.2% specificity, and an area under the curve of 0.981. Table 3 presents all metrics that are explicitly reported in the source manuscript; values not available in the source are marked NR rather than inferred.

Table 3

Classification performance across baseline and optimised model configurations

Model Configuration	Accuracy (%)	Loss (%)	Sensitivity (%)	Specificity (%)	F1 (%)	AUC
ResNet-50 (baseline)	86.70	22.24	93.73	78.62	87.13	0.9396
VGG-16 (baseline)	87.80	17.43	92.76	82.22	88.06	0.9398
Xception (baseline)	88.46	20.61	93.73	82.63	88.73	0.9402
Xception + augmentation	95.07	NR	NR	NR	NR	0.9880
Hybrid features + SVM	96.40	NR	95.10	97.20	NR	0.9810

3.3 Structural Delineation and Morphometric Findings

The source manuscript reports that vessel-guided localisation and active-contour refinement improve optic nerve head extraction under heterogeneous imaging conditions. It reports a localisation accuracy of 96.00% compared with 83.64% for a standard spatial-crop baseline. The derived vertical cup-to-disc ratio is used as the principal interpretable morphometric measure, and

the narrative results associate glaucomatous images with neuroretinal rim thinning, increased vertical cup-to-disc ratio, and reduced local vessel density.

3.4 Reported Confusion-Matrix Performance

The manuscript also reports class-level confusion counts for a 400-image evaluation subset. As shown in Table 4, 191 glaucoma images were true positives and 194 normal images were true negatives, with 9 false negatives and 6 false positives. These counts correspond to reported glaucoma precision of 96.9%, glaucoma recall of 95.5%, normal precision of 95.5%, and normal recall of 97.0%.

Table 4
Reported confusion-matrix counts and class-level evaluation

Actual Class	Predicted Glaucoma	Predicted Normal	Precision (%)	Recall (%)
Glaucoma (positive)	191	9	96.9	95.5
Normal (negative)	6	194	95.5	97.0
Class-level F1-score	96.2%	96.2%	N/A	N/A

3.5 Statistical Significance and Robustness

One-way analysis of variance was reported across six repetitions to compare model performance. Table 5 shows that all reported p-values are below 0.05, leading to rejection of the null hypothesis of equal performance among the evaluated model variants for the listed metrics. The source manuscript additionally reports paired t-tests with $p < 0.01$ for the proposed multi-stage system relative to standard convolutional baselines.

Table 5
Analysis of variance significance across reported performance metrics

Metric Evaluated	ANOVA p-value	Significant at $p < 0.05$
Accuracy	3.35×10^{-23}	Yes
Precision (control group)	7.52×10^{-17}	Yes
Recall (control group)	2.55×10^{-5}	Yes
F1-score (control group)	1.28×10^{-13}	Yes
Precision (glaucoma group)	1.58×10^{-15}	Yes
Recall (glaucoma group)	8.13×10^{-23}	Yes
F1-score (glaucoma group)	2.61×10^{-26}	Yes

The study further states that the knowledge-enhanced configuration retained stable behaviour under cross-dataset testing, although exact cross-domain accuracy, sensitivity, specificity, and area-under-the-curve values are not supplied. Consequently, the present revision retains the qualitative robustness claim but does not introduce unreported numerical results.

4. Discussion

4.1 Anatomical Priors and Vessel-Guided Localisation

The central result of DiRetina is that explicit retinal anatomy can be used as an inductive constraint before classification. Retinal vessels converge toward the optic disc and therefore provide a stable geometric cue for localisation. This approach reduces the search space presented to subsequent segmentation and feature-extraction stages. Comparable vessel-aware and disc-cup segmentation studies also emphasise the value of structured preprocessing for preserving clinically relevant boundaries [2,6,11]. By contrast, transformer-based approaches can learn long-range context without explicit cropping but require substantially greater representation capacity and may depend more heavily on large-scale pretraining [13].

4.2 Hybrid Texture and Deep Feature Fusion

The second major finding concerns the complementary role of handcrafted texture and deep representations. Local Binary Patterns explicitly encode local intensity variation, while convolutional backbones provide higher-level semantic features. Their combination is well aligned with the clinical character of early glaucomatous damage, where small changes in neuroretinal rim texture may coexist with broader optic nerve head morphology. The reported 96.4% accuracy of the hybrid feature-Support Vector Machine configuration is consistent with recent hybrid texture-learning work [9], while alternative studies show that multi-scale deep fusion or additional clinical measurements can also achieve strong performance [5,7].

4.3 Class Imbalance and Cross-Domain Generalisation

The baseline models demonstrate why accuracy alone is insufficient for imbalanced screening data. Although sensitivity remains above 92% across the three baseline convolutional networks, specificity falls to 78.62%-82.63%. The improvement obtained after augmentation, class weighting, hybrid features, and retinal knowledge integration suggests that model optimisation should address both representation quality and class distribution. This interpretation is compatible with recent work on optimisation and retinal knowledge-oriented attention [3,10]. Cross-domain robustness remains particularly important for images acquired with different cameras and for patients whose optic nerve head appearance is altered by high myopia [12]. Table 6 positions the reported DiRetina design against selected recent approaches already cited in the source manuscript. The comparison is intended to clarify methodological convergence and divergence rather than to claim direct superiority across studies with different datasets and evaluation conditions.

Table 6
Selected recent glaucoma-related approaches reported in the source manuscript

Study	Design Paradigm	Sample	Reported Metric	Interpretive Relevance
Medhi <i>et al.</i> , [2]	Vessel-aware U-Net++ localisation + active contour	3,608 images	Vessel-mask accuracy 96.56%	Supports anatomy-guided localisation
Kuna <i>et al.</i> , [9]	ResNet-50 + Local Binary Patterns + SVM	3,608 images	Accuracy 96.4%; AUC 0.981	Supports hybrid texture/deep fusion

Table 6 (Continued)

Study	Design Paradigm	Sample	Reported Metric	Interpretive Relevance
Gandhi <i>et al.</i> , [13]	Multimodal Vision Transformer	4,637 images	Accuracy 94.5%	Alternative global self-attention design
Guntreddi and V [8]	ResNet/VGG/Swin majority-voting ensemble	148 images	Accuracy 87.16%	Direct deep-ensemble baseline
Kilic [5]	InceptionV3 + DenseNet121 hybrid network	5,000 images	Accuracy 99.71%	High-capacity pure deep fusion
Makala and Kumar [7]	Deep imaging + clinical values + SVM	PAPILA dataset	Accuracy 98%; F1 97%	Uses physiological clinical priors

4.4 Limitations and Clinical Implications

Several limitations must be resolved before the manuscript can support clinical deployment claims. First, the initial raw image pool and source-specific image-resolution range are not supplied. Second, the reported confusion-matrix denominator does not match the stated 540-image test partition. Third, complete quantitative segmentation-overlap results and exact cross-domain validation metrics are absent. Fourth, the study is based on retrospective public datasets rather than a prospective clinical screening cohort. These gaps do not negate the reported classification findings, but they limit the strength of claims concerning generalisability, clinical workflow integration, and real-world reliability.

Within these constraints, the DiRetina architecture has practical value as an interpretable screening framework. Its multi-stage design explicitly exposes the anatomical localisation and morphometric steps that precede the final classification decision. Such visibility can support auditability and error analysis, particularly in resource-limited screening settings where a model should function as a referral-support tool rather than as a replacement for comprehensive ophthalmic assessment.

5. Conclusions

This study developed DiRetina as a multi-stage automated glaucoma screening framework that integrates retinal image enhancement, vessel-guided optic nerve head localisation, active-contour disc and cup delineation, vertical cup-to-disc ratio measurement, hybrid texture-deep feature fusion, and retinal knowledge-enhanced attention. The reported results show that baseline convolutional models achieved 86.70%-88.46% accuracy with lower specificity, whereas the optimised Xception configuration reached 95.07% accuracy and the hybrid feature-Support Vector Machine configuration reached 96.4% accuracy with 95.1% sensitivity and 97.2% specificity. These findings support the study objective by showing that explicit anatomical constraints and complementary feature representations can strengthen fundus-based glaucoma classification under heterogeneous and imbalanced data conditions. Before submission and clinical interpretation, the unresolved evaluation-denominator, missing segmentation metrics, and cross-domain reporting gaps identified in this revision should be verified. Future work should include prospective multi-centre validation and evaluation on low-power clinical or mobile screening platforms.

Acknowledgement

The authors acknowledge Management and Science University (MSU), Shah Alam, Malaysia, for institutional support.

References

- [1] Meedeniya, Dulani, Thisara Shyamalee, Gilbert Lim, and Pratheepan Yogarajah (2025). "Glaucoma identification with retinal fundus images: A systematic review of deep learning techniques." *Informatics in Medicine Unlocked* 56: 101644. <https://doi.org/10.1016/j.imu.2025.101644>
- [2] Medhi, Jyoti Prakash, Anchita Kakati, Prithi Deb, Rishiraj Mahanta, Kandarpa Kumar Sarma, Malaya Kumar Nath, Kirthika K., and Jayashree Natrajan (2026). "Precision in sight: A vessel-aware deep-learning framework for automated glaucoma screening." *Optik - International Journal for Light and Electron Optics* 349: 172728. <https://doi.org/10.1016/j.ijleo.2026.172728>
- [3] Shukla, Alok Kumar, Tejasya Vashisht, Shubhra Dwivedi, and Vikash Kumar Singh (2026). "Optimizing glaucoma detection using hybrid deep-evolutionary framework from retinal images." *Measurement: Sensors* 45: 102006. <https://doi.org/10.1016/j.measen.2026.102006>
- [4] Gul, Muhammed Furkan, Ozlem Polat, and Halit Bakir (2025). "TL-GWO: Fine-tuned transfer learning with grey wolf optimizer for retinal disease classification." *Experimental Eye Research* 260: 110598. <https://doi.org/10.1016/j.exer.2025.110598>
- [5] Kilic, Safak (2025). "HybridVisionNet: An advanced hybrid deep learning framework for automated eye disease classification." *Ain Shams Engineering Journal* 16 (1): 103594. <https://doi.org/10.1016/j.asej.2025.103594>
- [6] Chowdhury, Arindam, Ankit Lodh, Rohit Agarwal, Rahul Garai, Ahana Nandi, Narayan Murmu, Sumit Banerjee, and Debashis Nandi (2025). "Rim learning framework based on Teacher-Student Generative Adversarial Network (TS-GAN) for automated glaucoma screening." *Computers in Biology and Medicine* 187: 109752. <https://doi.org/10.1016/j.compbiomed.2025.109752>
- [7] Makala, Bindu Priya, and D. Manoj Kumar (2025). "An efficient glaucoma prediction and classification integrating digital imaging and clinical values." *Results in Engineering* 25: 104220. <https://doi.org/10.1016/j.rineng.2025.104220>
- [8] Guntreddi, Venkatesh, and Sivakumar V. (2025). "Deep learning-based glaucoma detection using majority voting ensemble of ResNet50, VGG16, and Swin Transformer." *Results in Engineering* 28: 107229. <https://doi.org/10.1016/j.rineng.2025.107229>
- [9] Kuna, Srinidhi, S. Jeeva, Ramesh Sekaran, and George Fernandez. "Hybrid Deep Texture Learning for Automated Glaucoma Detection from Fundus Images." *Procedia Computer Science* 282 (2026): 2364-2375. <https://doi.org/10.1016/j.procs.2026.04.074>
- [10] Liu, Chi, Yuzhuo Zhou, Sheng Shen, Zongyuan Ge, Fengshi Jing, Shiran Zhang, Yu Jiang, Anli Wang, Wenjian Liu, Feilong Yang, Tianqing Zhu, and Xiaotong Han (2026). "Fundus-image-based glaucoma screening via retinal knowledge-oriented multi-scale feature integration and attention learning." *Knowledge-Based Systems* 350: 116585. <https://doi.org/10.1016/j.knosys.2026.116585>
- [11] Kanwar, Akshay, Emjee Puthooran, and Vinod Kumar (2026). "Advanced glaucoma detection in obscured retinal images using Energy Mask R-CNN and Cyclic Code-based LSTM." *Biomedical Signal Processing and Control* 117: 109543. <https://doi.org/10.1016/j.bspc.2026.109543>
- [12] Goyanes, Elena, Joaquim de Moura, Antoine Lelong, Patricia Robles-Amor, Jose M. Martinez-de-la-Casa, Javier Moreno-Montanes, Francisco J. Munoz-Negrete, Francisco Ignacio Una, Jorge Novo, and Marcos Ortega (2026). "Domain-specific multi-transfer approaches for deep learning-based automated glaucoma detection in highly myopic eyes." *Expert Systems with Applications* 296: 129265. <https://doi.org/10.1016/j.eswa.2025.129265>
- [13] Gandhi, Vaibhav C., Priyesh P. Gandhi, Ali K. Abdul Raheem, Yasser Taha Alzubaidi, Kabul Khudaybergenov, and Mohammad Khishe (2026). "Advancing glaucoma diagnosis: Multi-modal deep learning with vision transformer architectures." *Intelligence-Based Medicine* 13: 100355. <https://doi.org/10.1016/j.ibmed.2026.100355>
- [14] Mayee, M. Kiran, M. Humera Khanam, and Shaik Lathifa Tabasum. "Proactive-bilstm: glaucoma detection using fundus and oct retinal images." *Expert Systems with Applications* 298 (2026): 129490. <https://doi.org/10.1016/j.eswa.2025.129490>
- [15] Tiwari, Shivam, Kalicharan Sahu, Surbhi Vijh, and Charu Awasthi. "Deep Learning-Based classification of ocular toxoplasmosis fundus images: A comparative study of CNN and SVM models." *Procedia Computer Science* 259 (2025): 1189-1197. <https://doi.org/10.1016/j.procs.2025.01.234>

