



HAL
open science

ExplAIn: Explanatory artificial intelligence for diabetic retinopathy diagnosis

Gwenolé Quellec, Hassan Al Hajj, Mathieu Lamard, Pierre-Henri Conze,
Pascale Massin, Beatrice Cochener

► **To cite this version:**

Gwenolé Quellec, Hassan Al Hajj, Mathieu Lamard, Pierre-Henri Conze, Pascale Massin, et al.. ExplAIn: Explanatory artificial intelligence for diabetic retinopathy diagnosis. Medical Image Analysis, In press, 10.1016/j.media.2021.102118 . hal-03230907

HAL Id: hal-03230907

<https://imt-atlantique.hal.science/hal-03230907>

Submitted on 13 Jun 2023

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 - NonCommercial 4.0 International License



ExplAIIn: Explanatory Artificial Intelligence for Diabetic Retinopathy Diagnosis

Gwenolé Quellec^{a,*}, Hassan Al Hajj^{b,a}, Mathieu Lamard^{b,a}, Pierre-Henri Conze^{c,a}, Pascale Massin^d, Béatrice Cochener^{b,a,e}

^aInserm, UMR 1101, Brest, F-29200 France

^bUniv Bretagne Occidentale, Brest, F-29200 France

^cIMT Atlantique, Brest, F-29200 France

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

^eService d'Ophthalmologie, CHRU Brest, Brest, F-29200 France

ARTICLE INFO

Article history:

explanatory artificial intelligence, self-supervised learning, diabetic retinopathy diagnosis

ABSTRACT

In recent years, Artificial Intelligence (AI) has proven its relevance for medical decision support. However, the “black-box” nature of successful AI algorithms still holds back their wide-spread deployment. In this paper, we describe an eXplanatory Artificial Intelligence (XAI) that reaches the same level of performance as black-box AI, for the task of classifying Diabetic Retinopathy (DR) severity using Color Fundus Photography (CFP). This algorithm, called ExplAIIn, learns to segment and categorize lesions in images; the final image-level classification directly derives from these multivariate lesion segmentations. The novelty of this explanatory framework is that it is trained from end to end, with image supervision only, just like black-box AI algorithms: the concepts of lesions and lesion categories emerge by themselves. For improved lesion localization, foreground/background separation is trained through self-supervision, in such a way that occluding foreground pixels transforms the input image into a healthy-looking image. The advantage of such an architecture is that automatic diagnoses can be explained simply by an image and/or a few sentences. ExplAIIn is evaluated at the image level and at the pixel level on various CFP image datasets. We expect this new framework, which jointly offers high classification performance and explainability, to facilitate AI deployment.

© 2021 Elsevier B. V. All rights reserved.

1. Introduction

Diabetic Retinopathy (DR) is a leading and growing cause of vision impairment and blindness: by 2040, around 600 million people throughout the world will have diabetes (Ogurtsova et al., 2017), a third of whom will have DR (Yau et al., 2012). Early diagnosis is key to slowing down the progression of DR and therefore preventing the occurrence of blindness. Annual retinal screening, generally using Color Fundus Photography (CFP), is thus recommended for all diabetic patients (Javitt and

Aiello, 1996). However, the goal of annual screening for all diabetic patients represents a huge burden on ophthalmologists and it is far from being achieved (Benoit et al., 2019). In order to improve DR screening programs, numerous Artificial Intelligence (AI) systems were thus developed to automate DR diagnosis using CFP (Ting et al., 2019b). However, due to the “black-box” nature of state-of-the-art AI, these systems still need to gain the trust of clinicians and patients.

To gain this trust, one solution investigated by Araújo et al. (2020) and Ayhan et al. (2020) is to design AI systems able to reliably estimate the uncertainty level of their predictions. This feature is expected to help clinicians know when AI predictions should be carefully reviewed and when they can be trusted. Another solution, investigated by Abramoff et al. (2016), is to

*LaTIM - IBRBS - 22, avenue Camille Desmoulins - 29200 Brest, France - Tel.: +33 2 98 01 81 29

e-mail: gwenole.quellec@inserm.fr (Gwenolé Quellec)

develop two-stage AI systems that 1) learn to detect or segment lesions considered relevant by ophthalmologists (microaneurysms, exudates, etc.) and 2) base the AI predictions on these detections. Because they mimic ophthalmologists' reasoning, clinicians are more likely to adopt them. A similar approach was investigated by Fauw et al. (2018) for the classification of optical coherence tomography images. However, these approaches cannot generalize easily to new imaging modalities or new decision problems, such as DR progression prediction (Arcadu et al., 2019), since relevant patterns are not fully known to ophthalmologists. Alternatively, another solution investigated by Quellec et al. (2017) and Sayres et al. (2019) is to help clinicians interpret AI predictions by highlighting image regions supposedly involved in AI predictions. If clinicians agree with highlighted areas, they will more likely trust the AI and eventually adopt it. Note that a similar approach was recently investigated for medical image segmentation (Wickstrøm et al., 2020). However, these visualization methods provide limited information: they tell us which pixels seem to play a role in the decision process, but they do not tell us precisely how. Although interpretability is an interesting feature, it may not be enough to gain the trust of clinicians. And it is certainly not enough to gain the trust of patients, which would rather have an explanation.

Gilpin et al. (2018) differentiate interpretability and explainability as follows: *interpretability* is the science of comprehending what a model did or might have done, while *explainability* is the ability to summarize the reasons for an AI behavior. Explainability implies interpretability, but the reverse is not always true. To gain the trust of patients and clinicians, explainability is desirable. EXplanatory Artificial Intelligence (XAI) is a growing field of research (Gilpin et al., 2018) motivated by potential AI users, worried about safety (Russell et al., 2015). It is also pushed by European regulations and others: the goal is to grant users the right for an explanation about algorithmic decisions that were made about them (Goodman and Flaxman, 2017). **In particular, explainability is an important consideration when clearing autonomous diagnostic AI products (Amann et al., 2020; Abrmoff et al., 2018).** An Explanatory AI system, called ExplAIIn, is presented and evaluated in this paper. Unlike visualization methods above (Quellec et al., 2017; Sayres et al., 2019), ExplAIIn does not attempt to retrospectively analyze a complex classification process. Instead, it modifies the classification process in such a way that it can be understood directly.

“Black-box” image classification AI algorithms are usually defined as Convolutional Neural Networks (CNNs) or, more generally, as ensembles of multiple CNNs (Ting et al., 2019a; Quellec et al., 2019). Each of these CNNs is supervised at the image level: given an image, one or several experts indicate which labels should be assigned to this image. To enable explainability, we propose to include a pixel-level classification step into the neural network. Pixel-level classification, also known as image segmentation, is generally performed by an Encoder-Decoder Network (EDN) (Ronneberger et al., 2015; Lin et al., 2017; Conze et al., 2020). EDNs, however, are supervised at the pixel level: given an image, one or several ex-

perts assign a label to each pixel in the image. ExplAIIn bridges the gap between the two paradigms: pixel-level classification and image-level classification are trained simultaneously, using image-level supervision only.

DR severity assessment is an ideal task for evaluating ExplAIIn. In the International Clinical Diabetic Retinopathy (ICDR) severity scale (Wilkinson et al., 2003), for instance, the severity level directly derives from the abnormalities observed in CFP images: the relationship between pixel-level and image-level classifications is known. We took advantage of this property for evaluation purposes. First, we trained ExplAIIn to automate the ICDR scale at the image level. Next, we checked whether or not 1) pixel-level classification and 2) rules linking pixel-level and image-level classifications are consistent with the ICDR scale.

The paper is organized as follows. Related machine learning frameworks are presented in section 2. The proposed ExplAIIn solution is described in section 3. This framework is applied to DR diagnosis in section 4. We end up with a discussion and conclusions in section 5.

2. Related Machine Learning Frameworks

In terms of purpose, ExplAIIn is related to existing algorithms for visualizing/interpreting what image classification CNNs have learnt. Given a trained classification CNN and an input image, these algorithms compute the influence of each pixel on CNN predictions. In the occlusion method, patches are occluded in the input image and the occluded image is run through the CNN: a drop in classification performance indicates that the occluded patch is relevant (Zeiler and Fergus, 2014). In sensitivity analysis, the backpropagation algorithm is used to compute the gradient of CNN predictions with respect to the value of each input pixel (Simonyan et al., 2014). Various improvements on sensitivity analysis, including layer-wise relevance propagation (Bach et al., 2015), also rely on backpropagated quantities to build a high-resolution heatmap showing the relevance of each pixel (Samek et al., 2017; Quellec et al., 2017). Next, Class Activation Mapping (CAM) was proposed for image classification CNNs containing a Global Average Pooling (GAP) layer after the last convolutional layer: CAM computes CNN predictions for each of the GAP's input locations rather than for the GAP's output, thus providing coarse-resolution class-specific activation maps (Zhou et al., 2016). Grad-CAM generalizes this idea to any classification CNN architecture (Selvaraju et al., 2017). These visualization methods attempt to retrospectively analyze a complex classification process. Another approach, investigated in this paper, is to replace the classification process with one that can be understood directly.

Because deep neural networks are hard to interpret, several authors have proposed to train deep neural decision trees (Yang et al., 2018) or deep neural decision forests (Kontschieder et al., 2016; Hehn et al., 2020) instead. These architectures are indeed based on rules that can be interpreted more easily by humans. Frosst and Hinton (2017) thus proposed a training procedure to derive a deep neural decision tree from a deep neural network,

so that it can be interpreted. However, understanding a deep decision tree is still a challenging task for machine learning agnostics (Hehn et al., 2020): such an algorithm is interpretable, but not explainable. To enable explanations, the solution investigated in ExplAIIn is rather to base the classification process on a segmentation and a categorization of the pathological signs, obtained solely through weak supervision.

In that sense, ExplAIIn is also related to Weakly-Supervised Semantic Segmentation (WSSS). In WSSS, the goal is to predict pixel classes using image labels only for supervision. In particular, no positional information, like manual segmentations or bounding boxes, is required for supervision. Note that segmentation is the end goal of WSSS, while it is an intermediate step in ExplAIIn. WSSS solutions can be classified into four categories (Chan et al., 2019). (1) Expectation-Maximization solutions use image annotations to initialize prior assumptions about the class distribution in images. Next, an EDN is trained to meet those constraints. Then, the prior assumption model is updated based on the EDN features, and the training cycle is repeated again until convergence (Papan-dreou et al., 2015; Pathak et al., 2015; Kervadec et al., 2019). (2) Multiple-Instance Learning (MIL) solutions train an image classification CNN with image-level supervision and then infer the image locations responsible for each class prediction: inference relies on the MIL assumption that an image belongs to one class if and only if at least one of its pixels does (Shimoda and Yanai, 2016; Durand et al., 2017). (3) Self-supervised learning solutions train an image classification CNN with image-level supervision to obtain a coarse-resolution CAM. Next, a segmentation EDN is trained using the CAM as ground truth to obtain a higher-resolution segmentation (Kolesnikov and Lampert, 2016; Zhang et al., 2020). Interestingly, the Reliable Region Mining (RRM) solution by Zhang et al. (2020) merges the two steps (CAM computation and segmentation) into one. (4) In a final category, object proposals are extracted first and the most probable class is assigned to each of them, using coarse-resolution CAMs obtained as above (Kwak et al., 2017; Zhou et al., 2018; Laradji et al., 2019). The self-supervised learning approach seems to be the most popular nowadays (Ahn et al., 2019; Lee et al., 2019; Jiang et al., 2019; Zhang et al., 2020; Wang et al., 2020). It should be noted that ExplAIIn is more general than WSSS solutions since the number of pixel labels can be different from the number of image labels; in particular, it can be larger. As a result, ExplAIIn is more semantically rich and it can lead to better image classification, as it allows for more general “pixel to image” label inference rules.

Note that the proposed framework is related to the GP-Unet solution by Dubost et al. (2020): both solutions include an EDN inside an image classification CNN, in order to generate high-resolution attention maps. ExplAIIn differs in that it generates multiple, complementary maps: one map is generated per type of discriminant patterns in images. It also differs in that image classification derives very simply from the pixel classifications in order to allow explainability, rather than interpretability. Finally, ExplAIIn introduces a new criterion, namely the generalized occlusion method, in order to optimize foreground/background separation in the pixel classification maps.

3. Explanatory Artificial Intelligence

3.1. Overview and Notations

This paper addresses multilabel image classification: given an input image I and N image-level labels, the goal is to predict whether or not experts would assign the n -th label to image I , $\forall n \in \{1, \dots, N\}$. Let $p_n \in [0; 1]$ denote the probabilistic prediction of ExplAIIn and let $\delta_{I,n} \in \{0, 1\}$ denote the ground truth: did experts actually assign the n -th label to image I ? Unlike multiclass classification, multilabel classification does not assume image-level labels to be mutually exclusive: $\sum_{n=1}^N \delta_{I,n} \in \{0, 1, \dots, N\}$. Images where $\sum_{n=1}^N \delta_{I,n} = 0$ are referred to as “background images”, i.e. images where experts did not annotate anything.

As an intermediate step, ExplAIIn also assigns a label to each (color or grayscale) pixel $I_{x,y}$ for explainability purposes. Let M denote the number of pixel-level labels:

- M can be smaller than, equal to, or larger than the number N of image-level labels,
- the first of these pixel-level labels represents “background pixels”.

A pixel-level prediction tensor P is thus calculated, where $P_{m,x,y} \in [0; 1]$ indicates the probability that pixel $I_{x,y}$ should be assigned the m -th pixel-level label, with $m \in \{1, \dots, M\}$. Unlike image-level labels, pixel-level labels are assumed to be mutually exclusive: $\sum_{m=1}^M P_{m,x,y} = 1$ (multiclass classification). Since our framework is trained with image supervision only, we assume that no ground truth is available for pixel-level labels.

As illustrated in Fig. 1, the pixel-level and image-level classification problems are solved jointly as follows:

- an Encoder-Decoder Network (EDN) s predicts P from I ,
- a classification head c predicts $\mathbf{p} = \{p_n, \forall n \in \{1, \dots, N\}\}$ from P ,
- additional branches are included during training, to compute auxiliary losses improving explainability.

3.2. Pixel-Level Label Prediction

In order to predict pixel-level labels $P_{m,x,y}$ for each pixel $I_{x,y}$, an EDN s is used: $P = s(I)$. This EDN is composed of an encoder network e and a decoder network d such that $s = d \circ e$.

The encoder part e of s was defined as an EfficientNet (Tan and Le, 2019). EfficientNet is a family of CNN architectures of increasing complexity. The smallest model, namely EfficientNet-B0, was obtained through neural architecture search. Larger models, EfficientNet-B1 to -B7, were then obtained by up-scaling EfficientNet-B0: depth, width and resolution were increased proportionally. Although any classification CNN may be used as backbone for s , EfficientNets were chosen for their clearly superior tradeoff between accuracy and complexity, compared to traditional deep networks.

The overall architecture of s was defined as a Feature Pyramid Network (FPN) by Lin et al. (2017). Like the well-know U-Net (Ronneberger et al., 2015), an FPN is a top-down architecture with skip-connections between e and d . Unlike U-Net, high-level semantic feature maps are produced at multiple

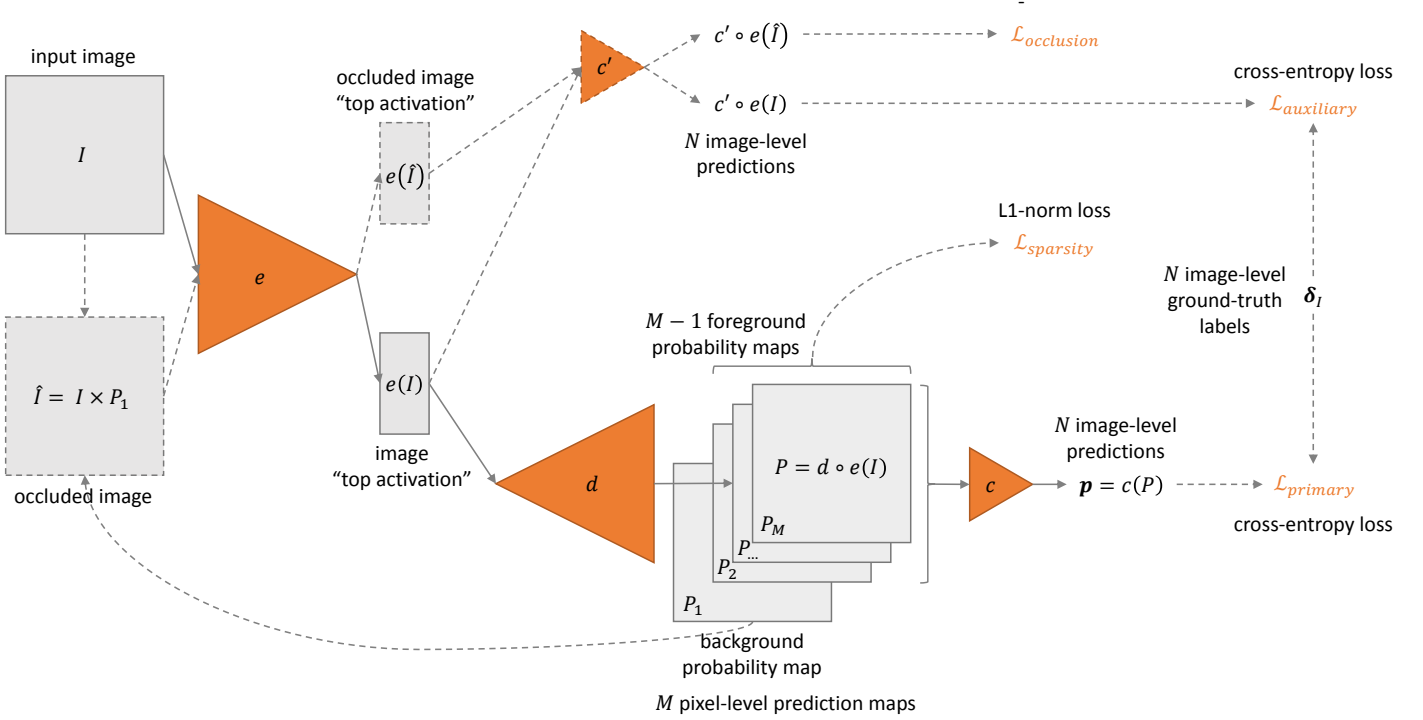


Fig. 1: Outline of the ExplAI framework. Dashed lines indicate operations performed during training only.

scales: these multi-scale feature maps are then 1) up-scaled to match the size of I , 2) concatenated and 3) processed by a final convolutional layer to obtain the output tensor P , following a deep supervision strategy (Xie and Tu, 2015). Although any EDN architecture may be used, FPN was chosen for its faster convergence.

Because pixel-level labels are mutually exclusive, a softmax operator μ was used as activation function for the last convolutional layer:

$$\mu(\mathbf{z}) = \left\{ \frac{e^{z_m}}{\sum_{i=1}^M e^{z_i}}, \forall m \in \{1, \dots, M\} \right\}, \quad (1)$$

where \mathbf{z} are the outputs of the M neurons at a given pixel location.

3.3. Image-Level Label Prediction

In order to predict the vector \mathbf{p} of image-level labels from the tensor P of pixel-level labels, a classification head c is used: $\mathbf{p} = c(P) = c \circ s(I)$. c has to be very simple to enable explainability. In practice, it consists of two layers only.

3.3.1. Summary Layer

The first layer, Π , summarizes each label prediction map P_m by a single scalar value. Two easily understandable summaries may be considered:

1. the average value, which is proportional to the surface covered by each pixel-level label,
2. the maximal value, which represents the strongest clue of presence for each pixel-level label.

In theory, both solutions enable explainability. However, in practice, the first option has one major drawback: it favors over-segmentation in foreground images, which limits explainability. Indeed, to improve separation between background and foreground images, the EDN is encouraged to increase the average predictions for foreground pixel-level labels in foreground images (and decrease them in background images): one way is to increase the number of foreground pixels in foreground images, i.e. to oversegment. **In the second option, the maximal value is independent from the number of foreground pixels. Therefore, this option does not favor over- or under-segmentation: this option was selected.** Π is thus defined as the global max pooling layer:

$$\Pi(P) = \left\{ \max_{x,y} P_{m,x,y}, \forall m \in \{2, \dots, M\} \right\}. \quad (2)$$

Note that the background probability map P_1 is not used for image-level prediction.

3.3.2. Classification Layer

The second layer, Δ , is a special class of dense layers, where neural weights are all positive. The positivity constraint improves explainability: the image-level prediction is defined as a sum of pixel-level clues, each clue being associated with a (positive) confidence level. Layer Δ is defined as follows:

$$\Delta(\mathbf{z}) = \left\{ \sigma \left(\sum_{m=2}^M z_m w_{m,n}^2 + b_n \right), \forall n \in \{1, \dots, N\} \right\}, \quad (3)$$

where $w_{m,n}^2$ are positive neural weights, b_n are biases and σ is an activation function. Because image-level labels are not mu-

tually exclusive, σ was defined as a sigmoid function:

$$\sigma(z) = \frac{1}{1 + \exp(-z)}. \quad (4)$$

3.4. Learning to Detect Background Pixels

In ExplAIn, pixel-level labels are not necessarily related to image-level labels. However, for localization purposes, “background pixels” are related to “background images”. We remind that: 1) background pixels are those associated with the first pixel-level label and 2) background images are those associated with no image-level label ($\sum_{n=1}^N \delta_{I,n} = 0$). To improve explainability, we propose that background images only contain background pixels.

To ensure this property, a generalization of the occlusion method (Zeiler and Fergus, 2014) is proposed. The original occlusion method generates multiple occluded versions of the input image I by zeroing all pixel intensities inside a sliding square window. These occluded versions are then run through a previously trained CNN in order to detect background (or conversely foreground) pixels. Instead, we propose to generate a single occluded version \hat{I} of the input image. Let $P_1 = \{P_{1,x,y}, \forall(x,y)\}$ denote the background probability map. The occluded image \hat{I} is defined as follows:

$$\hat{I} = I \times P_1 = \{I_{x,y} P_{1,x,y}, \forall(x,y)\}, \quad (5)$$

where \times is the element-wise product. If I is a color image, each color component of pixel $I_{x,y}$ is multiplied by the same value $P_{1,x,y}$. Unlike the original occlusion method, this generalized occlusion method is performed during training, in order to optimize the background probability map P_1 . The following two properties are optimized jointly:

Occlusion sensitivity: the occluded image $\hat{I} = I \times P_1$ should always be perceived as a background image, regardless of the ground-truth image-level labels. This indicates that all relevant pixels have been successfully occluded: occlusion is sensitive.

Occlusion specificity: the background P_1 should be as extended as possible or, conversely, the foreground image $1 - P_1$ should be as sparse as possible. This indicates that occlusion is specific to relevant pixels.

3.5. Auxiliary Classification Branch

In order to optimize the first property above, namely occlusion sensitivity, the occluded image $\hat{I} = I \times P_1$ should run through a classifier, and the background image P_1 should be optimized in such a way that \hat{I} is predicted as a background image. The $c \circ s$ classifier may be used for that purpose. However, optimizing $c \circ s(\hat{I})$ would not only alter background pixel detection, it will potentially alter the entire classifier. Therefore, an independent classification branch should be used instead of $c \circ s$.

For the related problem of CAM computation in WSSS, Zhang et al. (2020) used a completely independent classifier. However, we assume that foreground/background separation is mainly performed by the decoder part d of the EDN $s = d \circ e$

(see section 3.2). Therefore, we propose to reuse the encoder part e of s for this auxiliary classification branch. It has the advantage of significantly reducing training complexity, while allowing a more generic feature extraction. In that purpose, an auxiliary classification head c' is connected to the L -channel tensor $T = e(I)$, i.e. to the “top activation” layer of s . Following common practice, this classification head consists of a global average pooling layer, followed by a regular dense layer. Like c , this classification head has N non-mutually exclusive outputs:

$$c'(T) = \left\{ \sigma \left(\sum_{l=1}^L \frac{\sum_{x,y} T_{l,x,y}}{\sum_{x,y} 1} w'_{l,n} + b'_n \right), \forall n \in \{1, \dots, N\} \right\}, \quad (6)$$

where σ is the sigmoid function of Eq. (4), and where $w'_{l,n}$ and b'_n are neural weights and biases, respectively.

To summarize, $c' \circ e$ is the classification branch used to classify occluded images, for the purpose of optimizing background images.

3.6. Loss Functions

The entire neural architecture has been described. This section enumerates the loss functions to optimize both pixel-level and image-level classification performance.

3.6.1. Cross-Entropy Losses

The main goal of the proposed framework is to correctly classify image-level labels. Given network predictions $\mathbf{p} = c \circ s(I)$ and ground truth labels $\delta_I = \{\delta_{I,n}, \forall n \in \{1, \dots, N\}\}$, the primary loss function $\mathcal{L}_{primary}$ is thus defined as a cross-entropy loss function:

$$\mathcal{L}_{primary} = - \frac{1}{N} \sum_I \sum_{n=1}^N [\delta_{I,n} \log(c \circ s(I)_n) + (1 - \delta_{I,n}) \log(1 - c \circ s(I)_n)], \quad (7)$$

where $c \circ s(I)_n = p_n, n \in \{1, \dots, N\}$.

Because an auxiliary classification branch c' is defined in section 3.5, N auxiliary image-level predictions $\mathbf{p}' = c' \circ e(I)$ should also be optimized. An auxiliary loss function $\mathcal{L}_{auxiliary}$ is thus defined similarly to $\mathcal{L}_{primary}$, by replacing $c \circ s$ with $c' \circ e$ in equation 7.

3.6.2. Occlusion Loss

An additional loss function $\mathcal{L}_{occlusion}$ is defined to optimize the first property of the background probability map P_1 , namely that the occluded image $\hat{I} = I \times P_1$ should always be perceived as a background image (see section 3.4). Given the definition of a background image, namely $\sum_{n=1}^N \delta_{I,n} = 0$, $\mathcal{L}_{occlusion}$ is defined as the squared L2-norm of $c' \circ e(\hat{I})$ predictions:

$$\mathcal{L}_{occlusion} = \frac{\sum_I \sum_{n=1}^N [c' \circ e(\hat{I})_n]^2}{N \sum_I 1}. \quad (8)$$

3.6.3. Sparsity Loss

A final loss function $\mathcal{L}_{sparsity}$ is defined to optimize the second property of the background probability map P_1 , namely that the foreground map $1 - P_1$ should be sparse (see section 3.4). $\mathcal{L}_{sparsity}$ is defined as the L1-norm of $1 - P_1$ maps. Because a softmax activation function is used at the end of s , $\mathcal{L}_{sparsity}$ can be expressed as follows (see Eq. (1)):

$$\mathcal{L}_{sparsity} = \frac{\sum_I \sum_x \sum_y \left| \sum_{m=2}^M s(I)_{m,x,y} \right|}{\sum_I \sum_x \sum_y 1}. \quad (9)$$

Note that the modulus operator can be dropped since $0 \leq \sum_{m=2}^M s(I)_{m,x,y} \leq 1, \forall (x, y)$, due to the use of a softmax activation function in section 3.2.

3.6.4. Total Loss

Those four loss functions are combined linearly to obtain the total loss function \mathcal{L} that should be minimized during training:

$$\mathcal{L} = \mathcal{L}_{primary} + \alpha \mathcal{L}_{auxiliary} + \beta \mathcal{L}_{occlusion} + \gamma \mathcal{L}_{sparsity}, \quad (10)$$

where $\alpha \geq 0, \beta \geq 0$ and $\gamma \geq 0$.

3.6.5. Loss Function Competition

Ideally, convergence to a suitable classifier should not depend critically on the choice of α, β and γ weights. In other words, the four basic loss functions should not compete with one other.

First, $\mathcal{L}_{auxiliary}$ does not compete with $\mathcal{L}_{primary}$: whether we want to segment and classify pathological signs (through $\mathcal{L}_{primary}$) or directly classify images without segmentation (through $\mathcal{L}_{auxiliary}$), the objective of the shared encoder e is to encode the presence of pathological signs in images.

$\mathcal{L}_{occlusion}$ and $\mathcal{L}_{sparsity}$, on the other hand, are competing. By design, these loss functions optimize occlusion sensitivity and specificity (see section 3.4), two metrics that always need to be traded off. $\mathcal{L}_{sparsity}$ and $\mathcal{L}_{primary}$ are also competing: if all pixels are assigned to the background (high foreground sparsity), then no pathological sign can be detected and the classifier will always predict background images, leading to poor classification performance. Fortunately, assuming that decoder d ensures foreground/background separation, $\mathcal{L}_{auxiliary}$ and $\mathcal{L}_{sparsity}$ are independent. Therefore, encoder e , the largest part of the network, can always be trained through $\mathcal{L}_{auxiliary}$, even if the foreground is temporarily too sparse. Besides training complexity, this property also motivates the use of a partly independent classification branch for occlusion sensitivity maximization, over a completely independent classification branch (see section 3.5).

Finally, $\mathcal{L}_{occlusion}$ has no reason to compete with $\mathcal{L}_{primary}$ and $\mathcal{L}_{auxiliary}$. In conclusion, γ , the weight controlling $\mathcal{L}_{sparsity}$, is the most critical weight. It can be used to trade-off image-level classification performance and pixel-level classification quality (i.e. explainability).

3.7. Explanation Generation

Once training has converged, the proposed system can be used to infer automatic diagnoses for an unseen image I . In

addition to N image-level predictions, $M - 1$ pixel-level probability maps are obtained (see Fig. 1). The following procedure is proposed to explain automatic diagnoses for I :

1. The $M - 1$ pixels maximizing pixel-level prediction in each of the $M - 1$ foreground maps can be highlighted in I (see Eq. (2)).
2. Let i_2, i_3, \dots, i_M denote the (positive) intensity of those pixels. To explain the n -th image-level decision, those intensities should be multiplied, respectively, by the (positive) weights $w_{2,n}^2, w_{3,n}^2, \dots, w_{M,n}^2$ involved in c (see Eq. (3)). Product $i_m w_{m,n}^2$ indicates the strength of clue number m in the n -th decision.
3. Additionally, if experts can associate a keyword (typically lesion names) with each pixel-level label, then the previous step can be converted into a set of sentences.

4. Application to Diabetic Retinopathy Diagnosis

The proposed framework has been applied to the diagnosis of DR through the analysis of CFP images. The goal was to grade DR severity in one eye according to the ICDR scale (Wilkinson et al., 2003):

No apparent DR – no abnormalities.

Mild nonproliferative DR (NPDR) – microaneurysms only.

Moderate NPDR – more than just microaneurysms but less than severe NPDR.

Severe NPDR – any of the following: more than 20 intraretinal hemorrhages in each of 4 quadrants; definite venous beading in 2 quadrants; prominent intraretinal microvascular abnormalities in 1 quadrant and no signs of proliferative DR (PDR).

PDR – one or more of the following: neovascularization, vitreous/preretinal hemorrhage.

By design, this multiclass classification problem can be transformed into a four-label multilabel classification problem ($N = 4$): ‘at least mild NPDR’, ‘at least moderate NPDR’, ‘at least severe NPDR’ and ‘PDR’. Each of these labels is associated with the presence of specific abnormalities (“microaneurysms” for ‘at least mild NPDR’, “more than just microaneurysms” for ‘at least moderate NPDR’, etc.). Note that 1) background images correspond to ‘no apparent DR’ and that 2) detecting the ‘no apparent DR’ is the exact opposite of detecting ‘at least mild NPDR’.

4.1. Image Datasets

A model was developed using training and validation datasets with image-level labels. Next, this model was evaluated on multiple datasets with image-level labels and on one dataset with pixel-level labels. In all datasets, images were graded according to the ICDR scale. In practice, image-level labels were not assigned to single images but, more generally, to small sets of images (CFPs) from the same eye. These datasets have various origins: USA, France, India and China.

4.1.1. OPHDIAT Image-Level Evaluation Dataset

The first image-level evaluation dataset originates from the OPHDIAT DR screening network, which consists of 40 screening centers located in 22 diabetic wards of hospitals, 15 primary health-care centers and 3 prisons in the Ile-de-France area (Massin et al., 2008). Let $OPHDIAT_{eval}$ denote this dataset. This dataset consists of 21,576 CFPs from 9,734 eyes of 4,996 patients (~ 2 CFPs per eye). As part of the quality-assurance program of OPHDIAT, DR severity in these randomly-selected patients was graded by two ophthalmologists. In case of disagreement, images were read a third time by a senior ophthalmologist.

4.1.2. Training and Validation Datasets

The training and validation datasets originate from two DR screening programs: OPHDIAT and EyePACS (Cuadros and Bresnick, 2009).

Images from EyePACS were acquired in multiple primary care sites throughout California and elsewhere. A total of 88,702 CFPs from 88,702 eyes of 44,351 patients (one CFP per eye), released for Kaggle’s Diabetic Retinopathy Detection challenge¹, was used. Training and validation images from OPHDIAT consist of all images from patients that were not included in the image-level evaluation dataset: a total of 610,748 CFPs from 275,236 eyes of 142,145 patients (~ 2 CFPs per eye) were included. In both subsets, DR severity in each eye was graded by a single human reader.

Images were distributed as follows: 90% were used for training and 10% were used for validation. The same proportion of eyes from EyePACS and OPHDIAT (24.4% / 75.6%) was used in the training and validation datasets.

4.1.3. IDRiD Image- and Pixel-Level Evaluation Datasets

Performance was also evaluated on the Indian Diabetic Retinopathy Image Dataset (IDRiD)², which originates from an eye clinic located in Nanded, India. This dataset was collected for the purpose of an image analysis challenge (Porwal et al., 2019). For image-level evaluation, we used the 103 CFPs released for the onsite evaluation of algorithms competing in sub-challenge 2 (disease grading). Let $IDRiD_{image}$ denote this dataset. For pixel-level evaluation, we used the 143 CFPs released for training the algorithms competing in sub-challenge 1 (lesion segmentation). Let $IDRiD_{pixel}$ denote this dataset. Experts manually segmented four types of abnormalities related to DR in those images: microaneurysms, hemorrhages, soft exudates and hard exudates (Porwal et al., 2019). For each CFP, one binary segmentation map is thus available per lesion type. No training or fine-tuning was performed on IDRiD data, to ensure independent evaluation.

4.1.4. DeepDR Image-Level Evaluation Dataset

Finally, image-level performance was evaluated on images from the DeepDR challenge³. We used the 400 CFPs from

200 eyes of 100 patients (2 CFPs per eye) released for validating the algorithms competing in sub-challenge 1 (disease grading). Part of these images originate from screening programs in China: the Shanghai Diabetic Complication Screening Project (SDCSP), the Nienheng Diabetes Screening Project (NDSP), and Nationwide Screening for Complications of Diabetes (NSCD). Some images originate from an eye clinic located in the Department of Ophthalmology, Shanghai Jiao Tong University, China. Let $DeepDR_{valid}$ denote this dataset. The onsite evaluation dataset, denoted by $DeepDR_{onsite}$, is not publicly-available yet. No training or fine-tuning was performed on DeepDR data, to ensure independent evaluation.

4.2. Image Pre-processing

The image pre-processing procedure proposed by B. Graham for the “min-pooling” solution, which ranked first in the Kaggle Diabetic Retinopathy competition, was followed.⁴ First, to focus the analysis on the camera’s field of view, images were adaptively cropped and resized to 448×448 pixels. Second, to attenuate intensity variations throughout and across images, image intensity was normalized. Normalization was performed in each color channel independently: the background was estimated using a large Gaussian kernel and subtracted from the image (Quelleg et al., 2017).

Data augmentation was performed during training. Before feeding a pre-processed image to a CNN, the image was randomly rotated (range: $[0^\circ, 360^\circ]$), translated (range: $[-15\%, 15\%]$), scaled (range: $[80\%, 120\%]$) and horizontally flipped. Additionally, Gaussian noise (standard deviation: 5% of the intensity range) was added and the contrast was randomly scaled (range: $[75\%, 125\%]$).

4.3. Managing Eye-Level Labels

In OPHDIAT, which is used for training, validation and image-level evaluation, labels (DR severity) are assigned to eyes. To facilitate training, a single image was selected per eye at the beginning of each epoch: the one maximizing the sum of image-level predictions ($\sum_{n=1}^N p_n$), i.e. the most pathological image of each eye. For image-level evaluation, the maximal prediction among images of the same eye was considered for each label $n \in \{1, \dots, N\}$.

4.4. Baseline Methods

4.4.1. Reliable Region Mining (RRM)

The self-supervised WSSS solution by (Zhang et al., 2020), named Reliable Region Mining (RRM), was used as baseline. Like the proposed solution, RRM jointly trains a classification branch and a segmentation branch. The purpose of the classification branch in RRM is to compute the CAMs, which are used to supervise the segmentation branch, after post-processing by a dense Conditional Random Field (CRF). However, unlike the proposed solution, segmentations are not used for classification, so only segmentation results can be compared. RRM also

¹<https://www.kaggle.com/c/diabetic-retinopathy-detection/data>

²<https://idrid.grand-challenge.org>

³<https://isbi.deepdr.org>

⁴<https://www.kaggle.com/c/diabetic-retinopathy-detection/discussion/15801>

use FPN as EDN architecture. For a fair comparison, the same family of classification CNNs (EfficientNet) was used as backbones.

4.4.2. ExplAIn-CAM

In order to evaluate the proposed occlusion strategy (see section 3.4), a variation on ExplAIn and RRM, called ExplAIn-CAM, was also evaluated. In ExplAIn-CAM, the background probability map P_1 is not trained to maximize occlusion sensitivity and specificity. Instead, what is maximized is the Dice similarity coefficient between the foreground probability map $(1 - P_1)$ and the element-wise maximum of class-specific CAMs. The CAMs are computed as in RRM, using the classification branch.

4.4.3. “Black-Box” AI

Finally, in order to evaluate the impact of explainability on image-level performance, ExplAIn was compared to a previous “black-box” AI solution from our group, relying on an ensemble of multiple CNNs (Quellec et al., 2019).

4.5. Implementation Details

ExplAIn was implemented using the *Keras* API of the *TensorFlow* library⁵. Encoder-Decoder Networks (EDNs) were based on the *Segmentation Models* library⁶. The *OpenCV*⁷ and *imgaug*⁸ libraries were used for image preprocessing and data augmentation, respectively.

4.6. Image-Level Performance

4.6.1. Performance Metrics

First, we evaluated multilabel classification performance at the image level. Receiver-Operating Characteristics (ROC) curves were used: one ROC curve was obtained per image-level label $n \in \{1, \dots, N\}$. Each curve was obtained by varying a cutoff on the output probabilities p_n and by measuring the classification sensitivity and specificity for each cutoff; ROC curves were obtained by connecting (1-specificity, sensitivity) pairs. Classification performance for one label can be summarized by the Area Under the ROC Curve (AUC).

Next, we evaluated multiclass classification performance at the image level. In this scenario, one severity level s must be assigned to each image: either ‘no apparent DR’, ‘mild NPDR’, ‘moderate NPDR’, ‘severe NPDR’ or ‘PDR’. Given four probability cutoffs τ_1, τ_2, τ_3 and τ_4 , this severity level was defined as follows:

- if $p_4 \geq \tau_4$: $s = \text{‘PDR’}$,
- else if $p_3 \geq \tau_3$: $s = \text{‘severe NPDR’}$,
- else if $p_2 \geq \tau_2$: $s = \text{‘moderate NPDR’}$,
- else if $p_1 \geq \tau_1$: $s = \text{‘mild NPDR’}$,
- else: $s = \text{‘no apparent DR’}$.

⁵<https://www.tensorflow.org>

⁶https://github.com/qubvel/segmentation_models

⁷<https://github.com/opencv/opencv-python>

⁸<https://imgaug.readthedocs.io>

terms of AUC, in the OPHDIAT evaluation dataset. The “/Bx” suffix in the architecture names indicates the EfficientNet-“Bx” model used as backbone. The selected ExplAIn model, which maximizes the mean AUC (both in the validation subset and in the OPHDIAT evaluation dataset), is in bold. EfficientNet-B1 is the best backbone for ExplAIn-CAM in terms of mean AUC (both in the validation subset and in the OPHDIAT evaluation dataset). ‘Combination’ denotes the optimal combination of multiple ExplAIn models.

architecture	at least mild NPDR	at least moderate NPDR	at least severe NPDR	PDR
ExplAIn/B0	0.9500	0.9886	0.9968	0.9938
ExplAIn/B1	0.9586	0.9917	0.9970	0.9905
ExplAIn/B2	0.9602	0.9930	0.9971	0.9890
ExplAIn/B3	0.9535	0.9906	0.9972	0.9933
ExplAIn/B4	0.9634	0.9929	0.9972	0.9900
ExplAIn/B5	0.9682	0.9939	0.9978	0.9945
ExplAIn/B6	0.9627	0.9923	0.9972	0.9945
ExplAIn/B7	0.9620	0.9918	0.9941	0.9914
Combination	0.9691	0.9941	0.9977	0.9958
ExplAIn-CAM/B1	0.9523	0.9918	0.9970	0.9874
“black-box” AI	0.9702	0.9859	0.9969	0.9969

Performance was evaluated using the accuracy and the quadratic weighted Kappa (κ), a variation on Cohen’s Kappa for ordinal scales, which takes into account the amount of disagreement between observers (Cohen, 1968). The probability cutoffs were chosen to maximize κ on the validation set.

4.6.2. Hyperparameter Optimization

Hyperparameters of each algorithm were optimized at the image level using the validation dataset. A ROC analysis was performed: hyperparameters maximizing the average per-label AUC on the validation subset were selected. For ExplAIn, the following hyperparameters had to be set:

- the number M of pixel-level labels (see section 3.1),
- the CNN backbone (see section 3.2),
- training weights α, β and γ (see section 3.6).

Training weight γ , which trades off image-level and pixel-level classification quality (see section 3.6.5), was empirically set to 0.1. Then, M, α and β were chosen to maximize the average per-backbone and per-label AUC. Finally, the best CNN backbone, given the optimal M, α and β , was selected. A similar procedure was followed for each algorithm (γ or equivalent set to 0.1, CNN backbone chosen last).

4.6.3. Results

The hyperparameters of ExplAIn maximizing image-level classification performance on the validation dataset are as follows:

- $M = 6$ pixel-level labels (see section 3.1),

quadratic weighted Kappa (κ), in various datasets. The “proposed” approach is ExplAIIn with the EfficientNet-B5 backbone and $M = 6$ pixel-level labels. Results of the IDRiD challenge originate from Porwal et al. (2019). Results of the DeepDR challenge are from the challenge website: the 4 best solutions (out of 8) are reported (SJTU: Shanghai Jiao Tong University; CUHK: Chinese University of Hong Kong) — the onsite evaluation dataset is not public yet.

dataset	method	accuracy	κ
<i>IDRiD_{image}</i>	proposed	0.75 (0.7476)	0.8673
	LzyUNCC	0.75	
	VRT	0.59	
	Mammoth	0.54	
	HarangiM1	0.55	
	AVSASVA	0.55	
	HarangiM2	0.48	
<i>DeepDR_{valid}</i> <i>DeepDR_{onsite}</i> ⁹	proposed	0.7300	0.9243
	SJTU (M1)		0.9215
	SJTU (M2)		0.9211
	Vuno Inc.		0.9097
	CUHK		0.8845
<i>OPHDIAT_{eval}</i>	proposed	0.9389	0.8937

- EfficientNet-B5 backbone (see section 3.2),
- training weights $\alpha = \beta = 0.1$ (see section 3.6).

ROC curves obtained with these hyperparameters in the *OPHDIAT_{eval}* dataset are reported in Fig. 2. For comparison purposes, image-level performance achieved with different EfficientNet backbones and with the ExplAIIn-CAM and “black-box” AI baselines is summarized in Table 1. The combination of multiple ExplAIIn models was also investigated: eight models with different backbones (EfficientNet-B0 to -B7) were combined through logistic regression on the validation dataset. Results on *OPHDIAT_{eval}* are reported in Table 1. Combining multiple models improves image-level classification performance. However, the improvement is limited and it comes at the expense of reduced explainability. Therefore, a simple ExplAIIn model (with the optimal hyperparameters) is used in the remainder of the study.

Next, the multiclass classification performance of this solution, on multiple evaluation datasets, is reported in Table 2: ExplAIIn is compared with competing solutions of the IDRiD and DeepDR challenges.

4.7. Pixel-Level Performance

4.7.1. Performance Metrics

Classification performance at the pixel level can also be evaluated by ROC curves. However, because positive and negative pixels are highly unbalanced, a second evaluation method was used: Precision-Recall (PR) curves. PR curves are built similarly to ROC curves, but (recall, precision) pairs are used in place of (1-specificity, sensitivity) pairs; note that recall and sensitivity are synonyms. A PR curve can be summarized by

the mean Average Precision (mAP), defined as the area under the PR curve. Two types of pixel-level evaluations were performed:

lesion-specific: each foreground probability map P_m , $m \in \{2, \dots, M\}$, is compared to each lesion-specific segmentation map. The label $m \in \{2, \dots, M\}$ maximizing the AUC is retained. We measure how well each lesion type can be detected by a foreground probability map.

combined: the combined foreground probability map ($1 - P_1$) is compared to the union of all lesion segmentation maps. Here, we evaluate foreground/background separation.

4.7.2. Qualitative Pixel-Level Evaluation

For improved visualization, pixel labeling can be summarized by color codes. In ExplAIIn, each pixel $I_{x,y}$ is labeled by an M -dimensional vector: $P_{m,x,y}$, $m \in \{1, \dots, M\}$: this vector can be summarized by a 3-dimensional color vector in CIE L*a*b* color space. The lightness component L* represents foreground probability $1 - P_{1,x,y}$. Chromaticity components a* and b* represent the normalized ($N - 1$)-dimensional label vectors $P_{m,x,y}^*$:

$$P_{m,x,y}^* = \left\{ \frac{P_{m,x,y}}{1 - P_{1,x,y}}, \forall m \in \{2, \dots, M\} \right\}. \quad (11)$$

This normalized vector is summarized by a 2-D vector (a*, b*) using isometric mapping (Tenenbaum et al., 2000). Isometric mapping, or Isomap, is a manifold learning algorithm that seeks a lower-dimensional embedding maintaining geodesic distances between all input vectors.

4.7.3. Results

Classification performance at the pixel level in the *IDRiD_{pixel}* evaluation dataset is summarized in Fig. 3 (PR curves) and 4 (ROC curves) for all architectures.

Typical pixel-level classification maps, obtained with the optimal hyperparameters, are reported in Fig. 5. From a visual inspection, it appears that hemorrhages are targeted by map P_3 , exudates are targeted by map P_4 and microaneurysms are targeted by map P_5 . This is confirmed by the lesion-specific quantitative evaluation described in section 4.7.1, the results of which are summarized in Fig. 4 (e). Map P_6 seems to target signs of advanced DR, such as neovascularization. In a few images from the IDRiD dataset, we observe that P_6 also detects laser scars, a possible sign that neovascularization is being treated (IDRiD is not a screening dataset). Unfortunately, advanced DR signs are not manually segmented in IDRiD, so quantitative evaluation is impossible in this case. For improved visualization, color-coded pixel-level classification maps are reported in Fig. 6. For comparison purposes, typical pixel-level classification maps obtained with different EfficientNet backbones and with the baseline methods are reported in Fig. 7.

4.8. Analysis of the Extracted Rules

Pixel-level classification maps P_m have now been associated with lesion types. The next step is to verify that the relationships between these maps and image-level labeling are consistent with the ICDR scale, reported at the beginning of section 4.

⁹<https://isbi.deepdr.org/leaderboard.html>

Table 3 summarizes these relationships: it reports the $w_{m,n}^2$ coefficients, which are multiplied by the maximal $P_{m,x,y}$ values, in order to obtain the image-level probabilities p_n in Eq. (3). We see that the dominant map for detecting ‘at least mild NPDR’ is P_5 , which we associated with microaneurysms. The dominant map for detecting ‘at least moderate NPDR’ is P_4 , which we associated with exudates. The dominant map for detecting ‘at least severe NPDR’ is P_3 , which we associated with hemorrhages. Finally, the dominant map for detecting ‘PDR’ is P_6 , which we associated with advanced DR signs, including neovascularization in particular. Overall, these relationships are consistent with the ICDR scale: the targeted lesions are compatible.

Map P_6 could not be validated at the pixel level, due to lack of manual annotations. However, Table 3 shows that P_6 is predominantly linked to PDR detection. In details, PDR detection relies on P_6 and, to a lesser extent, P_3 . One possible interpretation is that P_3 , which was shown to detect hemorrhages, detects ‘vitreous/preretinal hemorrhages’ and that, therefore, P_6 detects the other sign of PDR, namely neovascularization (Wilkinson et al., 2003).

5. Discussion and Conclusions

We have presented ExplAIIn, a novel eXplanatory Artificial Intelligence (XAI) framework for multilabel image classification. In addition to labeling images, ExplAIIn also classifies each pixel within images. Simple rules link pixel-level classification to image-level labeling. Consequently, image-level labeling can be explained simply by pixel-level classification. Unlike image-level labeling, pixel-level classification is self-supervised; a novel occlusion method is presented to ensure satisfactory foreground/background pixel separation and therefore meaningful explanations. This framework was applied to the diagnosis of Diabetic Retinopathy (DR) using Color Fundus Photography (CFP). Classification performance was evaluated both at the image level and at the pixel level.

ExplAIIn models were trained to classify DR at the image level, using data from French and American DR screening programs (OPHDIAI and EyePACS): the goal was to grade DR severity in one eye according to the ICDR scale. On an Indian dataset (IDRiD), we found that manually-segmented DR lesions (microaneurysms, hemorrhages, soft exudates and hard exudates) could be detected well (see Fig. 4). Overall, these lesions could also be differentiated correctly (see Fig. 5). Using the optimal number of pixel-level labels ($M = 6$) and the optimal CNN backbone (EfficientNet-B5), unsupervised pixel-level classification can be schematized as follows. By design, label $m = 1$ was assigned to the background (see section 3.4). One pixel-level label was assigned to microaneurysms ($m = 5$), another one was assigned to hemorrhages ($m = 3$), another one to exudates ($m = 4$) and another one to advanced DR signs ($m = 6$). The two types of exudates (hard and soft), however, were not separated well using this optimal CNN backbone. Assigning pixel-level labels to lesion types makes sense, since the ICDR scale relies primarily on the types of DR lesions present in images (Wilkinson et al., 2003). Regarding the last label,

$m = 2$, it seems to group together false alarms of the occlusion method: pixels irrelevant for DR classification but nevertheless unusual for background pixels. Table 3 suggests that the learnt relationships between pixel-level classification and image-level labeling is consistent with the ICDR scale. For easier visualization, we propose to analyze all pixel-level labels jointly, using color codes obtained through dimension reduction (see section 4.7.2). One advantage is a more compact representation: one color-coded image is generated per input image, as opposed of $M = 6$ grayscale images. As illustrated in Fig. 7, where different CNN backbones are compared, it appears that changing model hyperparameters does not change image-level classification radically, although large CNN backbones (EfficientNet models B4 and higher) lead to semantically richer pixel-level classification: color codes are more diverse.

Pixel-level ROC curves in Fig. 4 indicate that lesions can be detected with a good pixel-level sensitivity: under-segmentation is limited. On the other hand, pixel-level PR curves in Fig. 3 indicate that pixel-level precision is lower: over-segmentation is more problematic. However, we believe this level of precision is enough for explainability purposes. In any case, ExplAIIn clearly is superior to ExplAIIn-CAM and RRM at the pixel level (see Fig. 7). In particular, precision is about 10 times higher (see Fig. 3). In ExplAIIn, increasing γ , the weight assigned to the $\mathcal{L}_{sparsity}$ loss in Eq. (10), increases pixel-level precision without dramatically impacting image-level classification. On the other hand, enforcing the sparsity of CAM, like in ExplAIIn-CAM, largely impacts image-level classification. This is due to the low resolution of CAM: a sparse CAM implies that large regions of images have to be completely ignored. This suggests that CAM is not precise enough for Weakly-Supervised Semantic Segmentation (WSSS) of DR-related lesions, even though this is currently the most frequently used technique for WSSS (see section 2).

We have shown that ExplAIIn detects and categorizes relevant abnormalities for grading DR severity. But the type of abnormalities is not the only relevant factor. In the ICDR scale, for instance, severe NPDR is only triggered when abnormalities are sufficiently numerous and sufficiently spread across the retina. How can ExplAIIn reproduce these requirements, given that it only looks at the maximal intensity in each pixel-level classification map? Different solutions were witnessed during the course of training. One solution was to assign one pixel-level label per quadrant, in order to capture spatial distribution. Another solution was to focus on large clusters of abnormalities and ignore isolated abnormalities, in an attempt to capture their number. The latter solution was retained in the ExplAIIn-CAM baseline, as illustrated in Fig. 7 (j) and (t). In ExplAIIn, however, none of these solutions was retained at the end of training. Each type of abnormalities is assigned to a single pixel-level label and most of the abnormalities are detected (see Fig. 5): pixel-level probabilities simply tend to increase with the number of abnormalities and with their spread across the image.

Although explainability is a desirable feature, it should not come at the expense of decreased classification performance. Fortunately, for the same task on the same evaluation dataset (OPHDIAI_{eval}), image-level classification performance of Ex-

maximal intensity in each map P_m to the image-level probability p_n , as well as the bias term b_n . For each rule, the dominant coefficient is in bold.

$p_n \setminus P_m$	$w_{m,n}^2$ coefficients					b_n bias
	P_2	P_3	P_4	P_5	P_6	
at least mild NPDR (p_1)	3.009	1.673	1.070	3.892	0.417	-2.765
at least moderate NPDR (p_2)	1.288	2.024	3.159	2.417	0.915	-4.024
at least severe NPDR (p_3)	0.257	4.294	0.137	2.248	1.268	-5.318
PDR (p_4)	0.000	2.141	0.000	0.592	4.235	-5.885

plAIIn (see Fig. 2) is similar to previously reported results for a “black-box” AI solution from our group (Quellec et al., 2019). Clearly, this “black-box” solution was designed without explainability constraints: in particular, it consisted of an ensemble of multiple CNNs while, for explainability purposes, ExplAIIn necessarily consists of a single CNN. In details, two classification criteria have improved in ExplAIIn (see Table 1):

- “at least moderate NPDR”, the most important criterion in screening applications: AUC increases from 0.9859 to 0.9939. Besides, for a 100% sensitivity, specificity increases from 12.5% to 88.1%,
- “at least severe NPDR”: AUC increases from 0.9969 to 0.9978. Besides, for a 100% sensitivity, specificity increases from 92.4% to 95.6%.

For the remaining two criteria, performance in terms of AUC decreases overall. However, the optimal sensitivity/specificity endpoint is improved or unchanged:

- “at least mild NPDR”: sensitivity = 90%, specificity = 93.3% (“black-box” AI) or 94.1% (ExplAIIn),
- “PDR”: sensitivity = 100%, specificity = 98.9% (in both solutions).

We note that ExplAIIn is also more efficient than ExplAIIn-CAM for classification at the image level (see Table 1), although the difference is more subtle than at the pixel level. ExplAIIn was also evaluated on two independent datasets: the Indian IRDiD dataset and the Chinese DeepDR dataset. The proposed algorithm seems to compare favorably with baseline algorithms, even though these baseline algorithms were trained or fine-tuned on non-independent training datasets (see Table 2). Note that the comparison is not direct on DeepDR: the proposed and baseline algorithms were evaluated on slightly different datasets (validation and onsite evaluation datasets). As shown in Table 2, the accuracy of ExplAIIn varies across evaluation datasets. These variations can be explained by several factors: a different distribution of DR severity levels, different annotators (the OPHDIAT evaluation images were read by at least two graders), different population characteristics, different acquisition devices, etc. Variations in terms of quadratic weighted Kappa (κ) are more limited. Importantly, a large κ value was measured in all evaluation datasets considered in this paper ($0.8673 \leq \kappa \leq 0.9243$), indicating that disagreements between the AI and experts generally have a limited amplitude.

The very good classification performance of ExplAIIn at the image level was somehow unexpected. One possible explana-

tion for this success is that the proposed generalized occlusion method acts as a regularization operator. Enforcing foreground sparsity encourages the AI to look for local patterns in images rather than to analyze the image globally. More generally, by solving a multi-task problem (image-level and pixel-level classification), we favor the extraction of more general and relevant features.

In this paper, ExplAIIn has been applied to a well-known classification problem (DR severity classification). It allowed us to evaluate the relevance of the identified local patterns (namely DR lesions). However, ExplAIIn would be even more useful for a totally new classification problem (disease progression prediction, diagnosis of a new disease, etc.): it would allow knowledge acquisition. In medicine, for instance, it would help clinicians quickly identify new useful markers in images.

One limitation of the proposed solution for foreground/background separation (the generalized occlusion method) is that it is limited to binary or multilabel classification (zero, one or multiple labels per image): it is not directly applicable to multiclass classification (exactly one label per image). This is because it relies on the concept of “background image”, which does not generally apply to multiclass classification. This concept may apply to a subclass of multiclass classification problems where one of the classes corresponds to background images (with no relevant patterns visible); however, the occlusion loss would need a different formulation. Finally, the validation of ExplAIIn on the DR grading task also has one limitation: pixel-level evaluation relies on incomplete lesion segmentations. Only four types of lesions were manually-segmented in IDRiD: microaneurysms, hemorrhages, soft exudates and hard exudates (Porwal et al., 2019). This results in pessimistic precisions in Fig. 3.

In conclusion, we have presented a novel explanatory AI framework for multilabel image classification. Given an image, this framework 1) localizes and categorizes relevant patterns in the image, 2) classifies the image and 3) explains how exactly each pattern contributes to the classification. If experts can assign keywords to these patterns, then a textual report can also be generated. For the task of DR diagnosis using CFP, the new explainability feature comes without loss of classification performance. Thanks to this new feature, we expect healthcare AI systems to gain the trust of clinicians and patients more easily, which would facilitate their deployment.

References

- Abràmoff, M.D., Lou, Y., Erginay, A., Clarida, W., Amelon, R., Folk, J.C., Niemeijer, M., 2016. Improved automated detection of diabetic retinopathy on a publicly available dataset through integration of deep learning. *Invest Ophthalmol Vis Sci* 57, 5200–5206. doi:10.1167/iovs.16-19964.
- Abràmoff, M.D., Lavin, P.T., Birch, M., Shah, N., Folk, J.C., 2018. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *npj Digital Med* 1, 39. doi:10.1038/s41746-018-0040-6.
- Ahn, J., Cho, S., Kwak, S., 2019. Weakly supervised learning of instance segmentation with inter-pixel relations, in: *Proc IEEE CVPR, Long Beach, CA, USA*. pp. 2204–2213. doi:10.1109/CVPR.2019.00231.
- Amann, J., Blasimme, A., Vayena, E., Frey, D., Madai, V.I., the Precise4Q consortium, 2020. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Medical Informatics and Decision Making* 20, 310. doi:10.1186/s12911-020-01332-6.
- Araújo, T., Aresta, G., Mendonça, L., Penas, S., Maia, C., Carneiro, A., Mendonça, A.M., Campilho, A., 2020. DR|GRADUATE: Uncertainty-aware deep learning-based diabetic retinopathy grading in eye fundus images. *Med Image Anal* 63, 101715. doi:10.1016/j.media.2020.101715.
- Arcadu, F., Benmansour, F., Maunz, A., Willis, J., Haskova, Z., Prunotto, M., 2019. Deep learning algorithm predicts diabetic retinopathy progression in individual patients. *NPJ Digit Med* 2, 1–9. doi:10.1038/s41746-019-0172-3.
- Ayhan, M.S., Kühlewein, L., Aliyeva, G., Inhoffen, W., Ziemssen, F., Berens, P., 2020. Expert-validated estimation of diagnostic uncertainty for deep neural networks in diabetic retinopathy detection. *Med Image Anal* 64, 101724. doi:10.1016/j.media.2020.101724.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W., 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One* 10. doi:10.1371/journal.pone.0130140.
- Benoit, S.R., Swenor, B., Geiss, L.S., Gregg, E.W., Saaddine, J.B., 2019. Eye care utilization among insured people with diabetes in the U.S., 2010–2014. *Diabetes Care* 42, 427–433. doi:10.2337/dc18-0828.
- Chan, L., Hosseini, M.S., Plataniotis, K.N., 2019. A comprehensive analysis of weakly-supervised semantic segmentation in different image domains. Technical Report arXiv:1912.11186 [cs].
- Cohen, J., 1968. Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit. *Psychol Bull* 70, 213–220. doi:10.1037/h0026256.
- Conze, P.H., Kavur, A.E., Gall, E.C.L., Gezer, N.S., Meur, Y.L., Selver, M.A., Rousseau, F., 2020. Abdominal multi-organ segmentation with cascaded convolutional and adversarial deep networks. Technical Report arXiv:2001.09521 [cs, eess].
- Cuadros, J., Bresnick, G., 2009. EyePACS: an adaptable telemedicine system for diabetic retinopathy screening. *J Diabetes Sci Technol* 3, 509–516.
- Dubost, F., Adams, H., Yilmaz, P., Bortsova, G., Tulder, G., Ikram, M., Niessen, W., Vernooij, M., de Bruijne, M., 2020. Weakly supervised object detection with 2D and 3D regression neural networks. *Med Image Anal* 65, 101767. doi:10.1016/j.media.2020.101767.
- Durand, T., Mordan, T., Thome, N., Cord, M., 2017. WILDCAT: Weakly supervised learning of deep ConvNets for image classification, pointwise localization and segmentation, in: *Proc IEEE CVPR, Honolulu, HI, USA*. pp. 5957–5966. doi:10.1109/CVPR.2017.631.
- Fauw, J.D., Ledsam, J.R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., Askham, H., Glorot, X., ODonoghue, B., Visentin, D., Driessche, G.v.d., Lakshminarayanan, B., Meyer, C., Mackinder, F., Bouton, S., Ayoub, K., Chopra, R., King, D., Karthikesalingam, A., Hughes, C.O., Raine, R., Hughes, J., Sim, D.A., Egan, C., Tufail, A., Montgomery, H., Hassabis, D., Rees, G., Back, T., Khaw, P.T., Suleyman, M., Cornebise, J., Keane, P.A., Ronneberger, O., 2018. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat Med* 24, 1342. doi:10.1038/s41591-018-0107-6.
- Frosst, N., Hinton, G., 2017. Distilling a neural network into a soft decision tree, in: *Proc AI*IA Works CEX, Bari, Italy*.
- Gilpin, L.H., Bau, D., Yuan, B.Z., Bajwa, A., Specter, M., Kagal, L., 2018. Explaining explanations: An overview of interpretability of machine learning, in: *Proc IEEE DSAA, Turin, Italy*. pp. 80–89. doi:10.1109/DSAA.2018.00018.
- Goodman, B., Flaxman, S., 2017. European Union regulations on algorithmic decision-making and a right to explanation. *AI Magazine* 38, 50–57. doi:10.1609/aimag.v38i3.2741.
- Hehn, T.M., Kooij, J.F.P., Hamprecht, F.A., 2020. End-to-end learning of decision trees and forests. *Int J Comput Vis* 128, 997–1011. doi:10.1007/s11263-019-01237-6.
- Javitt, J.C., Aiello, L.P., 1996. Cost-effectiveness of detecting and treating diabetic retinopathy. *Ann Intern Med* 124, 164–169.
- Jiang, P.T., Hou, Q., Cao, Y., Cheng, M.M., Wei, Y., Xiong, H., 2019. Integral object mining via online attention accumulation, in: *Proc IEEE ICCV, Seoul, Korea*. pp. 2070–2079. doi:10.1109/ICCV.2019.00216.
- Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., Ayed, I.B., 2019. Constrained-CNN losses for weakly supervised segmentation. *Medical Image Analysis* 54, 88–99. doi:10.1016/j.media.2019.02.009.
- Kolesnikov, A., Lampert, C.H., 2016. Improving weakly-supervised object localization by micro-annotation. Technical Report arXiv:1605.05538 [cs].
- Kontschieder, P., Fiterau, M., Criminisi, A., Bul, S.R., 2016. Deep neural decision forests, in: *Proc IJCAI, New York, NY, USA*. pp. 4190–4194.
- Kwak, S., Hong, S., Han, B., 2017. Weakly supervised semantic segmentation using superpixel pooling network, in: *Proc AAAI, San Francisco, CA, USA*. pp. 4111–4117.
- Laradji, I.H., Vazquez, D., Schmidt, M., 2019. Where are the masks: Instance segmentation with image-level supervision, in: *Proc BMVC, Cardiff, UK*.
- Lee, J., Kim, E., Lee, S., Lee, J., Yoon, S., 2019. FickleNet: Weakly and semi-supervised semantic image segmentation using stochastic inference, in: *Proc IEEE CVPR, Long Beach, CA, USA*. pp. 5262–5271. doi:10.1109/CVPR.2019.00541.
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature pyramid networks for object detection, in: *Proc CVPR*, pp. 936–944. doi:10.1109/CVPR.2017.106.
- Massin, P., Chabouis, A., Erginay, A., Viens-Bitker, C., Lecleire-Collet, A., Meas, T., Guillausseau, P.J., Choupot, G., Andr, B., Denormandie, P., 2008. OPHDIAT: a telemedical network screening system for diabetic retinopathy in the Ile-de-France. *Diabetes Metab* 34, 227–234. doi:10.1016/j.diabet.2007.12.006.
- Ogurtsova, K., da Rocha Fernandes, J.D., Huang, Y., Linnenkamp, U., Guariguata, L., Cho, N.H., Cavan, D., Shaw, J.E., Makaroff, L.E., 2017. IDF Diabetes Atlas: Global estimates for the prevalence of diabetes for 2015 and 2040. *Diabetes Res Clin Pract* 128, 40–50. doi:10.1016/j.diabres.2017.03.024.
- Papandreou, G., Chen, L.C., Murphy, K., Yuille, A.L., 2015. Weakly- and semi-supervised learning of a DCNN for semantic image segmentation. Technical Report arXiv:1502.02734 [cs].
- Pathak, D., Krahenbuhl, P., Darrell, T., 2015. Constrained convolutional neural networks for weakly supervised segmentation, in: *Proc IEEE ICCV, Washington, DC, USA*. pp. 1796–1804. doi:10.1109/ICCV.2015.209.
- Porwal, P., Pachade, S., Kokare, M., Deshmukh, G., Son, J., Bae, W., Liu, L., Wang, J., Liu, X., Gao, L., Wu, T., Xiao, J., Wang, F., Yin, B., Wang, Y., Danala, G., He, L., Choi, Y.H., Lee, Y.C., Jung, S.H., Li, Z., Sui, X., Wu, J., Li, X., Zhou, T., Tth, J., Baran, A., Kori, A., Chennamsetty, S.S., Safwan, M., Alex, V., Lyu, X., Cheng, L., Chu, Q., Li, P., Ji, X., Zhang, S., Shen, Y., Dai, L., Saha, O., Sathish, R., Melo, T., Arajo, T., Harangi, B., Sheng, B., Fang, R., Sheet, D., Hajdu, A., Zheng, Y., Mendona, A.M., Zhang, S., Campilho, A., Zheng, B., Shen, D., Giancardo, L., Quelled, G., Mériaudeau, F., 2019. IDRiD: Diabetic retinopathy segmentation and grading challenge. *Medical Image Analysis*, 101561doi:10.1016/j.media.2019.101561.
- Quelled, G., Charrière, K., Boudi, Y., Cochener, B., Lamard, M., 2017. Deep image mining for diabetic retinopathy screening. *Med Image Anal* 39, 178–193. doi:10.1016/j.media.2017.04.012.
- Quelled, G., Lamard, M., Lay, B., Le Guilcher, A., Erginay, A., Cochener, B., Massin, P., 2019. Instant automatic diagnosis of diabetic retinopathy. Technical Report arXiv:1906.11875 [cs, eess].
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation, in: *Proc MICCAI, Munich, Germany*. pp. 234–241. doi:10.1007/978-3-319-24574-4_28.
- Russell, S., Dewey, D., Tegmark, M., 2015. Research priorities for robust and beneficial artificial intelligence. *AI Magazine* 36, 105–114. doi:10.1609/aimag.v36i4.2577.
- Samek, W., Binder, A., Montavon, G., Lapuschkin, S., Müller, K.R., 2017. Evaluating the visualization of what a deep neural network has learned. *IEEE Trans Neural Netw Learn Syst* 28, 2660–2673. doi:10.1109/TNNLS.2016.2599820.
- Sayres, R., Taly, A., Rahimy, E., Blumer, K., Coz, D., Hammel, N., Krause,

- J., Narayanaswamy, A., Rastegar, Z., Wu, D., Xu, S., Barb, S., Joseph, A., Shumski, M., Smith, J., Sood, A.B., Corrado, G.S., Peng, L., Webster, D.R., 2019. Using a deep learning algorithm and integrated gradients explanation to assist grading for diabetic retinopathy. *Ophthalmology* 126, 552–564. doi:10.1016/j.ophtha.2018.11.016.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-CAM: Visual explanations from deep networks via gradient-based localization, in: Proc IEEE ICCV, Venice, Italy. pp. 618–626. doi:10.1109/ICCV.2017.74.
- Shimoda, W., Yanai, K., 2016. Distinct class-specific saliency maps for weakly supervised semantic segmentation, in: Proc ECCV, Springer International Publishing, Amsterdam, The Netherlands. pp. 218–234. doi:10.1007/978-3-319-46493-0_14.
- Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps, in: ICLR Workshop, Calgary, Canada.
- Tan, M., Le, Q.V., 2019. EfficientNet: Rethinking model scaling for convolutional neural networks, in: Proc ICML, Long Beach, CA, USA.
- Tenenbaum, J.B., de Silva, V., Langford, J.C., 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–2323. doi:10.1126/science.290.5500.2319.
- Ting, D.S.W., Pasquale, L.R., Peng, L., Campbell, J.P., Lee, A.Y., Raman, R., Tan, G.S.W., Schmetterer, L., Keane, P.A., Wong, T.Y., 2019a. Artificial intelligence and deep learning in ophthalmology. *Br J Ophthalmol* 103, 167–175. doi:10.1136/bjophthalmol-2018-313173.
- Ting, D.S.W., Peng, L., Varadarajan, A.V., Keane, P.A., Burlina, P., Chiang, M.F., Schmetterer, L., Pasquale, L.R., Bressler, N.M., Webster, D.R., Abramoff, M., Wong, T.Y., 2019b. Deep learning in ophthalmology: The technical and clinical considerations. *Prog Retin Eye Res* 72, 100759. doi:10.1016/j.preteyeres.2019.04.003.
- Wang, Y., Zhang, J., Kan, M., Shan, S., Chen, X., 2020. Self-supervised Scale Equivariant Network for Weakly Supervised Semantic Segmentation, in: Proc IEEE CVPR.
- Wickström, K., Kampffmeyer, M., Jenssen, R., 2020. Uncertainty and interpretability in convolutional neural networks for semantic segmentation of colorectal polyps. *Med Image Anal* 60, 101619. doi:10.1016/j.media.2019.101619.
- Wilkinson, C.P., Ferris, F.L., Klein, R.E., Lee, P.P., Agardh, C.D., Davis, M., Dills, D., Kampik, A., Pararajasegaram, R., Verdager, J.T., 2003. Proposed international clinical diabetic retinopathy and diabetic macular edema disease severity scales. *Ophthalmology* 110, 1677–1682. doi:10.1016/S0161-6420(03)00475-5.
- Xie, S., Tu, Z., 2015. Holistically-nested edge detection, in: Proc IEEE ICCV, pp. 1395–1403. doi:10.1109/ICCV.2015.164.
- Yang, Y., Morillo, I.G., Hospedales, T.M., 2018. Deep neural decision trees, in: Proc ICML Works WHI, Stockholm, Sweden. pp. 34–40.
- Yau, J.W.Y., Rogers, S.L., Kawasaki, R., Lamoureux, E.L., Kowalski, J.W., Bek, T., Chen, S.J., Dekker, J.M., Fletcher, A., Grauslund, J., Haffner, S., Hamman, R.F., Ikram, M.K., Kayama, T., Klein, B.E.K., Klein, R., Krishniah, S., Mayurasakorn, K., O'Hare, J.P., Orchard, T.J., Porta, M., Rema, M., Roy, M.S., Sharma, T., Shaw, J., Taylor, H., Tielsch, J.M., Varma, R., Wang, J.J., Wang, N., West, S., Xu, L., Yasuda, M., Zhang, X., Mitchell, P., Wong, T.Y., Group, f.t.M.A.f.E.D.M.E.S., 2012. Global prevalence and major risk factors of diabetic retinopathy. *Diabetes Care* 35, 556–564. doi:10.2337/dc11-1909.
- Zeiler, M.D., Fergus, R., 2014. Visualizing and understanding convolutional networks, in: Proc ECCV, Zurich, Switzerland. pp. 818–833.
- Zhang, B., Xiao, J., Wei, Y., Sun, M., Huang, K., 2020. Reliability does matter: An end-to-end weakly supervised semantic segmentation approach, in: Proc AAAI, New York, NY, USA.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization, in: Proc IEEE CVPR, Las Vegas, NV, USA. pp. 2921–2929. doi:10.1109/CVPR.2016.319.
- Zhou, Y., Zhu, Y., Ye, Q., Qiu, Q., Jiao, J., 2018. Weakly supervised instance segmentation using class peak response, in: Proc IEEE CVPR, Salt Lake City, UT, USA. pp. 3791–3800. doi:10.1109/CVPR.2018.00399.

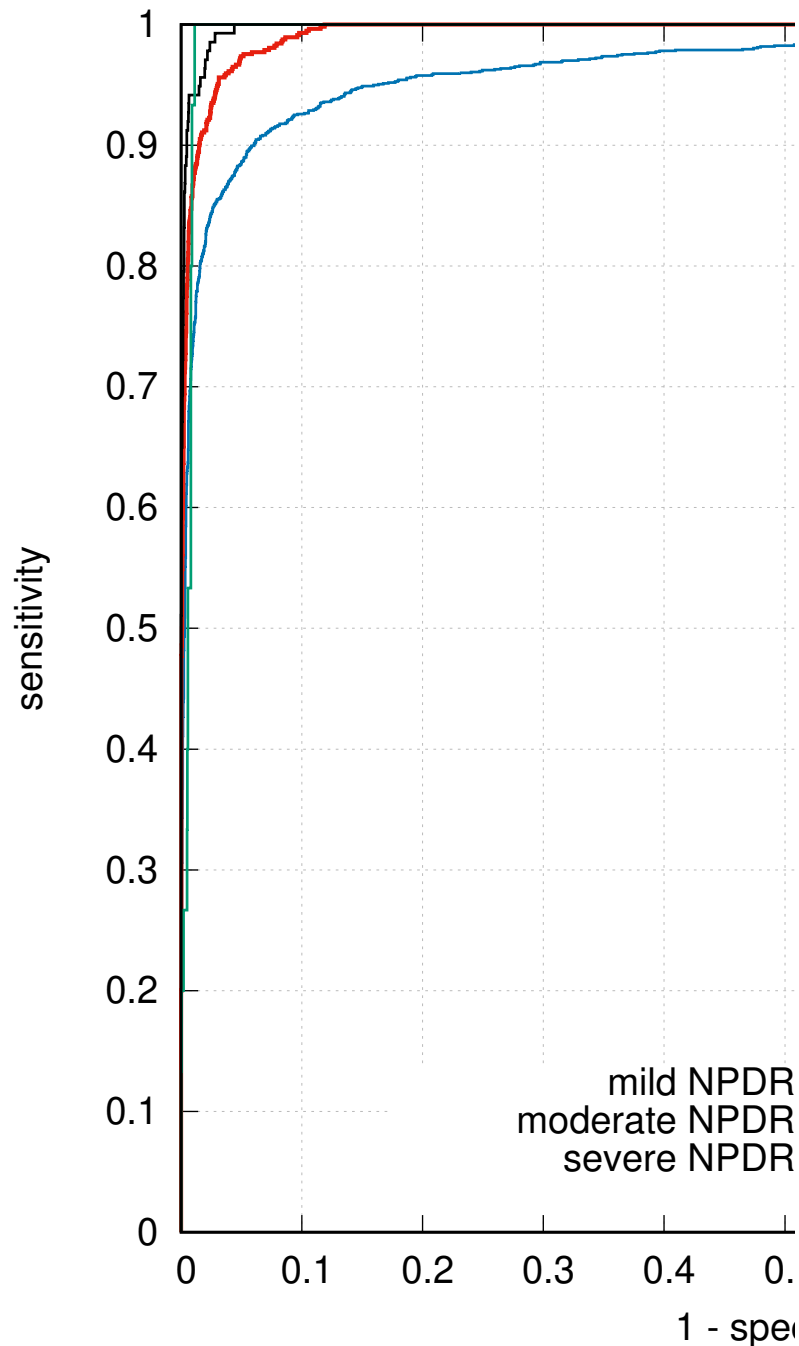


Fig. 2: Receiver-operating characteristics for DR severity assessment, i.e. image-level classification, in the OPHDIAT (Massin et al., 2008) evaluation dataset, using the EfficientNet-B5 backbone and $M = 6$ pixel-level labels.




Fig. 3: Precision-recall analysis for DR lesion detection, i.e. pixel-level classification, in the IDRiD (Porwal et al., 2019) evaluation dataset. The “/Bx” suffix in the architecture names indicates the EfficientNet-“Bx” model used as backbone.

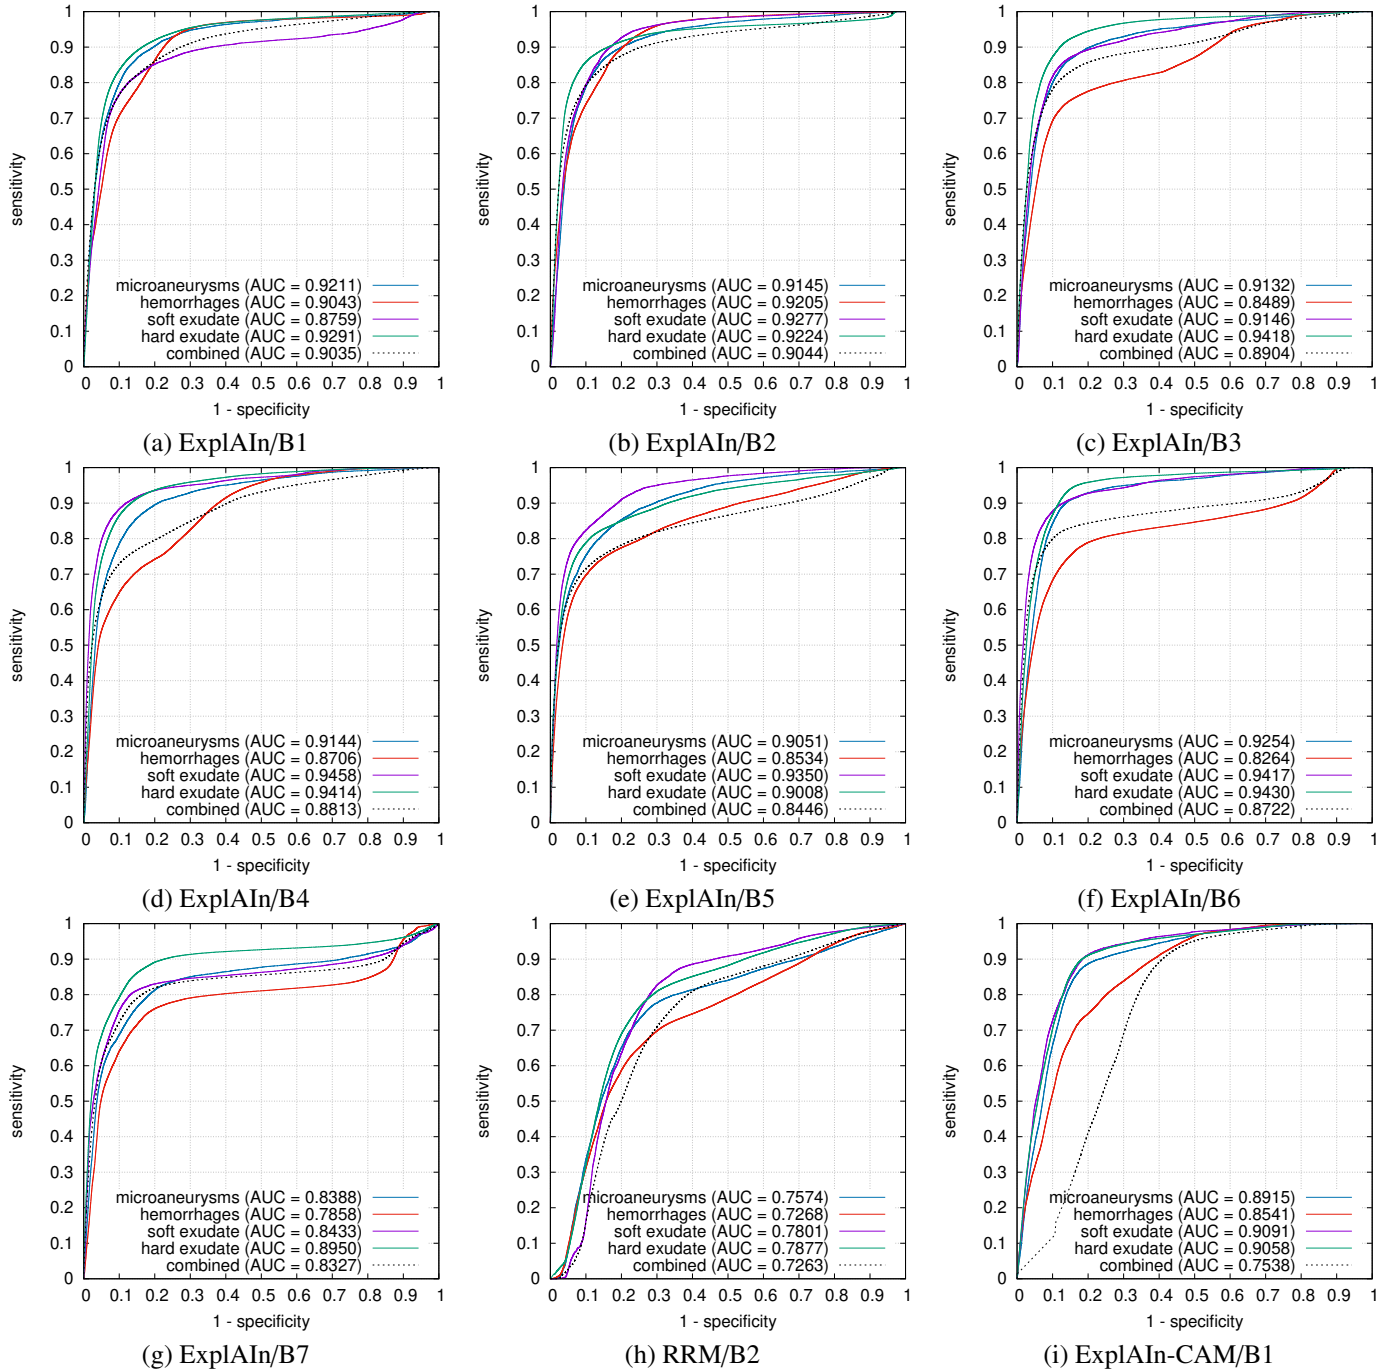


Fig. 4: Receiver operating characteristics for DR lesion detection, i.e. pixel-level classification, in the IDRiD evaluation dataset. The “/Bx” suffix in the architecture names indicates the EfficientNet-“Bx” model used as backbone.




Fig. 5: Pixel-level predictions for images from the IDRiD evaluation dataset, using [the EfficientNet-B5 backbone](#) and $M = 6$ pixel-level labels. For each pre-processed image (I), manual segmentations (hemorrhages, exudates and microaneurysms) are listed below, and the $M - 1$ foreground label maps (P_2 , P_3 , P_4 , P_5 and P_6) are listed on the right.




Fig. 6: Color-coded pixel-level predictions for images from the IDRiD evaluation dataset, using the EfficientNet-B5 backbone and $M = 6$ pixel-level labels. For each pre-processed image on the left, Color-coded pixel-level predictions (see section 4.7.2) are given on the right.





Fig. 7: Color-coded pixel-level predictions for two images from the IDRiD evaluation dataset using various architectures.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have ap-

peared to influence the work reported in this paper.

Graphical Abstract



ExplAIn: Explanatory Artificial Intelligence

