From Laboratory to SUS: Addressing Data Selection, Model Uncertainty, and Realism in Diabetic Retinopathy Lesion Detection Using a Brazilian Fundus Dataset

Laura Bernardes[1], Alejandro Pereira[2], Marcelo Dias[2], Marilton Aguiar[2], Daniel Welfer[3], Carlos Santos[1]

[1] Instituto Federal Farroupilha (IFFar) – Alegrete – RS – Brazil

[2]Universidade Federal de Pelotas (UFPel) – Pelotas – RS – Brazil

[3]Universidade Federal de Santa Maria (UFSM) – Santa Maria – RS – Brazil

laura.2023308995@aluno.iffar.edu.br

adspereira@inf.ufpel.edu.br

marcelo.sdias@inf.ufpel.edu.br

marilton@inf.ufpel.edu.br

daniel.welfer@ufsm.br

carlos.santos@iffarroupilha.edu.br

Abstract. *Diabetic retinopathy (DR) is a leading cause of vision loss in working-age adults. Early diagnosis depends on ophthalmological examination, but access remains limited in Brazil's public healthcare system (Sistema Único de Saúde — SUS) due to infrastructure and specialist shortages. Deep learning can support screening and triage, but progress is constrained by the scarcity of publicly labeled data. We curated and anonymized 13,131 color fundus images collected between 2012 and 2024 and selected 150 images to create an initial Brazilian dataset with DR grading and lesion annotations. We evaluated representative deep neural networks for classification, semantic segmentation, and lesion detection. RegNet Y 400MF achieved the best grading results (overall accuracy 0.6667, average accuracy 0.4634, Cohen's kappa 0.5179), R2U-Net showed the most balanced segmentation performance, and YOLOv11 obtained the highest mAP (0.3460 on validation; 0.3810 on test). The results highlight practical challenges, especially for microlesions, and support a SUS-oriented, deployment-aware perspective with attention to uncertainty, referral decisions, and cross-site robustness.*

1. Introduction

Diabetic Retinopathy (DR) is a major cause of vision loss in working-age adults. Around one-third (34.6%) of people with diabetes in the US, Europe, and Asia have DR. The incidence of diabetes is rising globally, leading to increased vision loss due to complications from the growing number of diabetics [Vocaturo and Zumpano 2020]. In Vocaturo et

al. [Vocaturo and Zumpano 2020], DR damages blood vessels in the retina and is a leading cause of vision loss for those aged 20 to 74. The longer someone has diabetes, the greater their risk of eye complications. DR is diagnosed by finding lesions on the retina, such as microaneurysms, hemorrhages, soft exudates, and hard exudates [Nayak et al. 2008].

Medical image analysis is vital in healthcare, aiding medical diagnoses for disease prevention and treatment. Early treatment can prevent vision loss from DR. However, early DR detection remains difficult. Diabetic patients typically visit endocrinology departments for fundus exams, which are time-consuming due to a limited number of daily exams. Additionally, there is a global shortage of ophthalmologists to meet increasing demands, especially in developing areas [Chakrabarti et al. 2012].

Therefore, developing computational tools to analyze and process fundus images can assist in screening and early identification of DR. Given the impact of DR and the complexity of its diagnosis, many automated methods using deep neural networks have been created to segment and detect these lesions. These methods face challenges due to the lack of public image datasets with expert-labeled lesions for training neural networks. This highlights the need for public datasets with annotated DR images. Such datasets can train and validate advanced deep neural networks, helping healthcare professionals achieve faster and more accurate DR diagnoses.

Beyond predictive performance, the translation of these methods into healthcare practice also depends on how predictions are embedded into clinical and organizational workflows. In public-health settings such as SUS, the practical value of a model is shaped not only by lesion-level accuracy, but also by how the system communicates uncertainty, supports triage, and helps define referral priorities. For this reason, this work is also positioned in the Information Systems domain, since it addresses how a dataset, model outputs, and decision points can be articulated in a screening-oriented information flow.

The primary contributions of this work include: (i) a new dataset for the classification, segmentation, and detection of retinal lesions associated with DR; (ii) the evaluation of this dataset through experiments with representative deep learning models; and (iii) a discussion of deployment-oriented issues involving uncertainty communication, annotation governance, referral logic, and the need for cross-site validation. It is important to note that this is only the first version of the dataset, and the goal is to increase the number of images and annotated lesions in future versions.

2. Related Work

This section reviews algorithms for detecting and segmenting lesions related to DR using deep learning and publicly available fundus image datasets.

Li et al. [Li et al. 2019] introduced the Dataset for Diabetic Retinopathy (DDR) and tested advanced deep learning models for classifying, segmenting, and detecting lesions. They collected 13,673 fundus images from 9,598 patients for clinical testing and annotated 757 diabetic images for fundus lesions: Microaneurysms, Hemorrhages, Hard Exudates, and Soft Exudates. The authors used the DDR dataset to evaluate deep learning

models, including segmentation and detection. A key limitation was recognizing microlesions.

Porwal et al. [Porwal et al. 2020] presented results from deep learning models for segmenting, classifying, and detecting fundus lesions in the IDRiD Diabetic Retinopathy Challenge. Their key contribution was the public release of the Indian Diabetic Retinopathy Image Dataset (IDRiD) dataset, which contains DR images. For lesion segmentation, teams primarily used the U-Net architecture [Ronneberger et al. 2015]. The detection challenge focused on identifying the Optic Disc and Fovea.

Deep neural networks classify, segment, and detect DR lesions. However, lesion size and the scarcity of annotated cases for training have limited deep learning methods. Despite their potential in medical image analysis, challenges remain, revealing a gap between research and clinical application. This issue involves the need for large training datasets for deep neural networks, as public DR image datasets have limited examples of various fundus lesions. Additionally, factors such as image quality, noise presence, and lesion annotation affect the accuracy of these methods.

From an Information Systems perspective, another gap concerns the way automated outputs are integrated into practical screening processes. Most prior studies emphasize model performance, while giving less attention to the organizational use of predictions, uncertainty handling, escalation policies, and referral support. In SUS-oriented screening scenarios, these aspects are essential for transforming model outputs into operational value.

3. Dataset Proposed Description

This work presents the Brazilian Diabetic Retinopathy Images Dataset (BDR-iD), created from 13,131 anonymized fundus images collected for public use at an ophthalmological clinic between 2012 and 2024, using a 45° FOV with a Canon CX1 retinograph and Canon BM7-0331 camera. Patient ages ranged from 0 to 99 years, with an average of 58; 58.82% were female and 41.18% male.

Because generating annotations for DR-associated lesions is complex, BDR-iD was limited to 150 fundus images: 150 for classification and 100 for detection and segmentation. We applied the following filters: (1) low resolution, (2) dark images, (3) glare, (4) blur, and (5) manual exclusion of unsuitable cases. The final set was classified and/or annotated for DR grading, lesion segmentation, and lesion detection. DR was categorized into: (i) Non-Proliferative Diabetic Retinopathy (NPDR) – mild, moderate, and severe; and (ii) Proliferative Diabetic Retinopathy (PDR). Table 1(a) shows the class distribution.

Retinal lesions for detection and segmentation include Hard Exudates (EX), Soft Exudates (SE), Microaneurysms (MA), and Hemorrhages (HE). The segmentation subset contains 100 images, of which 38 have DR lesions annotated at the pixel level. Table 1(b) summarizes lesion annotations.

The dataset includes image-level DR grading, pixel-level masks, and bounding-box annotations. Initial annotations were generated automatically and later reviewed by a

Table 1. Distribution of images across classes in BDR-iD for (a) classification and (b) segmentation/detection.

Classes	Quantity
No DR	54
Mild NPDR	23
Moderate NPDR	26
Severe NPDR	15
PDR	24
Unclassifiable	8
Total	150

(a)

Lesions	Quantity
EX	974
HE	318
SE	17
MA	307
Total	1616

(b)

specialist. In this version, a medical expert validated 100% of the automatically annotated EX and SE lesions. For the experiments, we used all available annotations, including cases still under validation, since restricting the analysis to only fully validated lesions would make segmentation and detection evaluation impractical at this stage. This should be interpreted as an incremental curation strategy rather than a final annotation protocol.

Future updates will aim to validate all lesion annotations and document inter-annotator variability on a duplicated-review subset, helping quantify agreement and better separate annotation uncertainty from model error, especially for MA and HE.

An ophthalmology specialist performed classification annotations on selected images, all backed by medical reports. A total of 88 DR images, 54 healthy eyes, and 8 undefined images were gathered, including 23 mild NPDR, 26 moderate NPDR, 15 severe NPDR, and 24 PDR (Table 1(a)).

We initially used models from Pereira et al. [Pereira et al. 2024] for binary mask generation based on R2U-Net. These models predict masks at a resolution of $256\times256\times3$, requiring resizing to the original image size, which we performed with the Upscayl tool using Generative Adversarial Networks (GANs). The selected model was Real-ESRGAN [Wang et al. 2021], which has shown utility in medical image enhancement [Aghelan and Rouhani 2022, Rashid et al. 2022, Nandal et al. 2024].

After generating binary masks, we converted them into an annotation format suitable for polygon editing and sharing. Edge detection was used to identify mask regions, and the final file was generated in COCO format. The annotations were then validated and corrected in CVAT. New binary masks were produced from the revised polygons to enable semantic segmentation, and bounding boxes were automatically derived. **Dataset publicly available**: BDR-iD dataset repository on GitHub.

4. Experiments

The evaluation of BDR-iD included three computer vision tasks: image classification, semantic segmentation, and object detection.

4.1. DR Grading

The dataset split was train 50%, validation 20%, test 30% (Table 2), with 512×512 inputs, augmentation (rotations up to 30°, flips, ColorJitter up to 20%, random resized crop), batch 16, Weighted Random Sampler, Adam (lr=0.001), 200 epochs. Metrics: per-class accuracy, OA, AA and Kappa.

Table 2. Class distribution across the train, validation, and test sets in the DR classification task.

Split	No DR	Mild NPDR	Moderate NPDR	Severe NPDR	PDR	Unclassifiable
Train	25	10	14	8	13	5
Validation	8	5	6	4	6	1
Test	21	8	6	3	5	2

Table 3. Validation (a) and test (b) results for grading.

Models	No DR	Mild NPDR	Moderate NPDR	Severe NPDR	PDR	Unclassifiable	OA	AA	Kappa
VGG-16	0.2700	0.0000	0.0000	0.0000	0.0000	0.0000	0.2667	0.0444	0.0000
ResNet-18	0.8600	0.6700	0.6700	0.4000	0.6200	0.0000	0.6333	0.5359	0.5461
GoogLeNet	0.8000	1.0000	0.6700	0.7500	0.5700	1.0000	0.7333	0.7980	0.6643
DenseNet-121	0.7300	0.0000	0.5000	0.3800	0.8000	0.0000	0.6000	0.4004	0.4958
EfficientNet B0	1.0000	1.0000	0.6700	0.0000	0.6200	0.0000	0.8000	0.5486	0.7461
RegNet Y 400MF	1.0000	1.0000	0.5000	1.0000	0.6000	0.0000	0.7333	0.6833	0.6667
SqueezeNet	0.4400	0.0000	0.0000	0.5000	0.3000	0.0000	0.4000	0.2074	0.2151
SwinTransformer	0.4700	0.0000	0.0000	0.0000	0.3300	0.0000	0.3667	0.1340	0.1926

(a)

Models	No DR	Mild NPDR	Moderate NPDR	Severe NPDR	PDR	Unclassifiable	OA	AA	Kappa
VGG-16	0.4700	0.0000	0.0000	0.0000	0.0000	0.0000	0.4667	0.0778	0.0000
ResNet-18	0.7200	0.5000	0.5000	0.3300	0.4300	0.0000	0.6222	0.4143	0.4239
GoogLeNet	0.7800	0.5000	0.0000	0.1700	0.4000	0.0000	0.6000	0.3074	0.4048
DenseNet-121	0.7300	1.0000	0.3300	0.1200	0.4000	0.0000	0.4889	0.4309	0.2930
EfficientNet B0	0.7900	0.5700	0.6000	0.3300	0.4000	0.0000	0.6444	0.4494	0.4853
RegNet Y 400MF	0.8300	0.4500	0.3300	0.5000	0.6700	0.0000	0.6667	0.4634	0.5179
SqueezeNet	0.6000	0.0000	0.0000	0.0000	0.2100	0.0000	0.3556	0.1351	0.1265
SwinTransformer	0.5000	0.0000	0.0000	0.0000	0.1700	0.0000	0.2889	0.1111	0.0400

(b)

Overall, the models underperformed on the test set, suggesting overfitting and highlighting that classes with fewer samples hinder learning. RegNet Y 400MF led the test (OA 0.6667, AA 0.4634, Kappa 0.5179), especially in PDR (0.67) and Severe NPDR (0.50).

4.2. DR Lesion Segmentation

We trained U-Net, Attention U-Net and R2U-Net (256×256, batch 4, 50 epochs, Adam 0.001, ReLU/Sigmoid, ImageNet weights) with augmentations (flip H/V, elastic transform, grid distortion, optical distortion). Metrics: Acc, Sen, Pre, DC, IoU.

4.3. DR Lesion Detection

We trained YOLOv9e, YOLOv10-X and YOLO11x (640×640, batch 8, 50 epochs, early stopping 15, RAdam 0.001) with horizontal/vertical flips, elastic/grid distortion and ran-

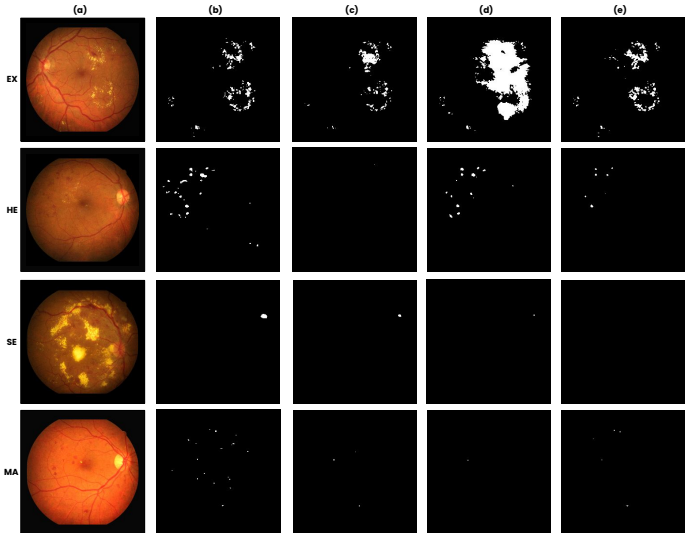

Figure 1. Visual comparison between the fundus lesion segmentations performed by the models with ground truth in images from the test set of the BDR-iD dataset. (a) Original images; (b) Ground Truth; (c) U-Net prediction; (d) Attention U-Net prediction; and (e) R2U-Net prediction.

Table 4. Segmentation results for validation (a) and test (b).

	EX					**HE**				
Models	*Acc*	*Sen*	*Pre*	*DC*	*IoU*	*Acc*	*Sen*	*Pre*	*DC*	*IoU*
U-Net	0.9973	0.3902	0.2806	0.2226	0.8701	0.9997	0.0387	0.5000	0.0520	0.9032
Attention U-net	0.9763	0.8003	0.0145	0.0257	0.6174	0.9993	0.4651	0.1259	0.1107	0.6775
R2U-Net	0.9898	0.3201	0.3925	0.2100	0.8666	0.9997	0.3004	0.4907	0.2229	0.8757
	SE					**MA**				
Models	*Acc*	*Sen*	*Pre*	*DC*	*IoU*	*Acc*	*Sen*	*Pre*	*DC*	*IoU*
U-Net	0.9995	0.0000	0.0000	0.0000	0.8997	0.9999	0.1214	0.3750	0.1445	0.9111
Attention U-net	0.9989	0.0022	0.0057	0.0008	0.8745	0.9999	0.0000	0.0000	0.0000	0.9249
R2U-Net	0.9995	0.0968	0.1557	0.0292	0.8524	0.9999	0.1456	0.5411	0.1751	0.9117

(a)

	EX					**HE**				
Models	*Acc*	*Sen*	*Pre*	*DC*	*IoU*	*Acc*	*Sen*	*Pre*	*DC*	*IoU*
U-Net	0.9958	0.3701	0.3875	0.2484	0.7902	0.9988	0.0406	0.4536	0.0515	0.8376
Attention U-net	0.9675	0.7207	0.0523	0.0722	0.6000	0.9986	0.3894	0.1784	0.1433	0.6979
R2U-Net	0.9959	0.3841	0.4706	0.2899	0.8245	0.9989	0.0861	0.7463	0.1188	0.8460
	SE					**MA**				
Models	*Acc*	*Sen*	*Pre*	*DC*	*IoU*	*Acc*	*Sen*	*Pre*	*DC*	*IoU*
U-Net	0.9999	0.2765	0.4814	0.2148	0.9713	0.9998	0.0857	0.2301	0.0632	0.8182
Attention U-net	0.9999	0.0638	1.0000	0.1200	0.9844	0.9998	0.0035	0.2222	0.0060	0.8499
R2U-Net	0.9998	0.0000	0.0000	0.0000	0.9166	0.9998	0.0762	0.3154	0.0901	0.8261

(b)

dom brightness. Metric: mAP@50.

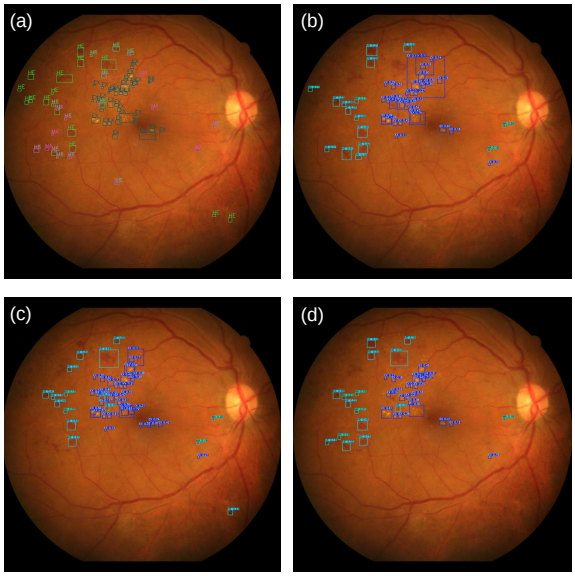

Figure 2. Fundus lesion detection performed by YOLOv9, YOLOv10, and YOLOv11 models on an image from the test set of the BDR-iD dataset. The annotated lesions are shown in (a), and the lesions detected by YOLOv9 in (b), YOLOv10 in (c), and YOLOv11 in (d), respectively.

Table 5. Detection mAP@50 for validation (a) and test (b).

Models	Yolov9	Yolov10	Yolov11
EX	0.0914	0.0967	0.1170
HE	0.3990	0.3980	0.3790
SE	0.5010	0.0000	0.6670
MA	0.2780	0.2580	0.2220
mAP	0.3170	0.1880	0.3460

(a)

Models	Yolov9	Yolov10	Yolov11
EX	0.2480	0.1530	0.2210
HE	0.2460	0.3640	0.3560
SE	0.5340	0.0000	0.7500
MA	0.1560	0.2290	0.1980
mAP	0.2960	0.1870	0.3810

(b)

5. Reflections and Provocations

This dataset foregrounds practical tensions between methodological ambition and clinical constraints. We surface reflections and productive provocations grounded in our evidence.

(i) On data sufficiency and class granularity

Our results suggest that intermediate grades (Moderate/Severe NPDR) are harder to separate than extremes. **Provocation:** rather than seeking larger end-to-end models, should DR pipelines adopt adaptive taxonomies where ambiguous cases are explicitly routed to human adjudication, turning uncertainty from failure into triage?

(ii) On microlesions and resolution

SE and MA remain challenging due to scarcity and scale. **Provocation:** is semantic segmentation the right target for microlesions in fundus? Would detection-first with lesion-centric crops and super-resolution pretext tasks provide better supervision than pixel masks at native resolution?

(iii) On generalization vs. overfitting

The gap between validation and test indicates instability. **Provocation:** should benchmarks emphasize cross-site validation even at small scale, rewarding robustness over absolute metrics, especially in early dataset versions and under different cameras or acquisition routines?

(iv) On annotation policy and governance

Automatic masks reviewed by experts accelerated curation, but create dependencies. **Provocation:** can hybrid protocols quantify model-in-the-loop bias, incorporate inter-annotator assessment, and require uncertainty reporting alongside COCO annotations to guide downstream users?

(v) On deployment reality

Given workforce shortages, models that reliably abstain may be more valuable than occasionally overconfident ones. **Provocation:** should DR systems be evaluated by *clinical utility under selective prediction* (coverage–risk curves) rather than only by mAP/OA?

6. Discussion

A realistic use scenario in SUS is a screening-and-referral workflow rather than autonomous diagnosis. In this scenario, fundus images may be acquired in primary care, diabetes care programs, or teleophthalmology actions by trained staff using standardized capture protocols. The images are then uploaded to an information system that stores the exam, links it to the patient record, and runs the DR analysis pipeline.

The system output should not be interpreted as a final diagnosis. Instead, it should support triage through three possible outcomes: (i) no evidence of referable findings, suggesting routine follow-up; (ii) suspected DR lesions or advanced grading patterns, indicating referral to specialist assessment; and (iii) low-confidence or poor-quality cases, which should be explicitly flagged for manual review or image reacquisition. In this workflow, the relevant actors include image acquisition staff, primary-care teams, telehealth operators, and ophthalmologists, while the critical decision points involve image quality control, uncertainty-aware triage, and referral prioritization.

Under this view, the contribution of the proposed approach is not to replace ophthalmologists, but to organize information flows that help SUS allocate specialist attention more efficiently, especially in contexts of limited infrastructure and workforce shortage.

Classification. Differences between validation and test support the hypothesis of overfitting and class imbalance effects. Intermediate stages (Moderate/Severe NPDR) concentrate errors, consistent with visual overlap in lesion burden and heterogeneous image quality. Beyond more data, two design levers emerge from our evidence: (i) *curriculum strategies* that gradually increase class granularity (binary DR/no-DR $\rightarrow$ NPDR/PDR $\rightarrow$ fine grades) and (ii) *selective prediction*, enabling abstention on borderline grades to preserve precision in operational settings. In a SUS-oriented workflow, uncertainty should

be communicated explicitly as part of the system output, for example by combining predicted class, confidence score, and a referral recommendation. Low-confidence cases should not be silently forced into a hard label; instead, they should trigger manual review, reacquisition, or direct referral depending on local protocol.

Segmentation. High accuracy with low Dice/IoU indicates background dominance; Attention U-Net raises sensitivity but inflates false positives, while R2U-Net yields a more stable balance for EX. The very low representation of SE and MA limits learning, pointing to the value of targeted enrichment through lesion-centric crops, weak supervision strategies, and consistency regularization to stabilize tiny-object boundaries.

Detection. YOLOv11 improves mAP mainly via SE gains, but MA remains the hardest class, suggesting a resolution–anchor mismatch. Our configuration hints that higher effective resolution or multi-scale heads tuned to sub-50px instances could yield disproportionate returns without changing labels.

Threats to validity. (i) Single-source imagery in v1 risks site bias and may limit transferability to other clinical environments; (ii) automatic-to-manual annotation pipelines may retain shape artifacts or model-induced bias; (iii) limited counts of SE and MA constrain conclusions for these classes; (iv) metrics averaged across classes can mask clinically meaningful regressions in difficult grades; and (v) the present study does not yet quantify inter-annotator variability on a duplicated subset, which is important for understanding label uncertainty in subtle lesions.

Reproducibility and reuse. To mitigate small-n effects while preserving anonymity, we documented splits and transformations and provide masks/boxes in a standard format. We encourage users to report results under our split and a cross-validated variant to assess sensitivity to sampling.

Baseline diversification. This study focused on deep learning baselines aligned with current literature. Future versions should also compare these models with more conventional computer vision baselines to better characterize the effective gains of data-driven methods under small-sample conditions.

7. Conclusions and Future Work

This paper presents the BDR-iD dataset and baselines for DR grading, segmentation, and detection. The results indicate that segmentation and detection are more challenging than grading, mainly due to the scarcity of small and rare lesions such as SE and MA. RegNet Y 400MF achieved the best grading results on test (OA 0.6667, AA 0.4634, Kappa 0.5179), R2U-Net showed a more balanced segmentation behavior, and YOLOv11 obtained the highest mAP and better generalization.

Beyond predictive performance, this work is also framed as an Information Systems contribution, since the dataset and models are treated as components of a SUS-oriented screening and triage workflow, in which uncertainty communication, referral logic, and human oversight are central. In this setting, the most realistic role of the system is to support prioritization and screening rather than replace specialist diagnosis.

We will expand the dataset and expert-validated lesions, quantify inter-annotator variability, study cross-site validation under different cameras and acquisition conditions, incorporate uncertainty-aware abstention and referral policies, and compare deep models with conventional computer vision baselines.

Ethics Approval

The proposed research project has already been processed on the Plataforma Brasil platform and has received the Certificate of Ethical Appreciation Presentation (CAAE) No. 7723523.1.0000.5346. The Research Ethics Committee (CEP) of the Universidade Federal de Santa Maria (UFSM) has approved the project, according to Consolidated Opinion No. 5.959.406.

References

Aghelan, A. and Rouhani, M. (2022). Fine-tuned generative adversarial network-based model for medical image super-resolution. *arXiv preprint arXiv:2211.00577*.

Chakrabarti, R., Harper, C. A., and Keeffe, J. E. (2012). Diabetic retinopathy management guidelines. *Expert Review of Ophthalmology*, 7(5):417–439.

Li, T., Gao, Y., Wang, K., Guo, S., Liu, H., and Kang, H. (2019). Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Information Sciences*, 501:511–522.

Nandal, P., Pahal, S., Khanna, A., and Pinheiro, P. R. (2024). Super-resolution of medical images using real esrgan. *IEEE Access*.

Nayak, J., Bhat, P. S., Acharya U, R., Lim, C. M., and Kagathi, M. (2008). Automated identification of diabetic retinopathy stages using digital fundus images. *Journal of Medical Systems*, 32(2):107–115.

Pereira, A., Santos, C., Aguiar, M., Welfer, D., Dias, M., de Menezes, R., and Santana, D. (2024). Enhanced semantic segmentation of retinal microlesions through r2u-net architecture. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 13–24, Goiânia, GO. SBC.

Porwal, P., Pachade, S., Kokare, M., Deshmukh, G., Son, J., Bae, W., Liu, L., Wang, J., and Liu, X. (2020). IDRiD: Diabetic Retinopathy – Segmentation and Grading Challenge. *Medical Image Analysis*, 59.

Rashid, S. I., Shakibapour, E., and Ebrahimi, M. (2022). Single mr image super-resolution using generative adversarial network.

Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9351:234–241.

Vocaturo, E. and Zumpano, E. (2020). The contribution of AI in the detection of the Diabetic Retinopathy. In *Proceedings of the IEEE International Conference on Bioin-*

formatics and Biomedicine, BIBM 2020, pages 1516–1519, Seoul, Korea, 16–19 December 2020. IEEE.

Wang, X., Xie, L., Dong, C., and Shan, Y. (2021). Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *International Conference on Computer Vision Workshops (ICCVW)*.

Author Biographies

Laura Bernardes is an undergraduate student in Systems Analysis and Development at the Federal Institute of Education, Science and Technology Farroupilha (IFFAR), Alegrete, Brazil. A research fellow within the project, her current work focuses on the application of artificial intelligence (AI) in healthcare.

Alejandro Pereira received his M.Sc. in Computing and B.Sc. in Computer Science from the Federal University of Pelotas (UFPEL). Between 2021 and 2022, he was a CNPq research scholar, investigating the design of cyber-physical systems for pest management. His research expertise lies in AI, with an emphasis on computer vision applied to medical diagnostics and health-related challenges.

Marcelo Dias holds an M.Sc. in Computing and a B.Sc. in Computer Science from the Federal University of Pelotas (UFPEL). Previously a PROBIC/FAPERGS research fellow, his academic interests center on artificial intelligence and object detection for medical imaging and diagnostic support.

Marilton Aguiar is a Professor at the Federal University of Pelotas (UFPEL), where he lectures in Computer Science and Engineering and contributes to the Graduate Program in Computing. He formerly served as the Director of the Center for Technological Development (CDTec) at UFPEL. His research investigates AI-driven solutions for environmental and public health sectors.

Daniel Welfer is a Professor in the Department of Applied Computing at the Federal University of Santa Maria (UFSM) and a permanent faculty member of the Graduate Program in Computer Science (PPGCC). His primary research areas include digital image processing, artificial intelligence, and software engineering.

Carlos Santos earned his Ph.D. in Computer Science from the Federal University of Pelotas (UFPEL) and is currently a Professor at the Federal Institute of Education, Science and Technology Farroupilha (IFFAR). His research focuses on AI applications in medical imaging, specifically Computer-Aided Detection (CAD) and Computer-Aided Diagnosis (CADx) systems.