



HAL
open science

DISCOVER: 2-D multiview summarization of optical coherence tomography angiography for automatic diabetic retinopathy diagnosis

Mostafa El Habib Daho, Yihao Li, Rachid Zeghlache, Hugo Le Boité, Pierre Deman, Laurent Borderie, Hugang Ren, Niranchana Mannivanan, Capucine Lepicard, Béatrice Cochener, et al.

► **To cite this version:**

Mostafa El Habib Daho, Yihao Li, Rachid Zeghlache, Hugo Le Boité, Pierre Deman, et al.. DISCOVER: 2-D multiview summarization of optical coherence tomography angiography for automatic diabetic retinopathy diagnosis. *Artificial Intelligence in Medicine*, 2024, 149 (March), pp.102803. 10.1016/j.artmed.2024.102803 . hal-04437753

HAL Id: hal-04437753

<https://imt-atlantique.hal.science/hal-04437753>

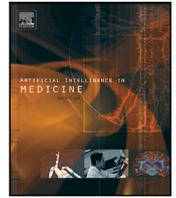
Submitted on 27 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Research paper



DISCOVER: 2-D multiview summarization of Optical Coherence Tomography Angiography for automatic diabetic retinopathy diagnosis

Mostafa El Habib Daho ^{a,b}, Yihao Li ^{a,b}, Rachid Zeghlache ^{a,b}, Hugo Le Boité ^{c,d}, Pierre Deman ^{e,f}, Laurent Borderie ^f, Hugang Ren ^g, Niranchana Mannivanan ^g, Capucine Lepicard ^d, Béatrice Cochener ^{a,h,b}, Aude Couturier ^d, Ramin Tadayoni ^{i,d}, Pierre-Henri Conze ^{j,b}, Mathieu Lamard ^{a,b}, Gwenolé Quellec ^{b,*}

^a Univ Bretagne Occidentale, Brest, F-29200, France

^b Inserm, UMR 1101, Brest, F-29200, France

^c Sorbonne University, Paris, F-75006, France

^d Service d'Ophthalmologie, Hôpital Lariboisière, APHP, Paris, F-75475, France

^e ADCIS, Saint-Contest, F-14280, France

^f Evolucare Technologies, Le Pecq, F-78230, France

^g Carl Zeiss Meditec, Dublin, CA 94568, USA

^h Service d'Ophthalmologie, CHRU Brest, Brest, F-29200, France

ⁱ Paris Cité University, Paris, F-75006, France

^j IMT Atlantique, Brest, F-29200, France

ARTICLE INFO

Keywords:

Diabetic retinopathy
Optical Coherence Tomography Angiography (OCTA)
Deep learning
Interpretability

ABSTRACT

Diabetic Retinopathy (DR), an ocular complication of diabetes, is a leading cause of blindness worldwide. Traditionally, DR is monitored using Color Fundus Photography (CFP), a widespread 2-D imaging modality. However, DR classifications based on CFP have poor predictive power, resulting in suboptimal DR management. Optical Coherence Tomography Angiography (OCTA) is a recent 3-D imaging modality offering enhanced structural and functional information (blood flow) with a wider field of view. This paper investigates automatic DR severity assessment using 3-D OCTA. A straightforward solution to this task is a 3-D neural network classifier. However, 3-D architectures have numerous parameters and typically require many training samples. A lighter solution consists in using 2-D neural network classifiers processing 2-D en-face (or frontal) projections and/or 2-D cross-sectional slices. Such an approach mimics the way ophthalmologists analyze OCTA acquisitions: (1) en-face flow maps are often used to detect avascular zones and neovascularization, and (2) cross-sectional slices are commonly analyzed to detect macular edemas, for instance. However, arbitrary data reduction or selection might result in information loss. Two complementary strategies are thus proposed to optimally summarize OCTA volumes with 2-D images: (1) a parametric en-face projection optimized through deep learning and (2) a cross-sectional slice selection process controlled through gradient-based attribution. The full summarization and DR classification pipeline is trained from end to end. The automatic 2-D summary can be displayed in a viewer or printed in a report to support the decision. We show that the proposed 2-D summarization and classification pipeline outperforms direct 3-D classification with the advantage of improved interpretability.

1. Introduction

Diabetic Retinopathy (DR), a complication of diabetes, is a major and growing cause of vision impairment and blindness. By 2040, around 600 million people throughout the world will have diabetes [1], a third of whom will have DR [2]. One major problem in the management of DR is its reliance on an older imaging technique, namely Color

Fundus Photography (CFP). Various classifications based on CFP were proposed over the years [3–5]. Unfortunately, decisions based on these classifications have poor predictive power. For instance, a severe non-proliferative DR case evolves to proliferative complication in 51.5% of cases, with only 17.1% evolving to a high risk of blindness [4]. This makes the management of DR challenging: clinicians often err

* Correspondence to: LaTIM - IBRBS - 22, avenue Camille Desmoulins 29200 Brest, France.
E-mail address: gwenole.quellec@inserm.fr (G. Quellec).

<https://doi.org/10.1016/j.artmed.2024.102803>

Received 24 August 2023; Received in revised form 19 December 2023; Accepted 3 February 2024

Available online 6 February 2024

0933-3657/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

on the side of caution and treat all those patients to mitigate the risk of complications. In the past decades, a significant number of studies have relied on CFP images for the automatic assessment of DR. The application of machine learning techniques, particularly deep learning, to these images has shown promising results in the detection and categorization of DR [6,7]. While these advancements are noteworthy, such techniques will always suffer from the poor predictive power of CFP. Fortunately, new imaging modalities are emerging that may improve predictions.

Optical Coherence Tomography (OCT) is a non-invasive imaging technique that uses interferometric information of partially coherent light to create cross-sectional (2-D B-scans) and three-dimensional (C-scans) structural images of biological tissues. Within optical scattering media, it can penetrate a few millimeters in depth, with micrometer resolution, and is therefore particularly well-suited to image the retina [8]. OCT Angiography (OCTA) is a motion-sensitive extension of OCT enabled by fast OCT acquisitions: it was shown to contrast the retinal vasculature and can provide quantitative blood flow information [9]. An OCTA acquisition can be summarized by two volumes: a *structure volume*, obtained by averaging consecutive 3-D scans, and a *flow volume*, describing the amplitude of local intensity variations across those consecutive 3-D scans [10]. To analyze the blood flow in specific vascular plexuses, clinicians generally inspect en-face (or frontal) Maximal Intensity Projections (MIP) of the flow volume in the corresponding retinal or choroidal layers [10]. In such 2-D projections, each 1-D A-scan is replaced with the maximal intensity value (throughout the entire A-scan or within the considered layers only). Recent OCTA devices enable ultra-widefield acquisitions of the retina (90°) [11]. In summary, OCTA can capture ultra-widefield volumetric structural and functional (flow) images of the retina: it is, therefore, a promising technique to diagnose various ocular pathologies [12–15]. In particular, DR can clearly benefit from OCTA [16]: (1) the structure volume allows objective and quantitative assessment of diabetic macular edema (DME), (2) flow MIPs allow quantification of retinal vascular plexuses, non-perfusion and vessel density as well as the identification of damage; Vujosevic et al. [17] lists the various biomarkers of DR and DME in OCTA acquisitions.

Computer-aided DR diagnosis using OCTA is an emerging field of research: it is motivated by the above promises (i.e., useful biomarkers) and the challenge of integrating large amounts of data (i.e., 3-D ultra-widefield structural and flow images). In particular, various quantitative metrics were automated to assist in early detection, staging, and progression of DR [18]. Those metrics quantify retinal fluid volumes [19], retinal vasculature features (e.g., density, tortuosity) [12, 20, 21], avascular zones [19, 22], including the Foveal Avascular Zone (FAZ) [12, 23], and proliferative DR features such as neovascularization [24].

Through a radiomics approach, these features were used for automatic DR severity assessment [21, 25, 26]. Various methods were also investigated to assess DR severity directly from OCTA images. Some authors classified 2-D en-face MIP images with 2-D Convolutional Neural Networks (CNN): [20, 26–29] classified one en-face flow MIP (superficial plexus, deep plexus or full retina), Heisler et al. [30] jointly classified two en-face flow MIPs (superficial and deep plexus) and the corresponding en-face structure MIPs, Yasser et al. [31] and Li et al. [23] jointly classified en-face flow MIPs and 2-D feature maps derived from feature segmentation. Other authors classified 3-D images with 3-D CNNs: Zang et al. [32] classified one 2-channel (structure and flow) 3-D image, Li et al. [33] jointly classified one 2-channel 3-D images and one 2-D Line-Scanning Ophthalmoscopy (LSO) localizer, and Li et al. [34] jointly classified two 2-channel 3-D images acquired with different fields of view ($6 \times 6 \text{ mm}^2$ and $15 \times 15 \text{ mm}^2$).

Theoretically, the radiomics approach and the 2-D classification approach are suboptimal in the sense that relevant features useful for classification may have been lost during the preprocessing (MIP) and feature extraction steps. However, the 3-D classification approach

also has some limitations. First, compared to their 3-D counterparts, 2-D neural architectures have better pre-trained weights (e.g., ImageNet weights) and have fewer parameters to optimize, thus requiring fewer training samples. Given the limited OCTA data sets available, this aspect is critical. Second, end-to-end 3-D classification lacks the interpretability power of the radiomics approach and, to a lesser extent, of the end-to-end 2-D classification approach. To alleviate these limitations, we propose an end-to-end 3-D image classification approach relying on 2-D views as intermediate steps, the 2-D view extraction process being trainable. This guarantees that: (1) 2-D neural architectures can be used at the end of the classification pipeline, while (2) relevant problem-specific features can be extracted at the beginning of the pipeline. In detail, two types of view are extracted:

1. 2-D en-face (or frontal) projections, generalizing the flow (or structure) MIP images used to assess OCTA flow features [16],
2. selected 2-D slices (B-scans), often used to assess structural OCT features [16]; attribution methods [35, 36] are used to identify the most relevant scan lines in the en-face projection, which are deemed worthy of further investigation in 2-D.

To maximize the interpretability of en-face projections, a novel “model dropout” mechanism is introduced. Projections are processed by an ensemble of 2-D image classifiers, and during training, classifiers in the ensemble are dropped randomly. The extracted 2-D features thus become more classifier-independent, i.e., more general and hopefully more meaningful to the human eye.

This paper aims to automate the most recent DR severity classification [5] from OCTA acquisitions, using the proposed interpretable classification approach, one step toward better management of DR.

2. Related methods

2.1. 3-D→2-D projection

In recent years, 3-D→2-D projection was used to solve various medical image analysis tasks. One such task is 2-D/3-D registration: in that case, the 2-D image is registered to a 2-D projection of the 3-D image. Fei et al. [37] used that strategy to register 2-D X-ray images to 3-D Computed Tomography (CT) images: the goal was to compare the ability of both modalities to detect cardiac calcifications. More recently, Schaffert et al. [38] and Jaganathan et al. [39] registered intra-operative 2-D X-rays to pre-operative 3-D CT images for navigation purposes during minimally invasive surgeries. Alternatively, Van Houtte et al. [40] registered intra-operative 2-D X-rays to a pre-operative 3-D atlas. For 2-D/3-D registration, 3-D→2-D projection is typically achieved by simulating a 2-D perspective projection from the 3-D image to the plane: in [38] and [39], a linear system of equations (point-to-plane correspondences) is solved; Van Houtte et al. [40] relies on projective spatial transformers. In Jaganathan et al. [39], an adversarial domain adaptation step is added.

3-D→2-D projection was also used for en-face segmentation of 3-D images. This is useful when ground truth masks are obtained using en-face projections of the 3-D images or using a different 2-D imaging modality. Li et al. [41] segmented blood vessels and foveal avascular zones in 3-D structural and flow OCTA images: the ground truth masks were obtained using the flow MIP between the Internal Limiting Membrane (ILM) layer and the Outer Plexiform Layer (OPL). Lachinov et al. [42] segmented geographic atrophies, a sign of age-related macular degeneration, in 3-D structural OCT images: the ground truth masks were obtained using 2D Fundus Auto Fluorescence (FAF) images. They also segmented blood vessels in 3-D OCT: the ground truth masks were obtained using en-face OCT images. Le et al. [43] segmented blood vessels in 3-D Adaptive Optics OCT (AO-OCT) images: ground truth masks were obtained using en-face AO-OCT images. For en-face segmentation of 3-D images, variations on the U-Net architecture are typically used: in the proposed architectures, the encoder part contains

3-D operations, and the decoder part contains 2-D operations [41,42]. Alternatively, an LSTM is combined with the encoder part of a 2-D U-Net in [43].

A variation on the previous task is the generation of high-quality 2-D images from 3-D images for visualization purposes or, optionally, for downstream segmentation tasks. Forsgren et al. [44] generated high-quality 2-D projections from low-quality 3-D fluorescence microscopy images, which can be acquired fast. In biology, Haertter et al. [45] projected curved 2-D manifolds from 3-D microscopy image stacks on a 2-D plane. Solutions to this problem are supervised: Haertter et al. [45] use a U-Net-like structure, Forsgren et al. [44] use conditional GANs with a U-Net-like backbone.

3-D→2-D projection was also used for fast feature detection or segmentation in 3-D images. In this case, ground truth masks are obtained from 3-D images: 2-D projections are simply used as intermediate steps. Shen et al. [46] detected 3-D junction points from various tree-like structures (blood vessels, neurons) in 2-D projections, followed by 2-D→3-D reverse mapping. Similarly, Wang et al. [47] segmented 3-D microvessels in 2-D projections of 3-D brain Magnetic Resonance Imaging (MRI), followed by 2-D→3-D reverse mapping. For this task, non-parametric 3-D→2-D projections are used, namely MIP [46,47]. Finally, Guo et al. [48] segmented blood vessels in OCTA images. A variation on MIP was used in that study: all voxel intensities in the same retinal layer are multiplied by a layer-specific trainable weight prior to MIP.

A final task explored in this paper is 3-D image classification. Statenko et al. [49] use 3-D→2-D projection to diagnose COVID-19-associated pneumonia in 3-D CT images. Gupta et al. [50] concatenates multiple 2-D projections of coronary arteries/branches to identify diseased coronary arteries/branches in 3-D Computed Tomography Angiography (CTA). Mandal et al. [51] use 3-D→2-D projection to differentiate lentigo maligna from atypical intraepidermal melanocytic proliferation, two melanoma subtypes, in 3-D Reflectance Confocal Microscopy (RCM). The motivation is to take advantage of 2-D neural architectures: compared to their 3-D counterparts, 2-D architectures have better pre-trained weights and fewer parameters to optimize [49,50]. Using 2-D architectures also reduces computational requirements [51]. For this task, non-parametric 3-D→2-D projections are also used: averaged 2-D slices [49], MIP [50,51].

Unlike previous 3-D image classification tasks, we propose a trainable parametric 3-D→2-D projection to allow the extraction of discriminant and interpretable problem-specific features. The proposed approach is more general than Guo et al. [48]’s solution for segmentation, which solely emphasizes specific retinal layers. Other trainable techniques proposed for (1) 2-D/3-D registration, (2) en-face segmentation of 3-D images, or (3) high-quality 2-D image generation assume that dense ground truth (e.g., 2-D images) is available to supervise 3-D→2-D projection training. This is not our case: the only supervision signal at our disposal is the DR diagnosis. This calls for a different architecture to ensure that feature localization is not lost in the en-face plane, as demonstrated in this paper.

2.2. Attribution methods

Over the past decade, with the growing popularity of deep learning, various methods were introduced to visualize what “black box” CNNs have learned. The purpose was notably to attribute an importance score to each image pixel for a given output target (e.g., a class prediction or a neuron activation). They can be classified into gradient-based or perturbation-based attribution methods.

The simplest gradient-based attribution method is the Saliency method by Simonyan et al. [52], which computes the gradient of the output target with respect to each image pixel. This approach is advantageous because it allows pixel-level attribution at the cost of a single gradient backpropagation. However, not all operations in a CNN are invertible, which often leads to unsatisfactory attributions. Zeiler

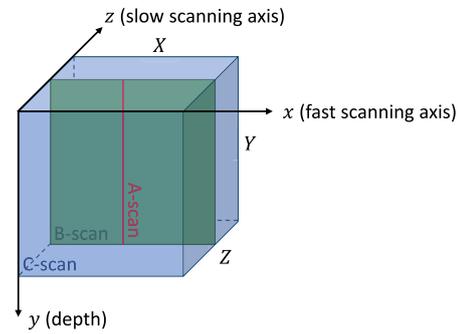


Fig. 1. Geometry of an Optical Coherence Tomography (OCT) acquisition. A 3-D B-scan consists of multiple 2-D B-scans, which in turn consist of multiple 1-D A-scans.

and Fergus [53] thus introduced a mechanism to correctly backpropagate attributions through MaxPool operations (Deconvolution method), and Springenberg et al. [54] introduced another way to correctly backpropagate them through ReLU activations (Guided Backprop method). These ideas were extended in Layer-Wise Relevance Propagation (LRP) and Deep Learning Important Features (DeepLIFT), with additional properties: LRP ensures that the magnitude of any output is conserved through the backpropagation process [55]; DeepLIFT bases the attributions on the difference between the activation of each neuron (for the current image) to its “reference activation” [56].

The simplest perturbation-based attribution method is the occlusion method by Zeiler and Fergus [53], which studies how the output target is impacted when parts of the image (square patches) are occluded. Ribeiro et al. [57] introduced Local Interpretable Model-agnostic Explanations (LIME), which looks for superpixels with the strongest association with the output target by successively turning superpixels off and on. Finally, Integrated Gradients can be regarded as both a gradient-based and a perturbation-based attribution method: (1) a set of images is generated by multiplying all pixel intensities in the target image by a constant factor ranging from 0 (the first image or reference) to 1 (the last image or target); (2) for each image in the series, a gradient-based attribution score is computed; (3) the final attribution is the integral of gradients over the series [35]. Intuitively, the most relevant features are detected early in the series and have, therefore, a higher integral. Unlike the other gradient-based attribution methods, multiple gradient backpropagation operations are needed. More generally, all perturbation-based attribution methods require multiple CNN evaluations, implying higher computation times and memory requirements.

To the best of our knowledge, this study is the first attempt to use attribution methods to extract and analyze relevant 2-D slices in a 3-D volume.

3. Proposed method

3.1. Overview and notations

The geometry of an OCT/OCTA acquisition is illustrated in Fig. 1. Let y denote the depth axis along which the partial coherent light penetrates the tissues (A-scan). Let x denote the fast scanning axis: a B-scan is thus indexed by x and y . Let z denote the slow scanning axis: a volume (C-scan) is thus indexed by x , y , and z and en-face projections by x and z . Let $X \times Y \times Z$ denote the size of the C-scans in voxels. Finally, a multi-channel volume (e.g., with a flow and a structure channel) is indexed by c , x , y , and z .

Given I , a preprocessed multi-channel OCTA acquisition (see Section 3.2), and N acquisition-level labels, the goal is to predict whether or not experts would assign the n th label to acquisition I , $n = 1..N$.

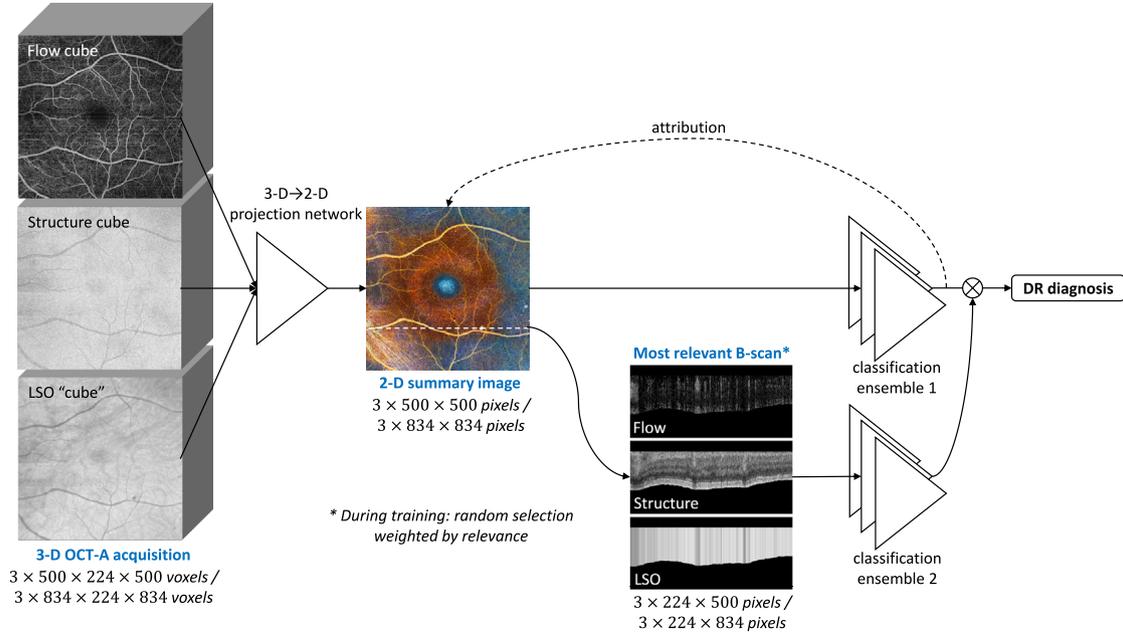


Fig. 2. Overview of the proposed approach. A multi-channel 3-D volume is summarized as a 2-D image through a 3-D→2-D projection network, detailed in Fig. 4. Next, a first classification branch classifies this 2-D summary image in order to produce a DR diagnosis. Through an attribution method, the most relevant 2-D B-scans are selected. Then, a second classification branch classifies the selected B-scans to improve the DR diagnosis. Each classification branch is an ensemble of classifiers, detailed in Fig. 5. In this figure, each 3-D channel in the input volume is represented by its Maximum Intensity Projection (MIP).

In the experiments (see Section 4), the goal is to assess DR severity in the patient's eye, according to the 5-level ICDR scale [5]: no DR (level 0), mild non-proliferative DR (NPDR, level 1), moderate NPDR (level 2), severe NPDR (level 3), proliferative DR (PDR, level 4); DR severity assessment is formulated as an N -label classification problem ($N = 4$): is DR severity $d(I)$ greater than or equal to level n ? Let $\mathbf{p}(I) = \{p_n(I) = p(d(I) \geq n) \in [0, 1], n = 1..N\}$ denote the probabilistic predictions and let $\lambda_n(I) \in \{0, 1\}$, $n = 1..N$, denote the ground truth labels.

As illustrated in Fig. 2, the preprocessed 3-D acquisition I is converted to a 2-D summary image $\Pi(I)$, defined as a parametric en-face projection of I (see details in Section 3.3). $\Pi(I)$ is defined as a color (3-channel) image for two reasons:

1. Interpretability: it can be displayed in a viewer or inserted in a report for human inspection.
2. Compatibility with off-the-shelf 2-D neural architectures with ImageNet pre-trained weights.

Next, $\Pi(I)$ is classified by a first ensemble C_1 of 2-D off-the-shelf image classifiers, as described in Section 3.4. A first estimation $\mathbf{p}^{(1)}(I)$ of the probabilistic prediction $\mathbf{p}(I)$ is given by $C_1 \circ \Pi(I)$.

Then, based on attributions derived from $\mathbf{p}^{(1)}(I)$, the most relevant B-scans $S(I)$ are selected, as described in Section 3.5. For interpretability purposes, the number of selected B-scans is limited to N : one per classification output. The motivations are:

1. A small number of B-scans is compatible with human inspection in a viewer or a report.
2. Each selected B-scan is associated with a severity cutoff and can, therefore, be used to document the course of action associated with that cutoff (treatment, follow-up, etc.).

Finally, the selected B-scans $S(I)$ are classified by a second ensemble C_2 of 2-D off-the-shelf image classifiers: a second estimation $\mathbf{p}^{(2)}(I)$ of the probabilistic prediction is given by $C_2 \circ S(I)$. It is combined with $\mathbf{p}^{(1)}(I)$ to obtain the final probabilistic prediction $\mathbf{p}(I)$, as described in Section 3.6.

3.2. Preprocessing

An OCTA acquisition is stored as two volumes:

1. a structure volume S , where the retinal layers and various retinal anomalies (e.g., fluid) are visible, among other structures (e.g., the choroid, below the retina, and the vitreous core, above it),
2. a flow volume F , where the blood vessels of the retina and the choroid are particularly highlighted.

Additionally, OCTA acquisitions are usually associated with:

1. A 2-D en-face localizer l , aligned with the OCTA data (size: $X \times Z$ pixels), to track eye motion. The PLEX Elite 9000 (Carl Zeiss Meditec Inc. Dublin, California, USA) device, for instance, is associated with a Line Scanning Ophthalmoscope (LSO) subsystem for that purpose [11].
2. Automatically-segmented surfaces delineating the vitreoretinal interface, namely the Inner Limiting Membrane (ILM), and the chorioretinal interface, below the Retinal Pigment Epithelium (RPE). Let s_{sup} and s_{inf} denote those two surfaces, respectively. They are stored as matrices of $X \times Z$ pixels: $s(x, z)$ represents the depth of surface s in the (x, z) A-scan of volumes S or F .

Although the LSO image is used primarily for motion tracking, it offers a complementary view on the retina (different optical properties, different wavelength, etc.), analogous to a grayscale fundus image [11]. We propose to analyze it jointly with the OCTA data. Therefore, we propose to preprocess an OCTA acquisition as follows (see Fig. 3).

First, an "LSO volume" L is created by duplicating the LSO localizer along the y-axis: $L(x, y, z) = l(x, z), \forall x, y, z$.

Second, a mask volume M of $X \times Y \times Z$ voxels is created: $M(x, y, z) = 1$ if $s_{inf}(x, z) \leq y \leq s_{sup}(x, z)$, 0 otherwise, $\forall x, y, z$. The flow, structure, and LSO volumes are multiplied by M , element-wise, to mask the choroid and vitreous core out. Let I' denote the 3-channel volume:

$$I' = [F \odot M, S \odot M, L \odot M]. \quad (1)$$

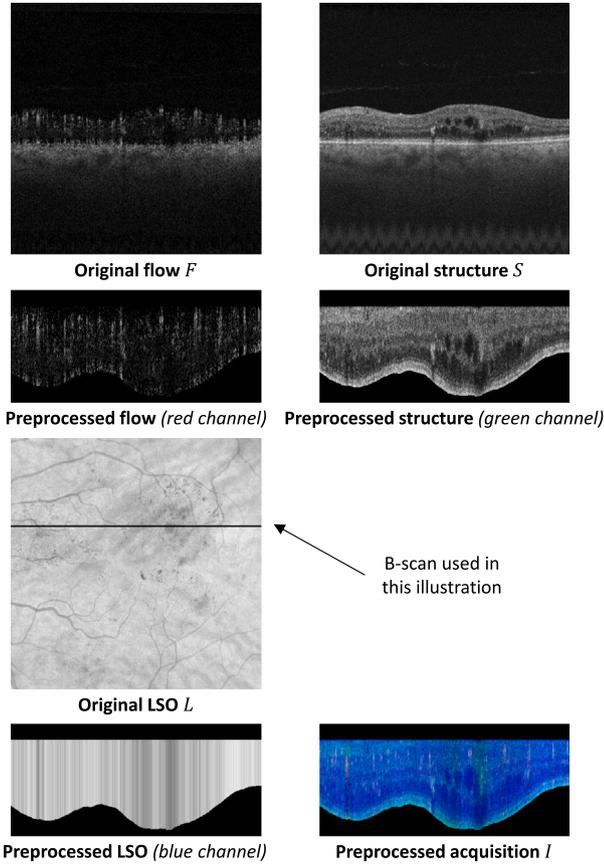


Fig. 3. Preprocessing pipeline for OCTA acquisitions (see Section 3.2) illustrated on one B-scan. Each original 2-D flow and structure B-scan is flattened, masked out, and cropped. The original 2-D LSO en-face localizer is transformed into a 3-D volume by duplicating pixel intensities along the depth axis (within the masked region). A 3-channel 3-D volume is obtained by stacking the resulting three volumes (flow, structure, LSO).

Third, the retinal region is flattened by shifting all voxels of I' along the y -axis, so that the ILM surface is set to a small constant depth $y = Y_0$. Let I'' denote the resulting volume:

$$I''(c, x, y + Y_0, z) = I'(c, x, y + s_{inf}(x, z), z), \forall c, x, y \in [0; Y - Y_0], z. \quad (2)$$

Parameter Y_0 is set to a non-zero value ($0 < Y_0 < Y$) to limit the loss of useful information during random data augmentation (see Section 3.6.1). Its value is determined as described in Section 4.3. This flattening process ensures that all the relevant information is concentrated at the top of I'' . The advantage of flattening the ILM surface, instead of the RPE (s_{inf}), is that its segmentation is less error-prone, due to a better contrast. This reduces the risk of alignment errors, which would lead to discontinuities between neighboring A-scans.

Fourth, I'' is cropped: all voxels with a depth $y > Y_1$ are discarded, $Y_0 \leq Y_1 \leq Y$. Parameter Y_1 is chosen to ensure the retinal region is never occluded. The final preprocessed image I denotes the cropped version of I'' .

3.3. 3-D→2-D projection

The preprocessed 3-D acquisition I is then converted to a 2-D summary image $\Pi(I)$ through a parametric 3-D→2-D en-face projection Π . Lachinov et al. [42] proposed a U-Net-like architecture for Π , where the encoder part contains 3-D operations and the decoder part contains 2-D operations. U-Net-like architectures have smaller and smaller activation maps as we go deeper into the encoder part, and their size

increases as we go deeper into the decoder part to finally reach the size of the input image. This contraction aims to increase the receptive field of deep encoder filters to better consider the context without increasing their size and, therefore, the number of network parameters. The drawback of this contraction is that small details are lost in the process. To recover those details, skip-connections are therefore introduced between encoder and decoder layers [58]. However, this trick assumes that the ground truth signal contains small details. In particular, it requires a dense supervision signal. For a classification task, the class labels are the only supervision signals available for training Π .

Our solution to the problem is to ensure that the details in the en-face plane are never lost throughout the 3-D→2-D projection process. In particular, we guarantee that the activation maps all have the same size in the en-face plane ($X \times Z$ pixels). Only the depth of these activation maps decreases as we go deeper in Π , to reach a final depth of 1 voxel (i.e., a 2-D image). To further prevent the loss of details in the en-face plane, we also limit the receptive field of the filters to one pixel in that plane. A simple solution based on 1-D operations, where each A-scan is processed independently, is presented hereafter. It should be noted that Li et al. [41]'s projection network for segmentation also ensures that activation maps all have the same size in the en-face plane. However, Li et al. [41] use 3-D convolution operators (kernel size: $3 \times 3 \times 3$ voxels): with a receptive field larger than one voxel in that plane, there is no guarantee that details are preserved without a dense supervision signal.

The proposed network, illustrated in Fig. 4, is divided into basic blocks containing:

1. a pooling operator,
2. two convolutional layers,
3. a batch normalization operator [59],
4. an optional skip-connection [60],
5. a ReLU activation.

Note that the pooling operator precedes the convolutional layers in order to limit network complexity: since no contraction in the en-face plane is performed, this is critical. An average pooling operator is used in the first block; otherwise, half of the voxels would never be used. However, a max pooling operator is used in the following blocks to add more nonlinearity. Following common practice [60], the number of convolutional filters increases as we go deeper into the network. Let Φ denote the number of filters per layer in the first block. The number of filters per layer in the i th block is set to $2^{i-1}\Phi$.

Each block reduces the depth by a factor of 4 (2 due to pooling \times 2 due to the stride in the first convolution layer). After three blocks, a global mean operator along the depth axis is performed to eliminate the depth dimension. Finally, a dense layer with sigmoid activation is applied to obtain a 2-D image with the desired number of channels, namely three channels (see Section 3.1). The sigmoid activation facilitates conversion to a bitmap image for visualization.

3.4. Classification of the 3-D→2-D projection

Now that a 2-D color image $\Pi(I)$ is obtained, any image classifier C_1 can be used to predict DR severity: a CNN, a transformer, an ensemble of CNNs and/or transformers, etc. However, two novel contributions specifically suited to this proposed framework are presented hereafter and illustrated in Fig. 5.

3.4.1. Ensuring 3-D→2-D projection generality

For interpretation purposes, we would like $\Pi(I)$ to be as independent from the classifier as possible. The rationale is as follows: if the projection is useful for any classifier, then we expect it to be informative for human experts as well. Various solutions can be considered:

- Following federated learning, multiple $\{\Pi^{(j)}, C_1^{(j)}\}$ couples can be trained in parallel, and the weights of the $\Pi^{(j)}$ instances can be aggregated at regular intervals.

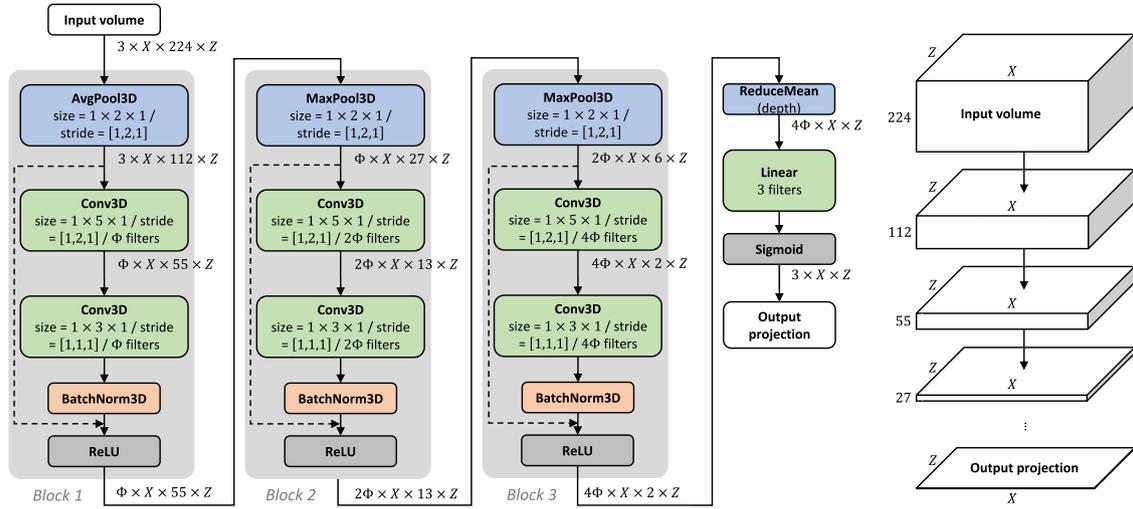


Fig. 4. Architecture of the 3-D \rightarrow 2-D projection network, detailed in Section 3.3. The input is the preprocessed acquisition, a 3-channel volume of $X \times Y_1 \times Z$ voxels, with $Y_1 = 224$ in this example. Parameter Φ , the number of filters in the first block, controls the complexity of the projection network. The figure on the right illustrates the size of the data tensors at the output of each block.

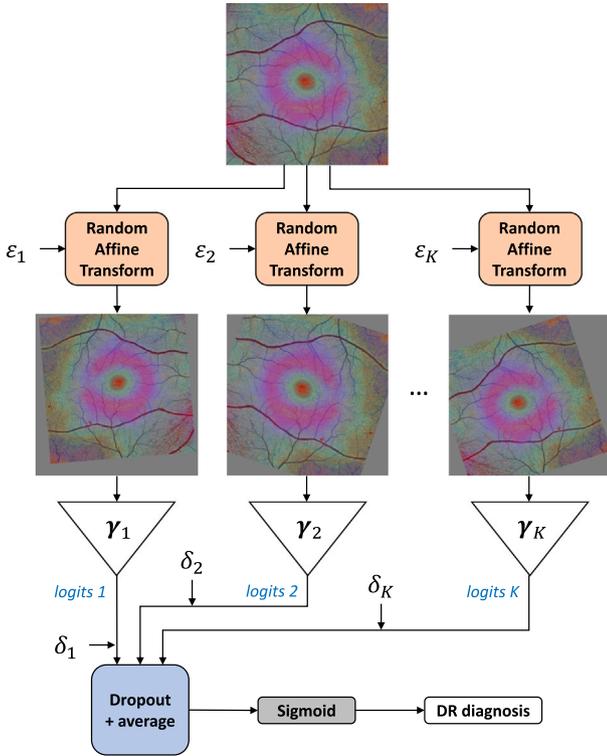


Fig. 5. Ensemble of classification networks $\{\gamma_k, k = 1..K\}$ with model dropout (controlled by random parameters $\delta_k, k = 1..K$) and differentiable random transformations (affine transformations and horizontal flips, controlled by random parameters $\epsilon_k, k = 1..K$). This pipeline, detailed in Section 3.4, is illustrated for the first classification branch of Fig. 2, in which the input images are 2-D summary images.

- Following continual learning, multiple classifiers $C_1^{(j)}$ can be trained sequentially, initializing training with the weights obtained for Π with the previous $C_1^{(j-1)}$ classifier.

However, such approaches imply longer training or require more resources. Instead, we propose to train one ensemble of classifiers, but with one trick that we call *model dropout*: for each mini-batch, a random subset of the classifiers is used for prediction. Like the other solutions mentioned above, this ensures that the 3-D \rightarrow 2-D projection Π does not

specialize for one specific classifier or one static ensemble of classifiers. Let $\gamma_k, k = 1..K$, denote the classifiers of the ensemble. We assume that these classifiers have no final activation function (i.e., they return logits). The ensemble prediction is given by:

$$C_1 \circ \Pi(I) = \sigma \left(\frac{\sum_{k=1}^K \delta_k \gamma_k(\Pi(I))}{\sum_{k=1}^K \delta_k} \right) \quad (3)$$

subject to: $\delta_k \in \{0, 1\}, k = 1..K$,

$$1 \leq \sum_{k=1}^K \delta_k \leq K,$$

where σ denotes the sigmoid function. The number of possible classifier combinations is given by $2^K - 1$. This process is equivalent to training $2^K - 1$ classifiers in random order, which is expected to improve the generality of Π . Model dropout is only used during training: the full ensemble is used during inference. The expected benefit of logit averaging (see Eq. (3)) is that the weight of poor classifiers in the ensemble can be reduced automatically by decreasing the amplitude of their logits. However, logit averaging may have drawbacks [61], so the traditional strategy, namely probability averaging, was also investigated.

3.4.2. Data augmentation

Data augmentation is typically performed by randomly transforming preprocessed images before feeding them to the neural network. For 2-D image classifiers, random transformations traditionally imply random affine transformations (random rotation, translation, and scaling) and random horizontal/vertical flips. However, our input preprocessed images are heavy 3-D volumes. Applying such random transformations to the 3-D volume takes a lot of time. Instead, we propose to apply them after the 3-D \rightarrow 2-D projection: applied to 2-D data, they are much faster. Besides, applying random spatial transformations before the projection is of limited interest since the proposed projection operator Π does not consider the context. Inserting random transformations inside the neural architecture is made possible by differentiable implementations of these transformations.¹

Since these random transformations can be inserted inside the neural network, we are able to generate one transformed version of $\Pi(I)$ for each classifier γ_k in the ensemble. This leads to a new definition for $C_1 \circ \Pi$:

$$C_1 \circ \Pi(I) = \sigma \left(\frac{\sum_{k=1}^K \delta_k \gamma_k(T(\Pi(I), \epsilon_k))}{\sum_{k=1}^K \delta_k} \right) \quad (4)$$

¹ <https://pytorch.org/vision/main/transforms.html>

subject to: $\delta_k \in \{0, 1\}, k = 1..K$,

$$1 \leq \sum_{k=1}^K \delta_k \leq K \text{ ,}$$

where T denotes the transformation operator and ϵ_k the random transformation parameters drawn for γ_k . As a way to generalize test-time data augmentation [62], random transformations are applied during both training and inference.

3.5. Relevant B-scan selection

A first estimation $\mathbf{p}^{(1)}(I) = C_1 \circ \Pi(I)$ of the probabilistic prediction $\mathbf{p}(I)$, based on the en-face 3-D→2-D projection Π , is now available. We propose to investigate further those B-scans of I which contribute the most to $\mathbf{p}^{(1)}(I)$. The idea is to find additional evidence to increase or decrease the confidence in this first estimation.

To detect the B-scans that contribute the most to $\mathbf{p}^{(1)}(I)$, we propose to use attribution methods presented in Section 2.2. Note that attributions are computed for one particular output prediction $p_n^{(1)}(I)$, i.e., for one DR severity cutoff. This aligns with our goal to collect additional evidence for each prediction: we will select one B-scan per prediction. As for the inputs, we can either:

- use the 3-D preprocessed acquisition I and accumulate voxel-wise attributions in the xy-plane.
- or use the 2-D projection $\Pi(I)$ and accumulate pixel-wise attribution along the x -axis (the fast scanning axis).

The second option was chosen for faster computations. Let $a_I(x, z, c, n)$ denote the attribution of pixel (x, z) , in the c th channel, for the n th prediction. A normalized attribution $\alpha_I(z, n)$ is defined for the z th B-scan, with respect to the n th prediction:

$$\alpha_I(z, n) = \frac{\sum_x \sum_c |a_I(x, z, c, n)|}{\sum_z \sum_x \sum_c |a_I(x, z, c, n)|} \text{ .} \quad (5)$$

Let z_n denote the index of the n th selected B-scan and let $B_n(I)$ denote its content:

$$B_n(I) = I(c, x, y, z_n), \forall c, x, y \text{ .} \quad (6)$$

For inference, the B-scans maximizing $\alpha_I(z, n)$, $n = 1..N$, are selected. However, for data augmentation purposes and to favor exploration, a random B-scan selection process is preferred during training: z_n is randomly drawn from the multinomial probability distribution defined by $\alpha_I(z, n)$:

$$z_n = \underset{z}{\operatorname{argmax}} \alpha_I(z, n) \quad \text{for inference ,} \quad (7)$$

$$z_n \sim M_Z(1; \alpha_I(1, n), \dots, \alpha_I(Z, n)) \quad \text{for training .} \quad (8)$$

It should be noted that B-scan selection is not impacted by random transformations from Section 3.4.2, which are an integral part of the C_1 classifier on which the attribution method operates (see Eq. (4)).

3.6. Final classification

3.6.1. Second classifier

Like classifier C_1 , the second classifier C_2 also requires data augmentation. Besides the random selection process described above, we propose to apply the same random transformation T as for classifier C_1 . More generally, we define C_2 very similarly to C_1 : an ensemble of classifiers γ'_k , $k = 1..K$, with random transformations (parameters: $\epsilon'_{k,n}$, $k = 1..K$, $n = 1..N$) and model dropout (parameters δ'_k , $k = 1..K$). Because their input images are of a different nature, no parameter sharing was set up between C_1 and C_2 .

By design, the n th selected B-scan $B_n(I)$ is meant to correct the confidence in the n th prediction. Therefore, we only consider the n th

prediction $\gamma'_{k,n}(B_n(I))$ of classifier γ' for B-scan $B_n(I)$. This leads to the following expression for the predictions $\mathbf{p}^{(2)}(I)$ of C_2 :

$$p_n^{(2)}(I) = \sigma \left(\frac{\sum_{k=1}^K \delta'_k \gamma'_{k,n}(T(B_n(I), \epsilon'_{k,n}))}{\sum_{k=1}^K \delta'_k} \right), n = 1..N \quad (9)$$

subject to: $\delta'_k \in \{0, 1\}, k = 1..K$,

$$1 \leq \sum_{k=1}^K \delta'_k \leq K \text{ .}$$

3.6.2. Final classifier

The second classifier C_2 is supposed to increase or decrease the confidence in the predictions of the first classifier C_1 . Therefore, the logits from both classifiers are combined linearly to obtain the final probabilistic prediction:

$$\mathbf{p}(I) = \sigma \left(\sigma^{-1}(\mathbf{p}^{(1)}(I)) + \sigma^{-1}(\mathbf{p}^{(2)}(I)) \right) \text{ ,} \quad (10)$$

where σ^{-1} is the logit function.

3.6.3. Training

The multi-label classifier thus defined is trained to minimize the binary cross-entropy \mathcal{L} between network predictions $p_n(I)$ and ground truth labels $\lambda_n(I)$, $n = 1..N$:

$$\mathcal{L} = -\frac{1}{N} \sum_I \sum_{n=1}^N \left[\lambda_n(I) \log(p_n(I)) + (1 - \lambda_n(I)) \log(1 - p_n(I)) \right] \text{ .} \quad (11)$$

Thanks to this loss function, the ordered nature of DR severity grades is taken into account (see Section 3.1). If a prediction is wrong by one severity level, then only one binary classifier will have an incorrect prediction. Whereas if a prediction is wrong by $n > 1$ severity levels, then n binary classifiers will have an incorrect prediction, and will therefore impact the global loss \mathcal{L} more.

We hypothesize that B-scan selection is most relevant when the first classifier C_1 is already well trained. Therefore, two training scenarios are investigated:

Two-step training: C_1 is trained alone until convergence. Then, its parameters are frozen, and C_2 is trained until convergence.

One-step training: C_1 and C_2 are trained jointly until convergence.

This concludes the presentation of the proposed framework.

4. Experiments

This framework is now evaluated for the task of automated DR severity assessment, according to the ICDR scale [5], using OCTA.

4.1. Dataset

In this study, we used OCTA images from the ‘‘Évaluation Inteligente de la Rétinopathie diabétique’’ (EviRed) project,² which comprises data collected between 2018 and 2022 from 14 hospitals and recruitment centers in France.

The PLEX Elite 9000, with a scanning frequency of 200 kHz, was employed to capture Swept-Source (SS) OCTA images. The ocular data in the EviRed dataset typically include two acquisition types: high-resolution $6 \times 3 \times 6 \text{ mm}^3$ SS-OCTA ($500 \times 1536 \times 500$ voxels) centered on the macula and lower-resolution $15 \times 6 \times 15 \text{ mm}^3$ ultra-widefield UWF-SS-OCTA ($834 \times 3072 \times 834$ voxels). Each OCTA volume contains 2-D en-face localizer, structural, and flow information. The EviRed dataset encompasses 811 eyes from 477 patients and is divided into training, validation, and test sets. It should be noted that, for a few eyes,

² <https://evired.org/>

Table 1

Dataset distribution among training, validation, and test sets for PLEX Elite $6 \times 6 \text{ mm}^2$ and PLEX Elite $15 \times 15 \text{ mm}^2$.

	PLEX Elite $6 \times 6 \text{ mm}^2$			PLEX Elite $15 \times 15 \text{ mm}^2$		
	Patients	Eyes	%	Patients	Eyes	%
Train	333	625	70.0	318	584	70.2
Val	82	159	17.7	78	147	17.7
Test	57	110	12.3	57	100	12.1

we have only $6 \times 6 \text{ mm}^2$ or $15 \times 15 \text{ mm}^2$ acquisitions. The distribution of patients and images in each set is presented in Table 1.

The partitioning of the EviRed dataset into training, validation, and test sets followed a systematic approach to ensure a robust evaluation of the models. The process was guided by the following criteria:

1. Patient independence: To minimize the risk of data leakage, each patient's data was assigned to only one of the sets (training, validation, or test). This approach prevents the model from learning any patient-specific characteristics that could lead to overfitting or an inflated performance metric.
2. Balanced distribution of disease severity: The dataset was divided in such a way that each set contained a similar proportion of images representing various stages of DR. This balanced distribution ensures that the model is exposed to a wide range of severity levels during the training process and provides a more accurate assessment of its performance during validation and testing.
3. Stratified sampling: To maintain consistency in the demographic characteristics and other factors, stratified sampling was employed when dividing the dataset. This approach not only ensures that each set contains a representative sample of the entire dataset but also respects the original class distribution in each split. By mirroring the class distribution of the whole dataset within each subset, we further reduce the risk of biased performance evaluation. Hence, we achieve a balanced representation of demographic characteristics and other factors across all subsets. This method gives us the confidence that the inferences drawn from our study will be robust and reliable.

Because $6 \times 6 \text{ mm}^2$ and $15 \times 15 \text{ mm}^2$ acquisitions have different sizes and resolutions, distinct models were built for these two types of acquisitions.

4.2. Implementation

The proposed framework was implemented using PyTorch³ Ignite for training and inference, MONAI⁴ for 3-D data handling, PyTorch Torchvision for differentiable data augmentation, PyTorch Image Models (timm)⁵ for 2-D image classification and Captum⁶ for attribution methods. Experiments were performed using two NVIDIA V100 GPUs (with 32 Gb of RAM each). One of the GPUs was dedicated to 3-D→2-D projection (IT), and the other one was dedicated to 2-D image classification (C_1 and C_2).

4.3. Hyperparameter optimization

Various hyperparameters need to be set:

- Architecture parameters: the use of C_1 alone or the joint use of C_1 and C_2 (see Section 3.1).

- Preprocessing parameters: the depth Y_0 at which the ILM is aligned and the depth Y_1 under which the volume is cropped (see Section 3.2).
- 3-D→2-D projection parameters: the number Φ of filters per convolutional layer in the first block and the use, or not, of skip-connections within blocks (see Section 3.3).
- Data augmentation parameters: the range of values for random affine transformations ϵ and ϵ' (see Sections 3.4 and 3.6.1).
- Ensemble parameters: the list of K classification backbones in C_1 and C_2 , and the use of logit or probability averaging (see Sections 3.4 and 3.6.1).
- B-scan selection parameters: the attribution method used (see Section 3.5).
- Training parameters: the general training parameters (optimizer, learning rate, etc.), training from scratch or the use of ImageNet pre-trained weights, and the problem-specific training schedule (one-step or two-steps — see Section 3.6.3).

The preprocessing and data augmentation parameters were set empirically through visual inspection: $Y_0 = 32$ voxels, $Y_1 = 224$ voxels, random rotation in the range $[-10; +10]$ degrees, random translation in the range $[-10; +10]$ percent of X and Y_1 or Z , random scaling in the range $[90; 110]$ percent of X and Y_1 or Z . The following parameters were restricted due to GPU memory limitations and computation time considerations:

- The number Φ of filters was limited to $\Phi \leq 32$ for $6 \times 6 \text{ mm}^2$ acquisitions, $\Phi \leq 16$ for $15 \times 15 \text{ mm}^2$ acquisitions.
- Integrated Gradients and perturbation-based attribution methods could not be used during training. The following attribution methods were investigated: Saliency, Deconvolution, Guided Backprop, and DeepLIFT.

The classifier backbones were chosen among CNNs: one advantage of most CNNs is that they can be applied to images of any size (in our case: 500×500 or 834×834 pixels) without adaptation. Two ensembles of classifiers were considered, with the following pre-trained weights (unless training from scratch is experimented):

Ensemble 1: The first set of classifiers was chosen from a single family, namely EfficientNet [63]. This family is popular for its good trade-off between complexity and performance. The $K = 4$ simplest networks were selected: namely EfficientNet-{B0, B1, B2, B3}, pre-trained on ImageNet-1K. This first ensemble was intended for quick experiments, hyper-parameter optimization, etc.

Ensemble 2: The second set of classifiers was chosen among the best-performing CNN families in ImageNet classification benchmarks.⁷ The set includes ConvNeXt Base, pre-trained on ImageNet-21K [64], ImageNet-V2 Small, pre-trained on ImageNet-21K [65] and SE-ResNet-152D, pre-trained on ImageNet-1K [66]. The set was limited to $K = 3$ due to GPU memory limitations.

The values of the other hyperparameters were then optimized using ensemble 1 for $6 \times 6 \text{ mm}^2$ acquisitions. Optimization relied on a Receiver Operating Characteristic (ROC) analysis in the validation set: hyperparameter values maximizing the Area Under the ROC Curve (AUC), averaged over the $N = 4$ binary classification tasks, were selected. The following hyperparameters were obtained:

- Full architecture (C_1 and C_2).

³ <https://pytorch.org/>

⁴ <https://monai.io/>

⁵ <https://github.com/fastai/timmdocs/>

⁶ <https://captum.ai/>

⁷ https://github.com/kentaroy47/timm_speed_benchmark — page accessed on January 2023.

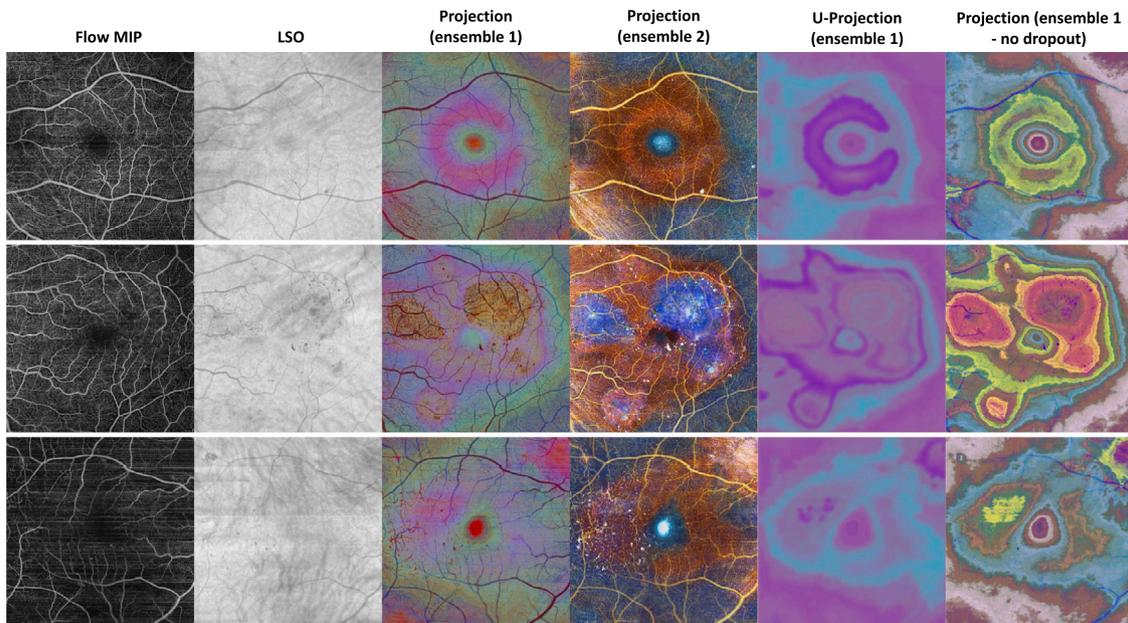


Fig. 6. Examples of projections for 6×6 mm² acquisitions from the test set. The projector has $\Phi = 8$ filters in the first block and the classifiers are EfficientNet-{B0, B1, B2, B3} for ensemble 1 and {ConvNeXt-base, Efficient-v2, SEResNet-152} for ensemble 2. The first row represents an eye graded as moderate NPDR. The second row represents an eye graded as moderate NPDR with macular edema. The third row represents an eye graded as proliferative DR.

- $\Phi = 8$ filters in the first block (and therefore $2\Phi = 16$ filters in the second and $2^2\Phi = 32$ filters in the third), without skip-connections.
- Logit averaging.
- Guided Backprop attribution method.
- Adam optimizer with an initial learning rate of 10^{-3} and an exponential learning rate scheduler (99% multiplicative decay at every epoch), using pre-trained weights.
- Two-step training.

4.4. Analysis of 3-D→2-D projections

Examples of 3-D→2-D projections obtained for 6×6 mm² acquisitions using both ensembles of classifiers are presented in Fig. 6. They are compared with projections obtained using the U-Net-like projection (or U-Projection) proposed by Lachinov et al. [42] for 3-D→2-D segmentation, and also with projections obtained without model dropout. Examples of 3-D→2-D projections obtained for 15×15 mm² acquisitions of the same eyes are presented in Fig. 7. Next, we present in Fig. 9 examples of A-scan-level attributions (see Eq. (5)) obtained for C_1 with various attribution methods. Attributions were obtained for 6×6 mm² acquisitions using ensemble 2. This figure illustrates that the Guided Backprop, DeepLIFT, and Integrated Gradients methods produce rather similar results. The Saliency and Deconvolution methods produce more noisy results.

4.5. Performance evaluation

Classification performance achieved in the test set for 6×6 mm² and 15×15 mm² acquisitions is presented in Table 2 (a). Results are reported for both ensembles and all $N = 4$ classification tasks. For a given ensemble and a given classification task, three AUC values are reported: the one obtained for the full classifier (predictions $\mathbf{p}(I)$) and the ones obtained for each branch C_1 and C_2 of the classifier separately (predictions $\mathbf{p}^{(1)}(I)$ and $\mathbf{p}^{(2)}(I)$, respectively). ROC curves for the full classifiers are reported in Fig. 10.

To show the relevance of the second branch (classifier C_2), we compared the AUC values obtained using (1) C_1 alone, (2) C_2 alone, or (3) C_1 and C_2 jointly. For each of these 3 scenarios, a set of 16 AUC

values is available in Table 2 (a) (2 acquisitions \times 2 CNN ensembles \times 4 decisions). Wilcoxon signed-rank tests were performed to compare two scenarios by confronting the corresponding 16 pairs of AUC values: results are reported in Table 2 (b).

4.6. Ablation study

Additional experiments were performed with 6×6 mm² acquisitions, using ensemble 1, to study the impact of the main hyperparameters. Results of experiments investigating the first classification branch C_1 only are reported in Table 3, and those investigating the second branch are reported in Table 4. Please note that results are reported on the test, not on the validation set (used for hyperparameter optimization), so the ranking of solutions may contradict the optimal hyperparameter values listed in Section 4.3 (associated with test results reported in Table 2).

In this ablation study, we only considered one acquisition and one CNN ensemble. With only four pairs of AUC values to compare, the previous statistical test is no longer suitable. Instead of comparing AUC values, we thus compared ROC curves directly: Delong tests were used for that purpose. Delong tests are designed for comparing two curves, not two sets of curves: each set of curves was thus micro-averaged [67] prior to Delong testing.

4.7. Comparison with a 3-D baseline

The proposed DISCOVER framework was also compared to a 3-D baseline model in terms of classification performance and inference times. The baseline model is a 3-D CNN processing the 3-channel 3-D preprocessed acquisition I , obtained as presented in Section 3.2. We ensured that the same splits were used for training, validation, and testing. Various backbones and hyperparameters were evaluated. The selection process, consistent with the one applied to our proposed method, involved choosing the optimal model based on average AUCs and the best checkpoint for each severity cut-off on the validation set. The most favorable results for the baseline were achieved using a 3-D ResNet50 [60,68] for both 6×6 mm² and 15×15 mm² acquisitions. This baseline and the proposed framework are compared in Table 5; For the proposed framework, ensemble 1 was used for 6×6 mm²

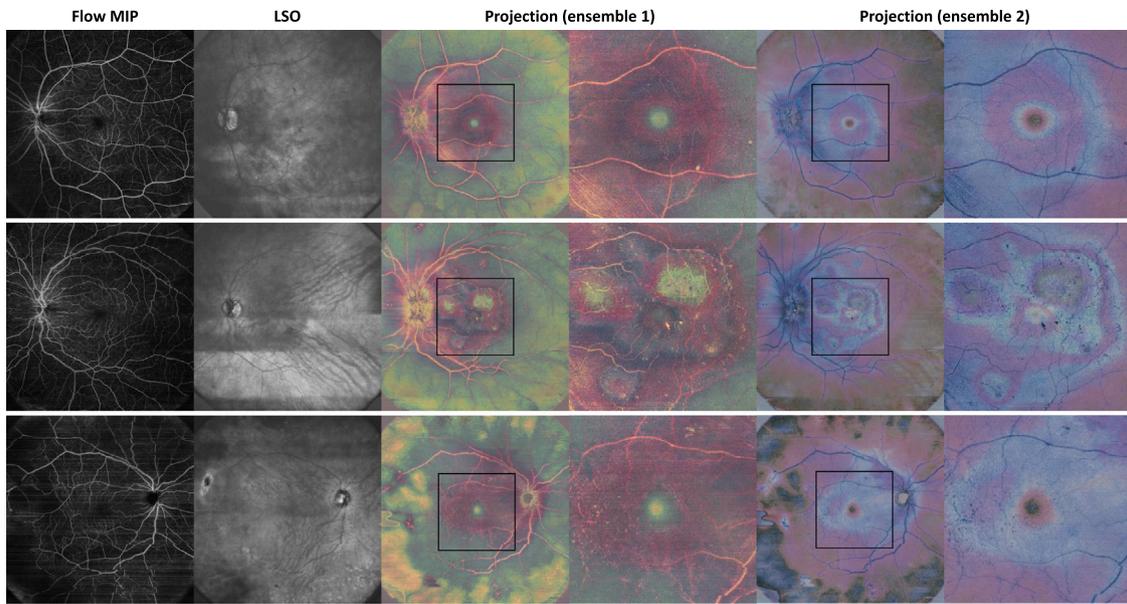


Fig. 7. Examples of projections for $15 \times 15 \text{ mm}^2$ acquisitions using the same hyperparameters as in Fig. 6, for the same eyes (the models are retrained for the new acquisition size). The squares indicate the zoomed areas, which were imaged in the $6 \times 6 \text{ mm}^2$ acquisitions.

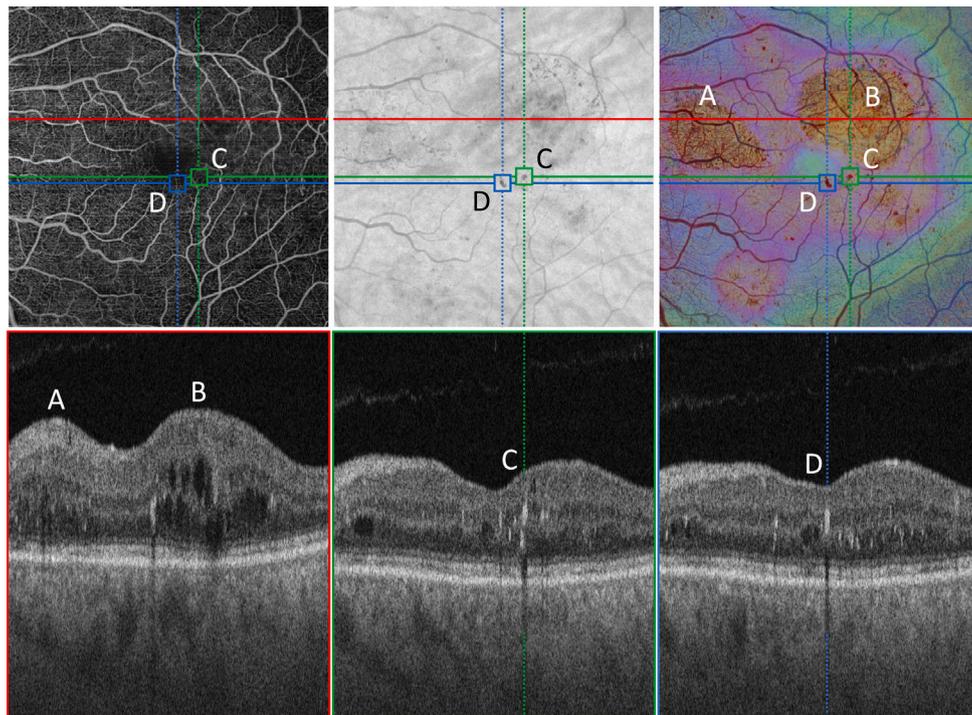


Fig. 8. Retinal lesions highlighted in the 2-D projections (for the second eye of Figs. 6 and 7). The first row shows, from left to right, the flow MIP, the LSO, and the projection (ensemble 1). The second row shows B-scans. Macular edemas (A and B) are visible in the first B-scan (horizontal red line). Large microaneurysms (C and D) are visible in the following B-scans (horizontal green and blue lines, respectively).

acquisitions and ensemble 2 was used for $15 \times 15 \text{ mm}^2$ acquisitions, as these are the best ensembles on the validation sets. Classification performances are compared using both Delong tests and a Wilcoxon signed-rank test.

5. Discussion

We have presented a general 3-D image classification framework, which combines a trainable 3-D→2-D en-face projection step, followed by a 2-D en-face image classification step. It is further complemented by

an auxiliary branch that extracts key 2-D cross-sectional slices (B-scans) and classifies them. The main purpose is to summarize 3-D information by complementary 2-D images (en-face and cross-sectional), for improved interpretability. This novel framework was applied to automated DR severity assessment using 3-D Optical Coherence Tomography Angiography (OCTA) acquisitions. This work aligns with our previous works on interpretable or explainable DR severity assessment [7,69]. But unlike these previous works, which operated on 2-D Color Fundus Photographs (CFP), this work is applied to 3-D OCTA acquisitions. The additional dimension complicates visual feedback to

Table 2

Performance of the proposed framework in terms of area under the ROC curve (AUC) in the test set (a). Impact of the classification branches on performance.

(a) Performance							(b) Comparison		
Acquisition	Backbones	Classification branches	\geq mild NPDR	\geq moderate NPDR	\geq severe NPDR	\geq PDR	Classification Branches		Wilcoxon test (p)
$6 \times 6 \text{ mm}^2$	EfficientNet-{B0, B1, B2, B3}	C_1	0.946	0.913	0.809	0.815	C_1	C_2	0.00031*
		C_2	0.798	0.764	0.675	0.589			
		$\{C_1, C_2\}$	0.958	0.920	0.808	0.821			
	{ConvNeXt-base, Efficient-v2, SEResNet-152}	C_1	0.935	0.849	0.768	0.749			
		C_2	0.824	0.737	0.734	0.816			
		$\{C_1, C_2\}$	0.951	0.862	0.812	0.796			
$15 \times 15 \text{ mm}^2$	EfficientNet-{B0, B1, B2, B3}	C_1	0.918	0.815	0.767	0.948			
		C_2	0.766	0.726	0.725	0.848			
		$\{C_1, C_2\}$	0.925	0.822	0.782	0.952			
	{ConvNeXt-base, Efficient-v2, SEResNet-152}	C_1	0.912	0.800	0.880	0.947			
		C_2	0.867	0.774	0.658	0.676			
		$\{C_1, C_2\}$	0.941	0.822	0.876	0.957			

* Significant difference, according to a Wilcoxon signed-rank test on the paired AUC values ($p < 0.05$).**Table 3**Influence of various hyperparameters on the area under the ROC curve (AUC) in the test set — analysis of the first classification branch C_1 only (no B-scan selection). Experiments were performed for $6 \times 6 \text{ mm}^2$ acquisitions and the EfficientNet-{B0, B1, B2, B3} backbones (ensemble 1). The first line corresponds to the optimal values (based on experiments on the validation set). Investigated parameters are in bold. U-Projection denotes the U-Net-like projection proposed by Lachinov et al. [42].

Φ	Number of filters	U-Projection?	Model dropout?	From scratch?	Probability averaging?	\geq mild NPDR	\geq moderate NPDR	\geq severe NPDR	\geq PDR	Delong test (p)
8	{8,16,32}	no	yes	no	no	0.946	0.913	0.809	0.815	
4	{4,8,16}	no	yes	no	no	0.865	0.799	0.755	0.860	0.0401*
16	{16,32,64}	no	yes	no	no	0.942	0.893	0.798	0.770	0.1416
32	{32,64,128}	no	yes	no	no	0.937	0.913	0.836	0.839	0.9100
8	{8,16,32}	yes	yes	no	no	0.851	0.774	0.754	0.769	0.0266*
8	{8,16,32}	no	no	no	no	0.856	0.834	0.718	0.731	0.0004*
8	{8,16,32}	no	yes	yes	no	0.804	0.752	0.725	0.854	< 0.0001*
8	{8,16,32}	no	yes	no	yes	0.921	0.867	0.836	0.817	0.0072*

* Significant difference with the reference (first line), according to a Delong test on the micro-averaged ROC curves ($p < 0.05$).**Table 4**Influence of various hyperparameters on the AUC in the test set — analysis of the second classification branch C_2 (jointly with C_1). Experiments were performed for $6 \times 6 \text{ mm}^2$ acquisitions and the EfficientNet-{B0, B1, B2, B3} backbones (ensemble 1). The first line corresponds to the optimal values (based on experiments on the validation set). Investigated parameters are in bold. Argmax B-scan selection refers to a deterministic selection of B-scans during training, like during inference (see Eq. (7)).

One-step?	B-scan selection	Attribution method	\geq mild NPDR	\geq moderate NPDR	\geq severe NPDR	\geq PDR	Delong test (p)
no	random	Guided backprop	0.958	0.920	0.808	0.821	
yes	random	Guided backprop	0.930	0.855	0.809	0.846	0.0004*
no	argmax	Guided backprop	0.945	0.914	0.813	0.813	0.2411
no	random	Saliency	0.954	0.920	0.808	0.815	0.8705
no	random	Deconvolution	0.953	0.918	0.810	0.816	0.8541
no	random	DeepLIFT	0.955	0.918	0.811	0.825	0.9723

* Significant difference with the reference (first line), according to a Delong test on the micro-averaged ROC curves ($p < 0.05$).

human readers: the relevant information needs to be summarized so that it can be displayed on a 2-D screen or printed in a report. The proposed framework mimics and generalizes how ophthalmologists analyze OCTA acquisitions: en-face projections are often used to inspect the blood flow; cross-sectional views are often used to inspect structural anomalies.

Besides improved interpretability, by design, we show that this framework guarantees improved classification performance, compared to direct 3-D image classification. This can be explained by the large volume of information in OCTA acquisitions, which complex 3-D neural architectures cannot process efficiently with limited dataset sizes and limited GPU memory capacities. Summarizing also helps for this

purpose. In particular, it allows access to a large collection of highly efficient 2-D neural architectures, with ImageNet pre-trained weights.

5.1. Analysis of 2-D projections

Two types of OCTA acquisitions were analyzed in this study: high-resolution acquisitions centered on the macula ($6 \times 6 \text{ mm}^2$) and lower-resolution ultra-widefield acquisitions ($15 \times 15 \text{ mm}^2$). Fig. 6 shows that normal and pathological retinal features are highly visible in en-face projections for $6 \times 6 \text{ mm}^2$ acquisitions (columns 3 and 4). This figure highlights the benefit of processing each A-scan independently

Table 5

Comparison between the proposed DISCOVER framework and a 3-D CNN baseline in terms of classification performance in the test set (AUC) and in terms of inference times. Inference times are given in seconds per volume, excluding preprocessing (see Section 3.2), which is common for both frameworks. For information, inference times are 0.130 seconds/volume for ensemble 2 on $6 \times 6 \text{ mm}^2$ acquisitions and 0.187 seconds/volume for ensemble 1 on $15 \times 15 \text{ mm}^2$ acquisitions.

Acquisition	Framework	Criterion	AUC	Delong test (p)	Wilcoxon test (p)	Inference times (seconds/volume)
$6 \times 6 \text{ mm}^2$	Baseline (3-D CNN)	\geq mild NPDR	0.865	0.1164	0.0078*	2.038
		\geq moderate NPDR	0.809			
		\geq severe NPDR	0.764			
		\geq PDR	0.753			
	Proposed (ensemble 1)	\geq mild NPDR	0.958			
		\geq moderate NPDR	0.920			
		\geq severe NPDR	0.808			
		\geq PDR	0.821			
$15 \times 15 \text{ mm}^2$	Baseline (3-D CNN)	\geq mild NPDR	0.820	0.0473*		2.876
		\geq moderate NPDR	0.786			
		\geq severe NPDR	0.765			
		\geq PDR	0.886			
	Proposed (ensemble 2)	\geq mild NPDR	0.941			
		\geq moderate NPDR	0.822			
		\geq severe NPDR	0.876			
		\geq PDR	0.957			

* Significant difference between the baseline and the proposed approach, according to a Delong test on the micro-averaged ROC curves or according to a Wilcoxon signed-rank test on the paired AUC values ($p < 0.05$).

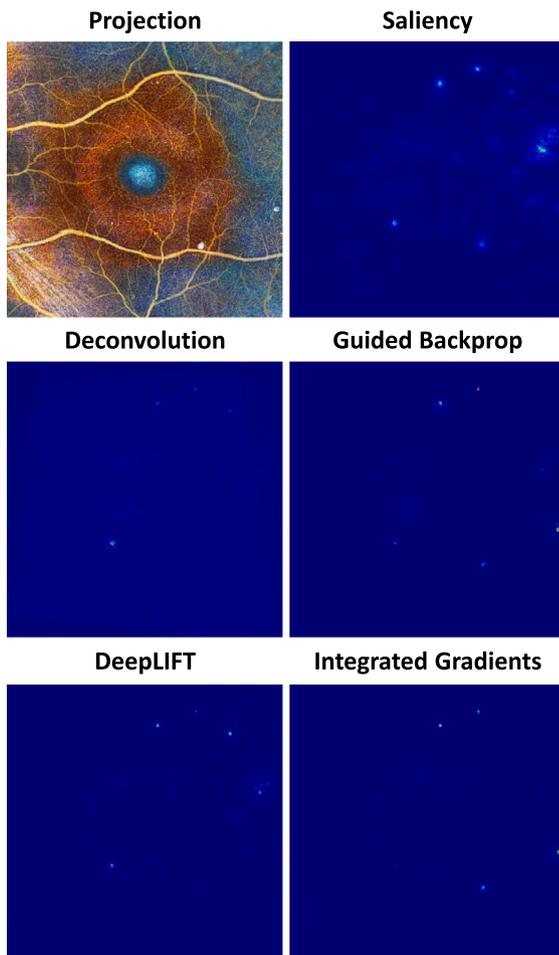


Fig. 9. Comparison of a few attribution methods described in Section 2.2, computed for a $6 \times 6 \text{ mm}^2$ acquisition ($\Phi = 8$, ensemble 2: {ConvNeXt-base, Efficient-v2, SEResNet-152}). Each pixel intensity represents the sum of attributions computed for the $N = 4$ model outputs (one output per DR severity cutoff), see Eq. (5). The ‘jet’ color map, ranging from blue (0) to red (255), is used.

through 1-D convolutions: all details are lost through U-Net-like projections (U-Projection in column 5). The last column also advocates the joint training of multiple 2-D classifiers through the proposed “model dropout” mechanism: by training a single classifier (or static ensemble of classifiers), relevant details may be lost, as the unique classifier may focus on a subset of discriminant features and let the 3-D→2-D projector ignore the others, while still obtaining good classification performance (see Table 3).

Through comparison of Figs. 6 and 7, it appears that $15 \times 15 \text{ mm}^2$ en-face projections offer a reduced level of details. This makes sense, given the reduced resolution of input images (along all three axes). It is also possible that the 3-D→2-D projector was not able to capture details equally well: the problem is not just a reduced resolution of normal and pathological retinal features, but also a reduced contrast between these features and the background. However, Fig. 7 suggests that large and peripheral pathological features are well captured in $15 \times 15 \text{ mm}^2$ acquisitions, explaining that advanced DR stages are detected well in those acquisitions (see Table 2).

By design, the proposed 3-D→2-D projection operator Π , which processes each A-scan independently, does not take the context of these A-scans into account. However, the classification branch C_1 does: by training $C_1 \circ \Pi$ jointly, Π can be trained to extract information allowing localization. For instance, we can see that color in Π varies with the retinal thickness (see Figs. 6 and 7): this could be useful to capture pathological features (like macular edemas), but certainly also to localize them relatively to the normal retinal features. Indeed, the clinicians collaborating in this study affirm that this method markedly enhances the visibility of retinal lesions, as shown in Fig. 8. As it is apparent, pathological features are well preserved and highlighted, while additional details from the B-scans complement the classifier, reinforcing the need for automated B-scan selection.

To validate the usefulness of proposed 3-D→2-D projections for decision support, the next step will be to compare the classification performance of clinicians when they use these projections for decision support versus when they do not. Although it is clear that the relevant abnormal structures stand out well in these images, the color-coding of retinal features derived from these projections may be disturbing for some clinicians. For instance, Fig. 6 leads to unusual nebula-like images using ensemble 2 (column 4). For such a validation study, it may be useful to reorder color channels before human inspection, for instance, to produce more conventional-looking images. However, these considerations are out of the scope of an artificial intelligence paper.

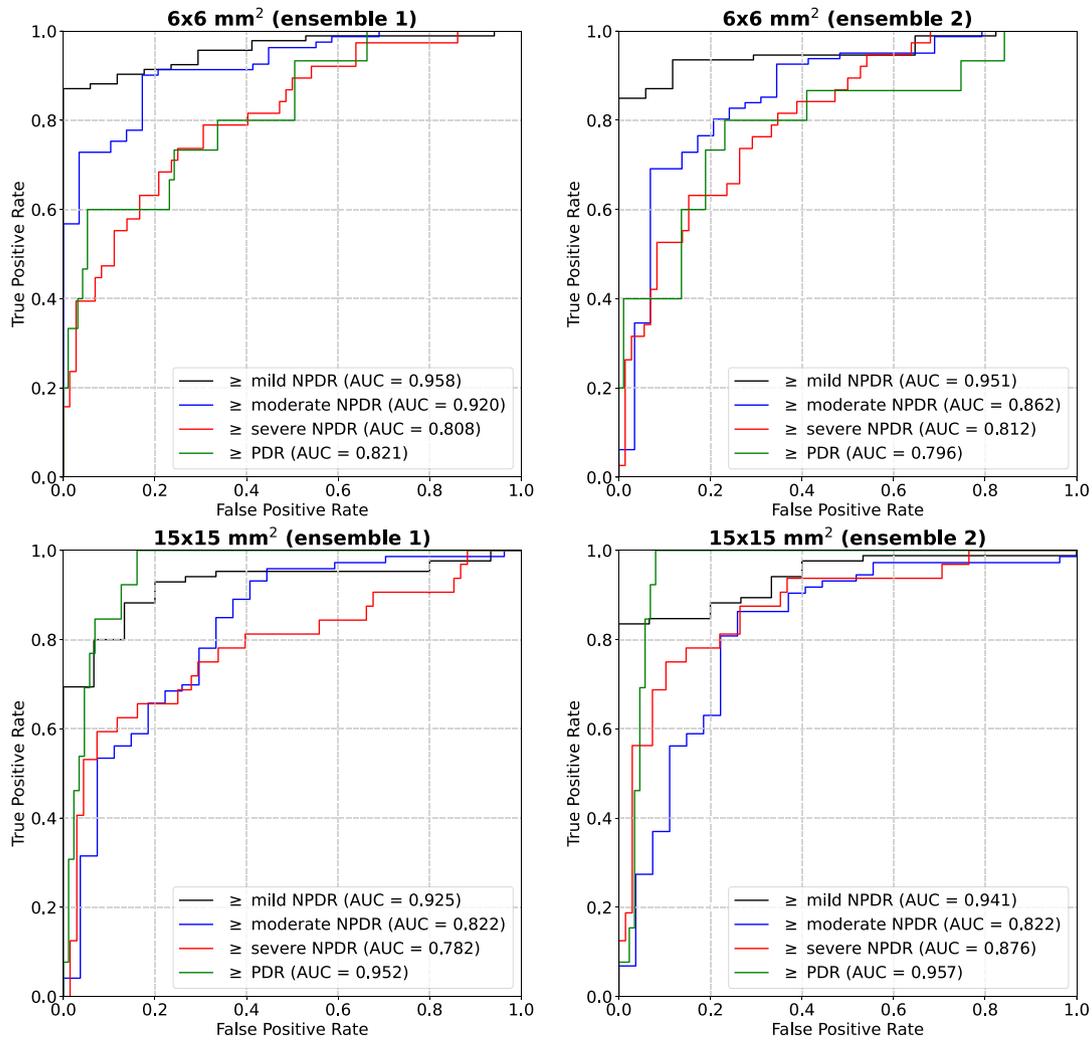


Fig. 10. Receiver Operating Characteristic of the proposed system in the test set using EfficientNet-{B0, B1, B2, B3} (ensemble 1) or {ConvNeXt-base, Efficient-v2, SEResNet-152} (ensemble 2) for $6 \times 6 \text{ mm}^2$ and $15 \times 15 \text{ mm}^2$ acquisitions.

One important property illustrated in Figs. 6 and 7 is that the same features (normal anatomical features like blood vessels, or pathological features like DR lesions) have consistent appearances in 2-D projections across patients (i.e., across lines of these figures). However, the appearance of these features clearly depends on the architecture of the projection operator (i.e., their appearance varies across columns of these figures). Similarly, their appearance depends on the weights trained for the projection operator. If one decides to retrain or fine-tune the proposed framework on a larger dataset, it may be beneficial to freeze the projection weights and only retrain the classification ensembles. This will ensure that the proposed 2-D visualization remains standardized.

5.2. Detailed analysis of the framework

Table 2 demonstrates that the first classification branch (C_1), which analyzes the 3-D \rightarrow 2-D projection, contributes the most to the final classification. This suggests that the proposed 3-D \rightarrow 2-D projections contain most of the discriminant information contained in 3-D OCTA acquisitions for the target classification task, which is good news for interpretability purposes: one single 2-D image conveys most of the relevant information. However, we have shown that combining the two branches leads to a significant increase in classification performance

($p = 0.00015$, see Table 2). The superiority of branch 1 is particularly true in the two-step training schedule, which was adopted: at first, branch 1 is trained independently, to maximize classification performance; branch 2 (C_2), which analyzes selected B-scans, is only trained afterward, to be complementary to the frozen branch 1. In other words, branch 2 was not trained to be discriminant on its own.

An ablation study was performed to analyze the benefits of most methodological choices in the proposed design. Table 3 suggests that classification performance is impacted by the complexity of the projector, driven by parameter Φ , but no significant difference was found with more than $\Phi = 8$. Next, the U-Net-like projector and the absence of model dropout, which clearly affect the quality of en-face projections, also decrease classification performance significantly. The use of ImageNet pre-trained weights and logit averaging also proved beneficial. Table 4 suggests that the performance of the second classification branch is little dependent on the attribution method used. As for the random selection of B-scans during training, it seems beneficial, but no significant difference was observed. The most influential parameter in Table 4 is the choice between one-step and two-step training ($p = 0.0004$). The two-step approach has the advantage of more stable training; in particular, one instance of training divergence was observed with the one-step approach. This may explain the increased performance with two steps. On the downside, two-step training implies longer training times (by a factor of two, approximately). Therefore,

Table 6

Comparison of DR severity assessment methods using OCTA. The following abbreviations are used: rDR = referable DR; vtDR = vision-threatening DR; PDR = proliferative DR; NPDR = non-proliferative DR; acc = accuracy; AUC = area under the ROC curve.

Reference	Dataset	Acquisition	Method	Criterion	Performance
Heisler et al. [30]	380 eyes (463 scans)	$3 \times 3 \text{ mm}^2$	2-D CNN	\geq rDR	acc = 0.92
Le et al. [27]	177 eyes (177 scans)	$6 \times 6 \text{ mm}^2$	2-D CNN	\geq mild NPDR	AUC = 0.97
Andreeva et al. [28]	113 patients (320 scans)	$3 \times 3 \text{ mm}^2$	2-D CNN	\geq mild NPDR	AUC = 0.83
Lo et al. [20]	700 eyes (700 scans)	$3 \times 3 \text{ mm}^2$	2-D CNN	\geq rDR	AUC = 0.954-0.960
Ryu et al. [29]	360 scans	$3 \times 3 \text{ mm}^2$	2-D CNN	\geq mild NPDR	AUC = 0.928-0.960
		$6 \times 6 \text{ mm}^2$		\geq rDR	AUC = 0.940
Ryu et al. [26]	918 eyes	$3 \times 3 \text{ mm}^2$	radiomics 2-D CNN	DR staging	acc = 0.574 acc = 0.684
	917 eyes	$6 \times 6 \text{ mm}^2$	radiomics 2-D CNN		acc = 0.531 acc = 0.728
Yasser et al. [31]	91 patients	$3 \times 3 \text{ mm}^2$	2-D CNN ^a	\geq mild NPDR	acc = 0.944
Zang et al. [32]	355 patients (355 scans)	$3 \times 3 \text{ mm}^2$	3-D CNN	\geq rDR \geq vtDR	AUC = 0.96 AUC = 0.92
Li et al. [23]	300 scans	$3 \times 3 \text{ mm}^2$ or $6 \times 6 \text{ mm}^2$	2-D CNN ^a	\geq mild NPDR	AUC = 0.92
Li et al. [33]	64 patients (151 scans)	$6 \times 6 \text{ mm}^2$	3-D CNN	\geq PDR	AUC = 0.911
Khalili Pour et al. [21]	78 patients/148 eyes	$6 \times 6 \text{ mm}^2$	radiomics	PDR v.s. NPDR	acc = 0.85
Li et al. [34]	432 patients/801 eyes	$6 \times 6 \text{ mm}^2$ and $15 \times 15 \text{ mm}^2$	3-D CNN	\geq mild NPDR	AUC = 0.912
				\geq moderate NPDR	AUC = 0.829
Proposed	472 patients/894 eyes	$6 \times 6 \text{ mm}^2$	3-D \rightarrow 2-D projection & 2-D CNN	\geq severe NPDR	AUC = 0.812
				\geq PDR	AUC = 0.900
				\geq mild NPDR	AUC = 0.958
				\geq moderate NPDR	AUC = 0.920
				\geq severe NPDR	AUC = 0.808
453 patients/831 eyes	$15 \times 15 \text{ mm}^2$	\geq PDR	AUC = 0.821		
		\geq mild NPDR	AUC = 0.941		
		\geq moderate NPDR	AUC = 0.822		
				\geq severe NPDR	AUC = 0.876
				\geq PDR	AUC = 0.957

^a Feature segmentations are used as inputs.

we recommend investigating both approaches when replicating these results.

5.3. Comparison with previous algorithms

Although DR severity assessment using OCTA is a recent topic, a few classification results have already been published, using a variety of DR severity cutoffs; those results are reported in Table 6. It should be noted that they were generally obtained on small datasets and that data collection often included a data selection process based on image quality. For instance, Ryu et al. [29] imaged 496 eyes, but only 360 scans were retained for further analysis, indicating a rejection rate of 27% (or more, if some patients were imaged more than once). The number of images rejected for quality reasons is not always mentioned: for instance, Le et al. [27] only indicated a quality threshold. These factors make comparisons between algorithms challenging. However, it appears that the proposed framework allows similar or better classification results than previously published algorithms, regardless of the type of analysis (2-D CNN, 3-D CNN, or radiomics). Two tasks are particularly well addressed by our framework: \geq mild NPDR detection (AUC = 0.958) and \geq PDR detection (AUC = 0.957). It should be noted that, besides results from our team [34], on a subset of the presented dataset, this is the first publication about ultra-widefield OCTA acquisitions ($15 \times 15 \text{ mm}^2$) on the topic. This type of acquisition seems promising for detecting advanced DR stages.

In order to thoroughly assess the performance of our proposed method, we have also conducted a direct comparison with a baseline 3-D CNN model, on the same test set (see Section 4.7). We have demonstrated the superiority of the proposed DISCOVER framework

over that baseline, both in terms of classification performance and in terms of inference times (see Table 5). In detail, when comparing ROC curves using Delong tests, a significant difference was found for $15 \times 15 \text{ mm}^2$ acquisitions, but not for $6 \times 6 \text{ mm}^2$ acquisitions ($p = 0.1164$). However, when looking solely at the AUC, using a Wilcoxon signed-rank test, a significant difference was found overall in favor of DISCOVER. As for inference times, they are up to 20 times faster using DISCOVER. We believe that our proposed pipeline, which employs a 3-D \rightarrow 2-D summarization in conjunction with a 2-D classification, surpasses the performance of direct 3-D classification models, such as the baseline, for several reasons:

1. By incorporating the strengths of both en-face projections and cross-sectional slices, our method captures more pertinent information from the OCTA volumes.
2. Our method features a lighter architecture compared to 3-D neural networks, resulting in a reduced propensity for overfitting and increased adaptability to smaller datasets. The utilization of pretrained 2-D backbones further bolsters this advantage.
3. The end-to-end training approach allows our pipeline to acquire more discriminative features, thereby enhancing DR severity assessment.

5.4. Limitations and future works

This study has a few limitations. First, due to long training times and high resource consumption, no cross-validation or advanced hyperparameter optimization strategy (like Bayesian optimization) was adopted, so chances are that hyperparameter optimization is suboptimal. Thus, in the test set, we observe a better performance with $\Phi = 32$

initial filters for two classification labels (\geq severe NPDR and \geq PDR), compared to the hyperparameter value $\Phi = 8$ selected in the validation set (see Table 3). A second limitation is that the dataset is limited in size, which impacts both training and performance evaluation. A third limitation is the use, for training and evaluation, of a DR severity scale known to be suboptimal for DR management due to its limited predictive power.

Addressing the latter two points is the purpose of the Evired project⁸: we are collecting longitudinal data from thousands of diabetic patients to define a more predictive DR severity scale. The end goal will be to predict the advent of two DR complications in the following year: proliferative DR and DME. Therefore, ultimately, the proposed framework will be trained to solve a novel 2-label prognosis task ($N = 2$). The interpretability of the proposed framework will be particularly useful in discovering which features are found to be predictive.

The use of Transformers could also be investigated in future works. CNNs were used in this paper for convenience: they can generally be applied to images of any size without adaptation. Besides, Transformers often require larger datasets, adding to their computational demands. Given our dataset and computational constraints, the use of CNNs thus proved to be a more practical choice, but this may change in the future.

6. Conclusion

This work presents a novel framework for 3-D image classification, with improved interpretability and classification performance. This framework is particularly suited to the analysis of OCT/OCTA images in ophthalmology: its usefulness was demonstrated for DR severity assessment. We expect it to be useful for other 3-D medical image classification tasks in the future.

Moving forward, the prospective validation of this framework through a clinical study is an essential next step. Such a study would evaluate whether the selected 2-D projections and B-scans that form the core of our approach indeed provide meaningful and actionable information for clinicians. This validation would support our contention that the methodology not only aids in image classification but also contributes significantly to the interpretability of the results, an aspect that is of high importance for doctors involved in diagnostic processes.

Moreover, while the proposed approach has demonstrated promising results, the generalizability of our method will need further assessment. Before our framework can be integrated into routine clinical practice, it is vital to evaluate its performance on an independent population. This evaluation would ensure that the framework is robust and reliable across different populations and clinical contexts and is not limited to the specific cases and data sets upon which it has been developed and tested.

CRedit authorship contribution statement

Mostafa El Habib Daho: Methodology, Writing – original draft, Writing – review & editing, Software, Visualization. **Yihao Li:** Software, Validation, Writing – review & editing. **Rachid Zeglache:** Software, Writing – review & editing. **Pierre Deman:** Data curation, Resources, Software. **Laurent Borderie:** Data curation, Resources. **Hugang Ren:** Resources, Software. **Niranchana Mannivanan:** Resources, Software. **Capucine Lepicard:** Project administration. **Béatrice Cochener:** Data curation, Resources. **Aude Couturier:** Data curation, Resources. **Ramin Tadayoni:** Funding acquisition, Project administration. **Pierre-Henri Conze:** Writing – review & editing. **Mathieu Lamard:** Writing – review & editing. **Gwenolé Quellec:** Conceptualization, Funding acquisition, Software, Supervision, Writing – original draft, Writing – review & editing, Investigation, Methodology.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ramin Tadayoni and Gwenolé Quellec report financial support was provided by French National Research Agency. Ramin Tadayoni and Béatrice Cochener report a relationship with Carl Zeiss Meditec that includes: consulting or advisory. Pierre Deman, Employee (ADCIS, Evolucare Technologies), Laurent Borderie, Employee (Evolucare Technologies), Hugang Ren and Niranchana Mannivanan, Employees (Carl Zeiss Meditec)

Acknowledgments

This work was funded in part by the French National Research Agency under the “Investissement d’Avenir” program (ANR-18-RHUS-0008 - EVIRED project) and under the LabCom program (ANR-19-LCV2-0005 - ADMIRE project).

References

- Ogurtsova K, da Rocha Fernandes JD, Huang Y, Linnenkamp U, Guariguata L, Cho NH, et al. IDF Diabetes Atlas: Global estimates for the prevalence of diabetes for 2015 and 2040. *Diabetes Res Clin Pract* 2017;128:40–50. <http://dx.doi.org/10.1016/j.diabres.2017.03.024>.
- Yau JWY, Rogers SL, Kawasaki R, Lamoureux EL, Kowalski JW, Bek T, et al. Global prevalence and major risk factors of diabetic retinopathy. *Diabetes Care* 2012;35(3):556–64. <http://dx.doi.org/10.2337/dc11-1909>.
- David M, Norton E, Myers F. The Airlie classification of diabetic retinopathy. In: Goldberg M, Fine S, editors. *Proc symp treatment of diabetic retinopathy*. Washington, DC, USA; 1969, p. 1890.
- ETDRS Research Group. Fundus photographic risk factors for progression of diabetic retinopathy: ETDRS report number 12. *Ophthalmology* 1991;98(5, Supplement):823–33. [http://dx.doi.org/10.1016/S0161-6420\(13\)38014-2](http://dx.doi.org/10.1016/S0161-6420(13)38014-2).
- Wilkinson CP, Ferris FL, Klein RE, Lee PP, Agardh CD, Davis M, et al. Proposed international clinical diabetic retinopathy and diabetic macular edema disease severity scales. *Ophthalmology* 2003;110(9):1677–82. [http://dx.doi.org/10.1016/S0161-6420\(03\)00475-5](http://dx.doi.org/10.1016/S0161-6420(03)00475-5).
- Ting DS, Peng L, Varadarajan AV, Keane PA, Burlina PM, Chiang MF, et al. Deep learning in ophthalmology: The technical and clinical considerations. *Prog Retin Eye Res* 2019;72:100759. <http://dx.doi.org/10.1016/j.preteyeres.2019.04.003>.
- Quellec G, Al Hajj H, Lamard M, Conze P-H, Massin P, Cochener B. ExplAIn: Explanatory artificial intelligence for diabetic retinopathy diagnosis. *Med Image Anal* 2021;72:102118. <http://dx.doi.org/10.1016/j.media.2021.102118>.
- Hitzenberger CK, Trost P, Lo P-W, Zhou Q. Three-dimensional imaging of the human retina by high-speed optical coherence tomography. *Opt Express* 2003;11(21):2753–61. <http://dx.doi.org/10.1364/OE.11.002753>.
- Baumann B, Pottsaid B, Kraus MF, Liu JJ, Huang D, Hornegger J, et al. Total retinal blood flow measurement with ultrahigh speed swept source/Fourier domain OCT. *Biomed Opt Express* 2011;2(6):1539–52. <http://dx.doi.org/10.1364/BOE.2.001539>.
- Gorczyńska I, Migacz JV, Zawadzki RJ, Capps AG, Werner JS. Comparison of amplitude-decorrelation, speckle-variance and phase-variance OCT angiography methods for imaging the human retina and choroid. *Biomed Opt Express* 2016;7(3):911–42. <http://dx.doi.org/10.1364/BOE.7.000911>.
- Niederleithner M, De Sisternes L, Stino H, Sedova A, Schlegl T, Bagherinia H, et al. Ultra-widefield OCT angiography. *IEEE Trans Med Imaging* 2023;42(4):1009–20. <http://dx.doi.org/10.1109/TMI.2022.3222638>.
- Alam M, Le D, Lim J, Chan R, Yao X. Supervised machine learning based multi-task artificial intelligence classification of retinopathies. *J Clin Med* 2019;8(6):872. <http://dx.doi.org/10.3390/jcm8060872>.
- Yang D, Ran A, Nguyen T, Lin T, Chen H, Lai T, et al. Deep learning in optical coherence tomography angiography: Current progress, challenges, and future directions. *Diagnostics* 2023;13(2):326. <http://dx.doi.org/10.3390/diagnostics13020326>.
- Xu X, Yang P, Wang H, Xiao Z, Xing G, Zhang X, et al. AV-casNet: Fully automatic arteriole-venule segmentation and differentiation in OCT angiography. *IEEE Trans Med Imaging* 2023;42(2):481–92. <http://dx.doi.org/10.1109/TMI.2022.3214291>.
- Zang P, Hormel T, Hwang T, Bailey S, Huang D, Jia Y. Deep-learning-aided diagnosis of diabetic retinopathy, age-related macular degeneration, and glaucoma based on structural and angiographic OCT. *Ophthalmol Sci* 2023;3(1):1000245. <http://dx.doi.org/10.1016/j.xops.2022.1000245>.
- Yang Z, Tan T-E, Shao Y, Wong T, Li X. Classification of diabetic retinopathy: Past, present and future. *Front Endocrinol* 2022;13:1079217. <http://dx.doi.org/10.3389/fendo.2022.1079217>.

⁸ <https://evired.org/>

- [17] Vujosevic S, Cunha-Vaz J, Figueira J, Löwenstein A, Midena E, Parravano M, et al. Standardization of optical coherence tomography angiography imaging biomarkers in diabetic retinal disease. *Ophthalmic Res* 2021;64:871–87.
- [18] Sun Z, Yang D, Tang Z, Ng D, Cheung C. Optical coherence tomography angiography in diabetic retinopathy: An updated review. *Eye* 2021;35(1):149–61. <http://dx.doi.org/10.1038/s41433-020-01233-y>.
- [19] Guo Y, Hormel T, Xiong H, Wang J, Hwang T, Jia Y. Automated segmentation of retinal fluid volumes from structural and angiographic optical coherence tomography using deep learning. *Transl Vis Sci Technol* 2020;9(2):54. <http://dx.doi.org/10.1167/tvst.9.2.54>.
- [20] Lo J, Yu T, Ma D, Zang P, Owen J, Zhang Q, et al. Federated learning for microvasculature segmentation and diabetic retinopathy classification of OCT data. *Ophthal Sci* 2021;1(4):100069. <http://dx.doi.org/10.1016/j.xops.2021.100069>.
- [21] Khalili Pour E, Rezaee K, Azimi H, Mirshahvalad S, Jafari B, Fadakar K, et al. Automated machine learning-based classification of proliferative and non-proliferative diabetic retinopathy using optical coherence tomography angiography vascular density maps. *Graefes Arch Clin Exp Ophthalmol* 2023;261(2):391–9. <http://dx.doi.org/10.1007/s00417-022-05818-z>.
- [22] Guo M, Zhao M, Cheong A, Dai H, Lam A, Zhou Y. Automatic quantification of superficial foveal avascular zone in optical coherence tomography angiography implemented with deep learning. *Vis Comput Ind Biomed Art* 2019;2(1):21. <http://dx.doi.org/10.1186/s42492-019-0031-8>.
- [23] Li Q, Zhu X-R, Sun G, Zhang L, Zhu M, Tian T, et al. Diagnosing diabetic retinopathy in OCTA images based on multilevel information fusion using a deep learning framework. *Comput Math Methods Med* 2022;2022:4316507. <http://dx.doi.org/10.1155/2022/4316507>.
- [24] Vaz-Pereira S, Morais-Sarmiento T, Engelbert M. Update on optical coherence tomography and optical coherence tomography angiography imaging in proliferative diabetic retinopathy. *Diagnostics* 2021;11(10):1869. <http://dx.doi.org/10.3390/diagnostics11101869>.
- [25] Carrera-Escalé L, Benali A, Rathert A-C, Martín-Pinardel R, Bernal-Morales C, Alé-Chilet A, et al. Radiomics-based assessment of OCT angiography images for diabetic retinopathy diagnosis. *Ophthalmol Sci* 2023;3(2):100259. <http://dx.doi.org/10.1016/j.xops.2022.100259>.
- [26] Ryu G, Lee K, Park D, Kim I, Park S, Sagong M. A deep learning algorithm for classifying diabetic retinopathy using optical coherence tomography angiography. *Transl Vis Sci Technol* 2022;11(2):39. <http://dx.doi.org/10.1167/tvst.11.2.39>.
- [27] Le D, Alam M, Yao C, Lim J, Hsieh Y-T, Chan R, et al. Transfer learning for automated OCTA detection of diabetic retinopathy. *Transl Vis Sci Technol* 2020;9(2):35. <http://dx.doi.org/10.1167/tvst.9.2.35>.
- [28] Andreeva R, Fontanella A, Giarratano Y, Bernabeu M. DR detection using Optical Coherence Tomography Angiography (OCTA): A transfer learning approach with robustness analysis. In: *Proc MICCAI OMIA works*, vol. 12069 LNCS, Lima, Peru; 2020, p. 11–20. http://dx.doi.org/10.1007/978-3-030-63419-3_2.
- [29] Ryu G, Lee K, Park D, Park S, Sagong M. A deep learning model for identifying diabetic retinopathy using optical coherence tomography angiography. *Sci Rep* 2021;11(1):23024. <http://dx.doi.org/10.1038/s41598-021-02479-6>.
- [30] Heisler M, Karst S, Lo J, Mammo Z, Yu T, Warner S, et al. Ensemble deep learning for diabetic retinopathy detection using optical coherence tomography angiography. *Transl Vis Sci Technol* 2020;9(2):20. <http://dx.doi.org/10.1167/tvst.9.2.20>.
- [31] Yasser I, Khalifa F, Abdeltawab H, Ghazal M, Sandhu H, El-Baz A. Automated diagnosis of optical coherence tomography angiography (OCTA) based on machine learning techniques. *Sensors* 2022;22(6):2342. <http://dx.doi.org/10.3390/s22062342>.
- [32] Zang P, Hormel T, Wang X, Tsuboi K, Huang D, Hwang T, et al. A diabetic retinopathy classification framework based on deep-learning analysis of OCT angiography. *Transl Vis Sci Technol* 2022;11(7):10. <http://dx.doi.org/10.1167/tvst.11.7.10>.
- [33] Li Y, El Habib Daho M, Conze P-H, Al Hajj H, Bonnin S, Ren H, et al. Multimodal information fusion for glaucoma and diabetic retinopathy classification. In: *Proc MICCAI OMIA works*, vol. 13576 LNCS, Singapore, Singapore; 2022, p. 53–62. http://dx.doi.org/10.1007/978-3-031-16525-2_6.
- [34] Li Y, El Habib Daho M, Conze P-H, Zeghlache R, Ren H, Lepicard C, et al. 3-D analysis of multiple OCTA acquisitions for the automatic diagnosis of diabetic retinopathy. In: *Proc ARVO*. New Orleans, LA, USA; 2023, p. C0116.
- [35] Sundararajan M, Taly A, Yan Q. Axiomatic attribution for deep networks. In: *Proc ICML*, vol. 70, Sydney, NSW, Australia: JMLR.org; 2017, p. 3319–28.
- [36] Samek W, Binder A, Montavon G, Lapuschkin S, Müller KR. Evaluating the visualization of what a deep neural network has learned. *IEEE Trans Neural Netw Learn Syst* 2017;28(11):2660–73. <http://dx.doi.org/10.1109/TNNLS.2016.2599820>.
- [37] Fei B, Chen X, Wang H, Sabol J, DuPont E, Gilkeson R. Automatic registration of CT volumes and dual-energy digital radiography for detection of cardiac and lung diseases. In: *Proc IEEE EMBC*, vol. 2006, New York, NY, USA; 2006, p. 1976–9. <http://dx.doi.org/10.1109/IEMBS.2006.259888>.
- [38] Schaffert R, Wang J, Fischer P, Borsdorf A, Maier A. Metric-driven learning of correspondence weighting for 2-D/3-D image registration. In: *Proc GCPR*, vol. 11269 LNCS, Stuttgart, Germany; 2018, p. 140–52. http://dx.doi.org/10.1007/978-3-030-12939-2_11.
- [39] Jaganathan S, Kukla M, Wang J, Shetty K, Maier A. Self-supervised 2D/3D registration for X-Ray to CT image fusion. In: *Proc IEEE/CVF WACV*. Waikoloa, HI, USA; 2023, p. 2787–97. <http://dx.doi.org/10.1109/WACV56688.2023.00281>.
- [40] Van Houtte J, Audenaert E, Zheng G, Sijbers J. Deep learning-based 2D/3D registration of an atlas to biplanar X-ray images. *Int J Comput Assist Radiol Surg* 2022;17(7):1333–42. <http://dx.doi.org/10.1007/s11548-022-02586-3>.
- [41] Li M, Chen Y, Ji Z, Xie K, Yuan S, Chen Q, et al. Image projection network: 3D to 2D image segmentation in OCTA images. *IEEE Trans Med Imaging* 2020;39(11):3343–54. <http://dx.doi.org/10.1109/TMI.2020.2992244>.
- [42] Lachinov D, Seeböck P, Mai J, Goldbach F, Schmidt-Erfurth U, Bogunovic H. Projective skip-connections for segmentation along a subset of dimensions in retinal OCT. In: de Bruijne M, Cattin PC, Cotin S, Padoy N, Speidel S, Zheng Y, Essert C, editors. *Proc MICCAI*. Lecture notes in computer science, Strasbourg, France: Springer International Publishing; 2021, p. 431–41. http://dx.doi.org/10.1007/978-3-030-87193-2_41.
- [43] Le C, Wang D, Villanueva R, Liu Z, Hammer D, Tao Y, et al. Novel application of long short-term memory network for 3D to 2D retinal vessel segmentation in adaptive optics—Optical coherence tomography volumes. *Appl Sci* 2021;11(20). <http://dx.doi.org/10.3390/app11209475>.
- [44] Forsgren E, Edlund C, Oliver M, Barnes K, Sjögren R, Jackson T. High-throughput widefield fluorescence imaging of 3D samples using deep learning for 2D projection image restoration. *PLoS One* 2022;17(5):e0264241. <http://dx.doi.org/10.1371/journal.pone.0264241>.
- [45] Haertter D, Wang X, Fogerson S, Ramkumar N, Crawford J, Poss K, et al. DeepProjection: Specific and robust projection of curved 2D tissue sheets from 3D microscopy using deep learning. *Dev (Camb)* 2022;149(21):dev200621. <http://dx.doi.org/10.1242/dev.200621>.
- [46] Shen L, Liu M, Wang C, Guo C, Meijering E, Wang Y. Efficient 3D junction detection in biomedical images based on a circular sampling model and reverse mapping. *IEEE J Biomed Health Inform* 2021;25(5):1612–23. <http://dx.doi.org/10.1109/JBHI.2020.3036743>.
- [47] Wang Y, Yan G, Zhu H, Buch S, Wang Y, Haacke E, et al. VC-Net: Deep volume-composition networks for segmentation and visualization of highly sparse and noisy image data. *IEEE Trans Comput Imaging* 2021;27(2):1301–11. <http://dx.doi.org/10.1109/TCVC.2020.3030374>.
- [48] Guo Y, Hormel TT, Pi S, Wei X, Gao M, Morrison JC, et al. An end-to-end network for segmenting the vasculature of three retinal capillary plexuses from OCT angiographic volumes. *Biomed Opt Express* 2021;12(8):4889–900. <http://dx.doi.org/10.1364/BOE.431888>, Publisher: Optica Publishing Group.
- [49] Statsenko Y, Habuza T, Talako T, Pazniak M, Likhorad E, Pazniak A, et al. Deep learning-based automatic assessment of lung impairment in COVID-19 pneumonia: Predicting markers of Hypoxia with computer vision. *Front Med* 2022;9:882190. <http://dx.doi.org/10.3389/fmed.2022.882190>.
- [50] Gupta V, Demirer M, Bigelow M, Little K, Candemir S, Prevedello L, White R, et al. Performance of a deep neural network algorithm based on a small medical image dataset: Incremental impact of 3D-to-2D reformation combined with novel data augmentation, photometric conversion, or transfer learning. *J Digit Imaging* 2020;33(2):431–8. <http://dx.doi.org/10.1007/s10278-019-00267-3>.
- [51] Mandal A, Priyam S, Chan H, Gouveia B, Guitera P, Song Y, et al. Computer-aided diagnosis of melanoma subtypes using reflectance confocal images. *Cancers* 2023;15(5):1428. <http://dx.doi.org/10.3390/cancers15051428>.
- [52] Simonyan K, Vedaldi A, Zisserman A. Deep inside convolutional networks: Visualising image classification models and saliency maps. In: *Proc ICLR works*. Calgary, Canada; 2014.
- [53] Zeiler MD, Fergus R. Visualizing and understanding convolutional networks. In: *Proc ECCV*. Zurich, Switzerland; 2014, p. 818–33.
- [54] Springenberg JT, Dosovitskiy A, Brox T, Riedmiller MA. Striving for simplicity: The all convolutional net. In: Bengio Y, LeCun Y, editors. *Proc ICLR works*. San Diego, CA, USA; 2015, URL: <http://arxiv.org/abs/1412.6806>.
- [55] Bach S, Binder A, Montavon G, Klauschen F, Müller K-R, Samek W. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One* 2015;10(7):e0130140. <http://dx.doi.org/10.1371/journal.pone.0130140>.
- [56] Shrikumar A, Greenside P, Kundaje A. Learning important features through propagating activation differences. In: *Proc ICML*, vol. 70, Sydney, NSW, Australia: JMLR.org; 2017, p. 3145–53.
- [57] Ribeiro MT, Singh S, Guestrin C. “Why should I trust you?”: Explaining the predictions of any classifier. In: *Proc KDD*. KDD '16, New York, NY, USA: Association for Computing Machinery; 2016, p. 1135–44. <http://dx.doi.org/10.1145/2939672.2939778>.
- [58] Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation. In: *Proc MICCAI*. Munich, Germany; 2015, p. 234–41. http://dx.doi.org/10.1007/978-3-319-24574-4_28.
- [59] Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Proc ICML*. ICML '15, Lille, France: JMLR.org; 2015, p. 448–56.
- [60] He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: *Proc CVPR*. Las Vegas, NV, USA; 2016, p. 770–8. <http://dx.doi.org/10.1109/CVPR.2016.90>.

- [61] Tassi CRN, Gawlikowski J, Fitri A, Triebel R. The impact of averaging logits over probabilities on ensembles of neural networks. In: Proc works AISafety. Vienna, Austria; 2022.
- [62] Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. In: Proc NIPS, vol. 25, Granada, Spain; 2012, p. 1097–105.
- [63] Tan M, Le QV. EfficientNet: Rethinking model scaling for convolutional neural networks. In: Proc ICML. Long Beach, CA, USA; 2019.
- [64] Liu Z, Mao H, Wu C, Feichtenhofer C, Darrell T, Xie S. A ConvNet for the 2020s. In: Proc CVPR. Los Alamitos, CA, USA: IEEE Computer Society; 2022, p. 11966–76. <http://dx.doi.org/10.1109/CVPR52688.2022.01167>.
- [65] Tan M, Le QV. EfficientNetV2: Smaller models and faster training. In: Proc ICML, vol. 139, Virtual; 2021, p. 10096–106, URL: <http://arxiv.org/abs/2104.00298>.
- [66] Hu J, Shen L, Albanie S, Sun G, Wu E. Squeeze-and-excitation networks. IEEE Trans Pattern Anal Mach Intell 2020;42(8):2011–23. <http://dx.doi.org/10.1109/TPAMI.2019.2913372>.
- [67] Aguilar-Ruiz JS, Michalak M. Multiclass classification performance curve. IEEE Access 2022;10:68915–21. <http://dx.doi.org/10.1109/ACCESS.2022.3186444>, Conference Name: IEEE Access.
- [68] Hara K, Kataoka H, Satoh Y. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? In: Proc CVPR. Los Alamitos, CA, USA: IEEE Computer Society; 2018, p. 6546–55. <http://dx.doi.org/10.1109/CVPR.2018.00685>.
- [69] Quellec G, Charrière K, Boudi Y, Cochener B, Lamard M. Deep image mining for diabetic retinopathy screening. Med Image Anal 2017;39:178–93. <http://dx.doi.org/10.1016/j.media.2017.04.012>.