



Normalization of series of fundus images to monitor the geographic atrophy growth in dry age-related macular degeneration

Florence Rossant, Michel Paques

► To cite this version:

Florence Rossant, Michel Paques. Normalization of series of fundus images to monitor the geographic atrophy growth in dry age-related macular degeneration. *Computer Methods and Programs in Biomedicine*, 2021, 208, 10.1016/j.cmpb.2021.106234 . hal-03155715v2

HAL Id: hal-03155715

<https://hal.science/hal-03155715v2>

Submitted on 30 Jun 2021

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.

Normalization of series of fundus images to monitor the geographic atrophy growth in dry age-related macular degeneration

Florence Rossant, Michel Paques

► To cite this version:

Florence Rossant, Michel Paques. Normalization of series of fundus images to monitor the geographic atrophy growth in dry age-related macular degeneration. Computer Methods and Programs in Biomedicine, Elsevier, In press. hal-03155715v2

HAL Id: hal-03155715

<https://hal.archives-ouvertes.fr/hal-03155715v2>

Submitted on 30 Jun 2021

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.

Normalization of series of fundus images to monitor the geographic atrophy growth in dry age-related macular degeneration

Florence Rossant^a and Michel Paques^b

^aISEP, Institut Supérieur d'Electronique de Paris, 10 rue de Vanves, 92130 Issy-les-Moulineaux, France, florence.rossant@isep.fr

^bClinical Investigation Center 1423, Quinze-Vingts Hospital, 28 rue de Charenton, 75012 Paris, France

ABSTRACT

Background and Objective: Age-related macular degeneration (ARMD) is a degenerative disease that affects the retina, and the leading cause of visual loss. In its dry form, the pathology is characterized by the progressive, centrifugal expansion of retinal lesions, called geographic atrophy (GA). In infrared eye fundus images, the GA appears as localized bright areas and its growth can be observed in series of images acquired at regular time intervals. However, illumination distortions between the images make impossible the direct comparison of intensities in order to study the GA progress. Here, we propose a new method to compensate for illumination distortion between images.

Methods: We process all images of the series so that any two images have comparable gray levels. Our approach relies on an illumination/reflectance model. We first estimate the pixel-wise illumination ratio between any two images of the series, in a recursive way; then we correct each image against all the others, based on those estimates. The algorithm is applied on a sliding temporal window to cope with large changes in reflectance. We also propose morphological processing to suppress illumination artefacts.

Results: The corrected illumination function is homogeneous in the series, enabling the direct comparison of grey-levels intensities in each pixel, and so the detection of the GA growth between any two images. To demonstrate that, we present numerous experiments performed on a dataset of 18 series (328 images), manually segmented by an ophthalmologist. First, we show that the normalization preprocessing dramatically increases the contrast of the GA growth areas. Secondly, we apply segmentation algorithms derived from Otsu's thresholding to detect automatically the GA total growth and the GA progress between consecutive images. We demonstrate qualitatively and quantitatively that these algorithms, although fully automatic, unsupervised and basic, already lead to interesting segmentation results when applied to the normalized images. Colored maps representing the GA evolution can be derived from the segmentations.

Conclusion: To our knowledge, the proposed method is the first one which corrects automatically and jointly the illumination inhomogeneity in a series of fundus images, regardless of the number of images, the size, shape and progression of lesion areas. This algorithm greatly facilitates the visual interpretation by the medical expert. It opens up the possibility of treating automatically each series as a whole (not just in pairs of images) to model the GA growth.

Keywords – eye fundus images, normalization of series of images, dry age-related macular degeneration (ARMD), geographic atrophy (GA), GA progression detection, GA progression representation.

1. Introduction

Age-related macular degeneration (ARMD) is the leading cause of visual loss in our countries. It affects the retina, the photosensitive layer located in the back of the eye. While the so-called "wet" form of ARMD can be managed by a specific treatment, there is currently no therapeutic options for its "dry" form. The latter is characterized by the progressive, centrifugal expansion of retinal degeneration, giving rise to what has been compared to a geographic-like lesion, hence the name geographic atrophy (GA). The intimate mechanisms of the growth of the lesions remain elusive.

In clinical routine, the diagnosis of ARMD requires optical coherence tomography (OCT) since this modality is very sensitive to reveal small atrophic areas [15]. However, our goal is not to diagnose but to analyze the GA progress. For that reason, we use confocal scanning laser ophthalmoscopy (cSLO), the resolution of which is higher than that of en-face OCT (limited to the interscan spacing, typically $50\mu\text{m}$). Moreover, in OCT, the delineation of the limits of atrophy can most conveniently be done when the scan is perpendicular to the border, while cSLO imaging does not have such constraint.

The disease is visible under the form of a localized change in the coloration of the fundus. Indeed, the affected layer is called the retinal pigment epithelium, which contains a pigment, melanin. Hence the disappearance of this layer causes a change in the coloration of the fundus, which can be observed most specifically in fundus imaging using long wavelengths such as red or infrared (IR) light (Fig. 1). This is why infrared imaging of the fundus through confocal scanning laser ophthalmoscopy (cSLO) is being widely used. This modality is comfortable for the patient, more robust and less invasive than fundus autofluorescence (FAF) imaging; it has higher resolution and higher contrast between normal and diseased areas than color imaging, an older technology much less used nowadays. In addition, the color photo is more often hampered by cataracts than infrared light. For all these reasons, we process cSLO images in IR in this project.

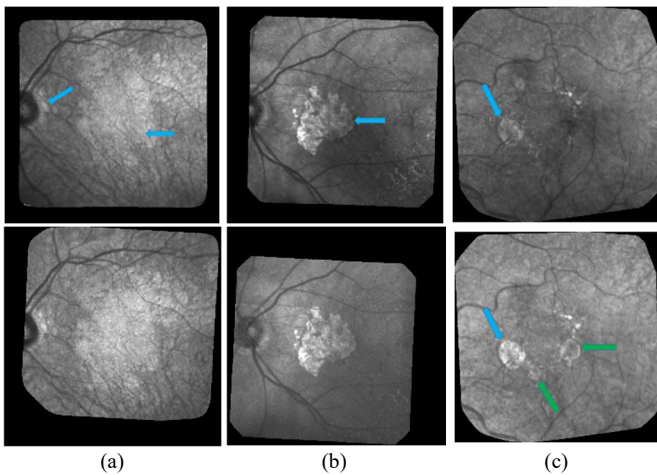


Fig. 1. Examples of cSLO fundus images, acquired in infrared from three patients with dry ARMD. Each column shows two consecutive images extracted from the series. The arrows point to the GA (green arrows to point to new GA zones, blue arrows otherwise).

Fig. 1 shows three pairs of consecutive images, taken at an interval of 6 months. The GA appears as brighter regions in the macula and around the optical disk. Automatic processing to follow up these areas is obviously very challenging given the quality of the images: uneven illumination inside images (a,b,c), even saturation to black or white, illumination distortion between

images (a,b,c), GA poorly contrasted (a) with retinal structures interfering (vessel, optical disk), blur, etc. The difficulty relies also in the high variability of the lesions in shape, size and number. The lesion boundary is quite smooth in some cases (b) and very irregular in others (a). At any time, new spots can appear (c) while older lesions can merge. All these features make the interpretation, in terms of GA delineation and growth, very difficult. On one hand, the segmentation task of the GA in every image is very challenging, especially long and tedious to perform manually, and even experts cannot always be sure of their manual delineation. On the other hand, differential analysis between images cannot be applied without first compensating for illumination distortion.

So, our objective is to automatically process series of eye fundus images acquired in infrared at regular time intervals, in order to detect the appearance of new atrophic areas and quantify the GA growth. The ultimate goal is to better understand the lesion evolution mechanisms and to propose predictive models of the disease progress. This article deals with the first step of this project, image preprocessing: the normalization of intensities between the images, which is a prerequisite to perform differential analysis, and the attenuation of artefacts thanks to morphological processing (Section 4). Experiments (Section 5) on a set of 18 series (328 images manually segmented by a medical expert, Section 3) enable us to evaluate the benefits of the proposed processing: first, in terms of homogeneity between images and of contrast between healthy and unhealthy areas; secondly in terms of ability to segment the GA growth. Concerning the second aspect, we applied standard pixel-wise, unsupervised and fully automatic algorithms, all based on Otsu's threshold, to show the advantages of the preprocessing method for interpreting variations in intensity along the time axis, and ultimately monitoring the GA growth. All these results are discussed in Section 6 before conclusion (Section 7).

2. State of the art

Automatic analysis of fundus images with dry ARMD has been an important research field for two decades, for diagnosis [11] or follow up [6] purposes. Ophthalmologists can observe pathologic features such as Drusen that occur in early stages of the ARMD, and GA progression at different stages of degeneration. This review covers the various aspects of the research that has been carried out so far on fundus images with the aim of automatically analyzing ARMD lesions. It includes the main modalities of eye fundus images in cSLO (color, IR, FAF).

In the literature, most works deal with segmentation of single images. Standard processing methods are applied, such as region growing [6], region oriented variational methods with level set implementation [3][8], texture analysis [7]. To cope with GA variability, Köse proposed to first segment all healthy regions to get the GA as the remaining areas [7]. This approach requires segmenting separately the blood vessels, which is known to be also a difficult task. The method involves many steps and parameters that can be monitored by an user. Besides, several researchers acknowledged the difficulty to design fully automatic tools reaching the required level of performance and they developed user-guided segmentation frameworks [1][9]. Machine learning has also been experimented, with random forest [2] or k-nearest neighbor classifiers [4]. Feature vectors include image intensity, local energy, texture descriptors, values derived from multi-scale analysis and distance to the image center. Nevertheless, these algorithms are supervised, they require training the classifier from annotated data, which brings us back to the difficulty of manually segmenting GA areas. On the contrary, Ramsey's algorithm is unsupervised, based on fuzzy c-means clustering [12]. However, if the method reaches good performances for FAF images, it

performs less well for color fundus photographs. In all these works, few preprocessing algorithms are applied prior to GA segmentation. Ramsey filtered the vessels in FAF images, since they interfere with the GA, thanks to a wavelet-based method [12]. Uneven illumination that occurs often in eye fundus images has been only managed in [2] and [8]. In [2], the author dealt with artifactual horizontal intensity variation by correcting the mean intensity of every image column with respect to the mean intensity of the entire image. In [8], Lee estimated the background illumination of FAF images with a non-linear adaptive smoothing operator and corrected it.

The difference between the segmentations obtained in consecutive images gives the GA growth. However, the previous study shows the high difficulty in segmenting accurately the GA without training or user intervention. That is why a few works proposed another approach, based on differential analysis, to catch the GA growth between two images. Troglio proposed an automatic algorithm to detect changes between two color fundus images, based on Kittler and Illingworth (K&I) method [13]. This method is unsupervised and allows to label pixels in two classes, “change” and “no change” in a Bayesian framework. The non-uniform illumination is corrected in each image using a homomorphic filtering technic, where the image intensity is modeled as a product of a luminance component and a reflectance component. The images are not corrected with respect to each other. Residual illumination variations are taken into account by applying the K&I algorithm in sub-windows and adopting a multiple classifier voting approach. The lack of joint processing induces extra complexity and computational cost. In contrast, Marrugo proposed to correct the images by pairs, by multiplying the second image by a polynomial surface of degree 4, whose parameters are estimated in the least-squares sense [10]. In this way, illumination distortion is compensated and the image difference enhances the areas of changes, which are segmented via a statistical test applied locally. This gives a mask of the change-free areas, which are used to estimate the point spread function (PSF) in each image and apply deconvolution. The illumination compensation method is limited to pairs of images and cannot process all of the images in a series jointly. On the contrary, the blind deconvolution can be extended and is very interesting to compensate for blur present in some eye fundus images.

This review shows that the issues mentioned in the introduction are not yet solved. To the best of our knowledge, the differential approach has not yet been fully explored. Especially, images have been processed in pairs and the proposed algorithms do not deal with large series. However, the differential analysis seems more promising than the segmentation approach, since it may enable us to better handle the high shape variability of the GA areas and the interference with retinal structures, as well as the lack of contrast. This requires designing 1) an efficient illumination normalization algorithm and 2) a differential algorithm robust to residual non-structural changes. The first step, which is the main contribution of this article, is crucial to enable the accurate segmentation of subtle and significant changes between images of the series. As highlighted by our review, the problem of illumination normalization has not been fully solved yet: solutions have been provided to correct uneven background illumination in single images [2][8][13] and joint processing is limited to pairs [10].

So, this article focuses on the joint processing of the whole series, to obtain comparable intensities in the whole stack of images and attenuate illumination artefacts (Section 4). This result is interesting on its own since it facilitates the visual interpretation by the medical expert. We also present many experiments assessing the benefit of the proposed image processing: in terms

of enhancement of GA growth areas and in terms of ability to segment these areas, even with basic automatic and unsupervised approaches (Sections 5). Our methods and results are discussed in Section 6. But before, we describe our dataset and the manually segmented subset used for the evaluation (Section 3).

3. Dataset

Our images were all acquired at the Quinze-Vingts National Ophthalmology Hospital in Paris, in cSLO with IR illumination. The study was carried out according to the tenets of the Declaration of Helsinki and followed international ethical requirements. Informed consent was obtained from all patients. These patients have been followed-up during a few years, hence we have series of retinal fundus images, often for both eyes, showing the progression of the GA. The images were acquired at a time interval of about six months. The average number of images in one series is 13 with a standard deviation of 12. All pictures are in grayscale and vary greatly in size, but the most common size is 650 x 650 pixels. As mentioned previously, we notice many imperfections such as blur, artifacts and, above all, non-uniform illumination inside the images and between them (Fig. 1). All images were spatially aligned with i2k software. In every image, the area of useful data does not cover the entire image and is surrounded by black borders. The automatic detection of these black zones in each image gives a mask of the useful data, and the intersection of all masks the common retinal region where changes can be searched for.

An ophthalmologist segmented manually a sub-database of eighteen series, for the quantitative evaluation of the proposed processing methods (Table 1). These series feature different characteristics in terms of disease progress, lesion shape, size, image quality. For 3 patients, we have series for both left and right eyes. The number of images per series ranges from 5 to 52 (18.2 ± 15.9 in average), with a total of 328 segmented images. We have images at the very beginning of the disease for 5 series. In this case, there is no apparent GA in the first image(s).

P. ID	Eye	N	P. ID	Eye	N	P. ID	Eye	N
001	R	24	008	L	47	089	R	10
001	L	15	010	R	6	109	R	16
003	R	5	016	L	31	110	L	11
005	R	9	018	R	9	112	R	8
005	L	9	019	R	11	115	R	9
008	R	50	020	L	52	117	L	6

Table 1. Sub-database used in the quantitative evaluation Patient ID, processed eye, left (L) or right (R), number of images in the series (N). Background in light gray when the acquisitions start with no GA (5 series).

We developed several user-guided segmentation tools to make the ground truth, based on classical segmentation algorithms: thresholding applied locally on a rectangle defined by the user, parametric active contour model initialized by the user, simple linear interpolation between points entered by the user. The user chooses the most appropriate tool to locally delimit the lesion border, and thus progresses step by step. Automatic thresholding or active contour algorithm initialized by the user leads to segmentations that depend less on the user than the only use of interpolation, and the expert was encouraged to apply these tools as often as possible. However, the segmentation remains mostly manual, user-dependent and tedious. An ophthalmologist realized all segmentations used in our experiments. This has been a very hard and long task: as previously mentioned, GA boundaries may be very indented; moreover, there is very often a strong ambiguity on the true position of the GA border and on the appearance or not

of new injured areas. Consequently, the expert had often to go through the images to make or review a decision. On average, it took 10 minutes to process one single image.

4. Methods

Let us denote by $X_n, n \in [1, N]$ the series of acquired images, to be processed. The grayscale pixels are coded by floating point numbers in the interval $[0, 1]$. Our method is made of two main steps: joint image normalization to compensate for illumination variation between images, and morphological processing to enhance the consistent changes between images and attenuate artefacts.

4.1. Joint normalization step

We consider the model where each acquired image $X_n, n \in [1, N]$, is the pixel-wise product of an image of reflectance R_n and an image of illumination I_n :

$$X_n = R_n I_n, n \in [1, N] \quad (1)$$

Our goal is to compensate for illumination variation (i.e. the fact that $I_n \neq I_m$ for $n \neq m$) to get homogeneous illumination between all processed images. In other words, we aim at calculating from X_n a new series of images, $Y_n = R_n J_n$, where J_n is the corrected illumination of the n^{th} image, with $J_n = J_m$ for all pairs. Note that the illumination can remain uneven over the image domain, but this inhomogeneity must be the same in all images of the series to allow differential analysis.

The illumination variation is generally smooth, and we can assume that the pixel-wise ratios I_{n+1}/I_n only contain low frequencies. As there are only few and only small structural changes between two consecutive images, we can also consider that the lowpass filtered version of the ratio X_{n+1}/X_n is a good estimate of I_{n+1}/I_n . So we calculate

$$U_n = \left(\frac{I_{n+1}}{I_n} \right) = \frac{X_{n+1}}{X_n} * f_\sigma, n \in [1, N-1] \quad (2)$$

where f_σ is a Gaussian kernel with standard deviation σ and $*$ the convolution operator. σ is a scale parameter, chosen to remove structural change areas between consecutive images. Then we estimate the ratio of illumination for any pair of images in a recursive way:

$$\begin{aligned} \forall (n, m) \in [1, N]^2, m \geq n, \\ V_{m,n} &= \left(\frac{I_m}{I_n} \right) = U_{m-1} V_{m-1,n} \\ V_{n,n} &= \frac{1}{V_{m,n}} \end{aligned} \quad (3)$$

So, the lowpass filter f_σ is only applied on ratios of consecutive images, so with little change from one to the other, and not on any other pairs where the reflectance is likely to be very different, with large areas of change. Our goal is to apply an algorithm that corrects the images, so that the new illumination ratio becomes close to 1 at each pixel, and that for any pair (m, n) . To achieve that, our preliminary idea was to propose an iterative algorithm that modifies each image against the $(N-1)$ others:

$$\forall (m, n), m > n, \begin{cases} Y_n^{(i+1)} = Y_n^{(i)} V_{m,n}^{(i) \frac{\alpha}{2}} \\ Y_m^{(i+1)} = Y_m^{(i)} V_{m,n}^{(i) - \frac{\alpha}{2}} \end{cases} \quad (4)$$

where i denotes the iteration, $Y_n^{(0)} = X_n$, $V_{m,n}^{(i)}$ is the updated estimate of the illumination ratio.

Let us assume that the illumination at a given pixel (x, y) is higher in $Y_m^{(i)}$ than in $Y_n^{(i)}$. Then its illumination in the n^{th} image will be multiply by a factor $V_{m,n}^{(i) \frac{\alpha}{2}}(x, y) > 1$ while its illumination in the m^{th} image will be divided by the same factor, resulting that the new illumination values in $Y_n^{(i+1)}$ and $Y_m^{(i+1)}$ become closer. Processing all pairs in this way enables us to correct every image against all the others. The parameter α controls the amount of correction brought by each image of the series to process another one. Thus, we can assume that the illumination distortion between images progressively decreases so that, after L iterations, we get a new series $Y_n = Y_n^{(L)}$ with same illumination in all images: $I_m^{(L)} = I_n^{(L)}$ for all pairs (n, m) . Note that the algorithm in equation (4) is equivalent to

$$\forall n, Y_n^{(i+1)} = Y_n^{(i)} \prod_{k=1}^N V_{k,n}^{(i) \frac{\alpha}{2}} \quad (5)$$

Let us now go farther in the interpretation by studying the convergence of the proposed algorithm. For that, we consider the particular case where all images have the same reflectance but were acquired under different illuminations. For any pair, we have $V_{k,n}^{(i)} = I_k^{(i)} / I_n^{(i)} = Y_k^{(i)} / Y_n^{(i)}$ and the correction algorithm becomes

$$\forall n, Y_n^{(i+1)} = Y_n^{(i)} \prod_{k=1}^N \left(\frac{Y_k^{(i)}}{Y_n^{(i)}} \right)^{\frac{\alpha}{2}}$$

We denote by $Z_n^{(i)}$ the logarithm of $Y_n^{(i)}$. We have

$$\forall n, Z_n^{(i+1)} = Z_n^{(i)} + \frac{\alpha}{2} \sum_{k=1}^N (Z_k^{(i)} - Z_n^{(i)}) \quad (6)$$

At iteration i , we define the difference between any two images:

$$\Delta_{n,m}^{(i)} = Z_n^{(i)} - Z_m^{(i)}$$

At the next iteration, we have

$$\Delta_{n,m}^{(i+1)} = Z_n^{(i+1)} - Z_m^{(i+1)} = \left(1 - \frac{\alpha N}{2} \right) \Delta_{n,m}^{(i)}$$

For $0 < \alpha < \frac{2}{N}$ we have $\lim_{i \rightarrow \infty} \Delta_{n,m}^{(i)} = 0$ for any pair of images, meaning that the algorithm converges to a common image. From equation (6) we notice that the mean value of all images is constant and equal to $\frac{1}{N} \sum_{n=1}^N \log X_n$. Consequently the final corrected image, after convergence, is the geometric average of all input images:

$$\forall n, \lim_{i \rightarrow \infty} Y_n^{(i)} = \left(\prod_{k=1}^N X_k \right)^{\frac{1}{N}}$$

Thus, our final corrected images, all equal to the geometric average of the input images, are given by

$$Y_n = X_n F_n \text{ with } F_n = \frac{\left(\prod_{k=1}^N X_k \right)^{\frac{1}{N}}}{X_n} = \prod_{k=1}^N \left(\frac{I_k}{I_n} \right)^{\frac{1}{N}} \quad (7)$$

As we are not in this ideal case where all images have the same reflectance, our strategy will be to estimate all illumination ratios as described by equations (2) and (3), and to calculate in just one pass the corrected images:

$$Y_n = X_n F_n \text{ with } F_n = \prod_{k=1}^N V_{k,n}^{\frac{1}{N}} \quad (8)$$

Finally, our correction algorithm is as follows:

Input : series of images $X_n, n \in [1, N]$

- Calculate U_n for $n \in [1, N - 1]$ (2)
- Calculate $V_{k,n}$ for all pairs (k, n) (3)
- Calculate new images Y_n (8)

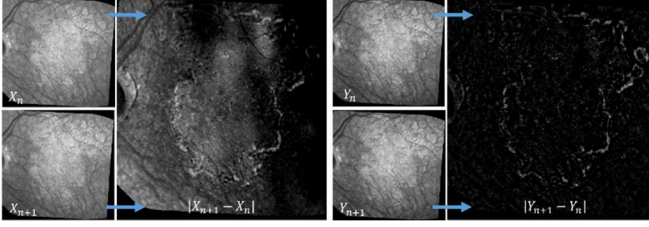


Fig. 2. Joint normalization applied on series 005-L. On the left, two consecutive source images X_n and X_{n+1} and their difference $|X_{n+1} - X_n|$. On the right, the same after normalization.

The low-pass filtering in equations (2) and (3) gives only an approximation of the illumination ratio between two images. Therefore, considering all pairs (m, n) of images, even when the two medical exams were done at a large time interval, with important reflectance changes, is probably not appropriate, even with our recursive approach. For that reason, we propose to implement the algorithm with a sliding temporal window: each image is corrected considering $2K + 1$ images, K images acquired before and K images acquired after the processed image. The overlap ensures that all images of the series will converge to the same illumination. Experiments showed that $2 \leq K \leq 7$ (5 to 15 images in the window) are good values for our dataset (Section 5.2). Fig. 2 shows an example of two consecutive images extracted from a series of 9 images. We observe that the initial absolute difference does not provide meaningful information, while it reveals the GA growth after normalization.

Finally, it is worth noting that the new illumination generally remains spatially uneven. It is certainly possible to compensate for this distortion when the series begins without any GA, thanks to classical background subtraction techniques. This has not been explored yet.

4.2. Enhancement step

The second step of our algorithm exploits prior knowledge about the studied disease: we know that the GA, which is brighter than the surrounding structures, grows continuously, without any chance of recovery. It means that the GA in image Y_n is included in the GA in $Y_{m>n}$. Therefore, we apply morphological reconstructions to all triplets Y_{n-1}, Y_n, Y_{n+1} as follows:

$$Y_n^{(Morph)} = \left(Y_n \ R_D^{Y_n} (Y_{n-1}) \ R_E^{Y_n} (Y_{n+1}) \right)^{\frac{1}{3}} \quad (9)$$

where $R_D^M(S)$ (resp. $R_E^M(S)$) denotes the reconstruction by dilation (resp. erosion) of the marker S in the mask M . The first reconstruction ensures that we keep only bright areas that were already present in the previous image. The second one suppresses bright areas that are not present in the next image. The pixel-wise product to the power $1/3$ merges all these conditions and enhances the meaningful bright areas in image n . The normalizing algorithm is applied again on the images $Y_n^{(Morph)}$, to ensure that the new series is properly normalized and to further reduce the influence of artefacts. We denote by Z_n the resulting images.

Fig. 3 shows an example of the whole processing on 3 consecutive images. The difference image is calculated by $\max\{|X_{n+1} - X_n|, |X_n - X_{n-1}|\}$. The first step (normalization) enables us to enhance the differences between the images, while the second step (morphological processing) attenuates the bright artefacts that are not new lesions.

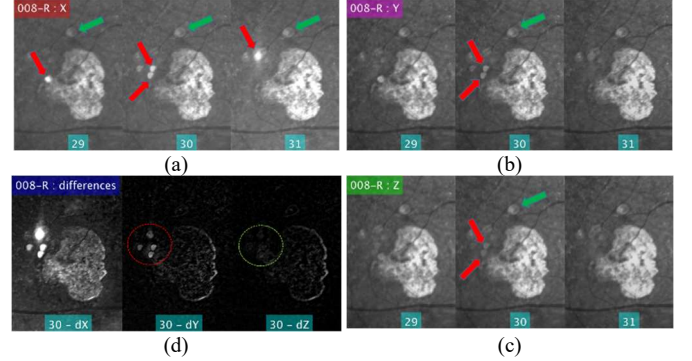


Fig. 3. Normalization and image enhancement. (a) Three consecutive images $X_{n-1}, X_n, X_{n+1}, n = 30$; the corresponding normalized images before (b) and after (c) morphological processing. The green arrows point to a part of the actual GA and the red arrows to bright artefacts; (d) corresponding difference images, emphasized by a factor 3, before normalization (left), after normalization (middle), after normalization and morphological processing (right). The normalization reveals the GA growth but light artefacts are also enhanced (inside the red circle); the morphological processing reduces the intensity of artefacts (green circle).

4.3. Pairwise correction based on polynomial fitting [10]

The proposed normalizing algorithm (Section 4.1) corrects each image with respect to all the others and this is our main contribution with respect to the literature. In [10], two images are processed to correct the second one with respect to the first one, based on surface fitting through a polynomial function. Here we extend the method to process the whole stack iteratively. Let us denote by $\tilde{X}_n, n \in [1, N]$, the resulting images, with $\tilde{X}_1 = X_1$. For every $n \in [2, N]$, a fourth-order polynomial function P_n (15 parameters) is estimated to minimize $\|\tilde{X}_{n-1} - P_n X_n\|^2$ where $\|\cdot\|$ is the L2 norm. So, we have

$$\begin{aligned} \tilde{X}_n &= P_n X_n \\ \text{with } P_n &= \arg \min_{Q_n} \|\tilde{X}_{n-1} - Q_n X_n\|^2 \end{aligned} \quad (10)$$

As our goal is to detect subtle changes between consecutive images, we plan to apply this processing on the series $Y_n^{(Morph)}$ (9), instead of applying the joint normalization a second time as described in Section 4.2. We denote by \tilde{Z}_n the resulting series:

$$\begin{aligned} \tilde{Z}_n &= P_n Y_n^{(Morph)} \\ \text{with } P_n &= \arg \min_{Q_n} \|\tilde{Z}_{n-1} - Q_n Y_n^{(Morph)}\|^2 \end{aligned} \quad (11)$$

So, in the end we propose two possible workflows that will be evaluated in Section 5:

$$\begin{aligned} X_n &\xrightarrow{\text{Joint norm. (8)}} Y_n \xrightarrow{\text{Morphological enhancement (9)}} Y_n^{(Morph)} \xrightarrow{\text{Joint norm. (8)}} Z_n \\ X_n &\xrightarrow{\text{Joint norm. (8)}} Y_n \xrightarrow{\text{Morphological enhancement (9)}} Y_n^{(Morph)} \xrightarrow{\text{Pairwise refinement (11)}} \tilde{Z}_n \end{aligned}$$

These workflows will be compared to the algorithm of [10] extended according to (10):

$$X_n \xrightarrow{\text{Pairwise refinement (10)}} \tilde{X}_n$$

5. Results

We rely on the manual segmentations performed by our expert (18 series, 328 images in total) to evaluate the proposed methods, in two different ways: though the estimation of the contrast between the GA and the background (i.e. surrounding tissues and anatomical structures, Sections 5.1 and 5.2) and through the evaluation of preliminary segmentation results (Sections 5.3 and 5.4). It is worth noting that all experiments apply to the whole series (not on single images or in pairs) since our main contribution is the joint processing of the entire stack of images. Furthermore, the proposed segmentation methods are basic and do not constitute a methodological contribution on its own; but as they are fully automatic, unsupervised and applied pixel-wise over the entire series, they help us to assess the relevance of the joint intensity normalization.

5.1. Ground truth

Let us denote by GT_n the binary image of the GA segmentation made by the medical expert in the n^{th} image of a series. The inside of the GA delineated by the expert is set to 1 and the outside to 0. To deal with possible errors and inconsistencies (such as $GT_n \not\subseteq GT_{m>n}$), we estimate the inside $L_n^{(in)}$ of the GA in image n and the outside $L_n^{(out)}$ by:

$$\begin{aligned} L_n^{(in)} &= \bigcup_{k=1}^n GT_k \\ L_n^{(out)} &= \overline{L_n^{(in)}} \end{aligned} \quad (12)$$

where \bar{E} denotes the complement of the set E . This choice ensures that $L_n^{(in)} \subseteq L_{n+1}^{(in)}$.

To deal with series where a GA is already present in the first image (13 series over 18), we define the total growth of the GA from the first image to the n^{th} by:

$$GTh_n = L_n^{(in)} \cap \bigcap_{k=1}^N \overline{GT_k} \quad (13)$$

Finally, the growth between two images is given by

$$gth_{n:m>n} = GTh_m \cap \overline{GTh_n} \quad (14)$$

These three equations gives us the ground truth for the GA segmentation (12), the growth from the first image to the n^{th} (13) and the growth between two images (14). We rely on these definitions to evaluate our methods.

We also estimate the contrast between two regions, denoted by R_1 and R_2 , in a given image X , by calculating the mean signal power $P(X, R)$ in these two regions, and then the difference $\Delta P(X, R_1, R_2)$ in dB:

$$P(X, R) = 10 \log_{10} \left[\frac{1}{|R|} \sum_{(x,y) \in R} (X(x,y))^2 \right] \quad (15)$$

where $|R|$ denotes the number of pixels in the region R .

$$\Delta P(X, R_1, R_2) = P(X, R_1) - P(X, R_2) \quad (16)$$

The higher $|\Delta P(X, R_1, R_2)|$, the higher the mean contrast in X between the two regions R_1 and R_2 and the easier the two regions

can be segmented. This measure will help us to evaluate the benefit of our algorithms by comparing the contrast between healthy and unhealthy areas in images derived from X_n (source images), Y_n (after normalization), and Z_n or \tilde{Z}_n (after morphological processing).

5.2. Evaluation from variance images calculated along the temporal axis

We first evaluate the joint normalization processing (Section 4.1) and the morphological enhancement (Section 4.2). These algorithms aim at correcting illumination variations between all images. As a result, the illumination may be still uneven in every image, but it should be the same in all images of the series. So, the gray level of a pixel belonging to a healthy region (outside the GA in all images) must be stable in the series (low variance), contrary to a pixel belonging to a region of GA growth whose gray level increases significantly at a given time. Consequently, we can evaluate the benefit of our normalization algorithm by estimating the contrast between these two regions in a variance image.

We consider a sliding temporal window of $M = 2m + 1$ consecutive images (Fig. 4). For each image $n \in [m + 1, N - m]$ we take the images X_{n-m} to X_{n+m} to calculate the variance of the gray levels at each pixel. We denote by $X_n^{(VAR)}$ the resulting image. We consider the healthy region $R_n^{(healt)} = L_{n+m}^{(out)}$ and the region of growth $R_n^{(growth)} = gth_{n-m:n+m}$, and we evaluate the contrast between both: $\Delta P(X_n^{(VAR)}, R_n^{(growth)}, R_n^{(healthy)})$ (15)(16). We average the results in the series and over all series to get mean values, denoted by $\Delta P(X^{VAR})$.

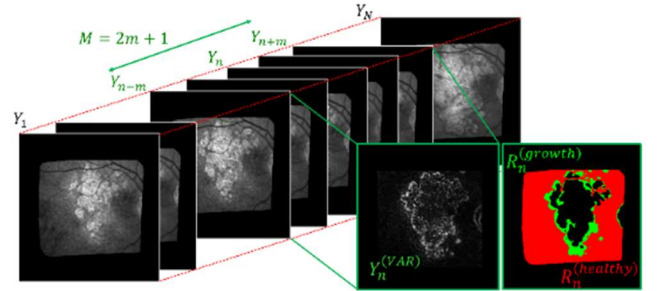


Fig. 4. Variance images calculated over $M = 2m + 1$ consecutive images. Mean power values are interpreted by considering the GA growth over this temporal window ($R_n^{(growth)}$) and the still healthy area ($R_n^{(healt)}$).

Fig. 5 displays the results for different sizes of sliding window in the normalization algorithm ($2K + 1$, Section 4.1), different sizes of sliding window to compute the variance images ($M = 2m + 1 \in \{5, 9, 11, 13, 15\}$), and at different steps of the algorithm: on the source images (X_n), after normalization (Y_n) and after morphological processing (Z_n). We consider only the 6 series having more than 15 images (maximum value for M) with GA, for the sake of comparison. These results demonstrate the benefits of our algorithm with great improvements of the contrast whatever the parameter K and the length of the analysis (M): the gain ranges from 11 to 15dB (Fig. 5d) after the whole processing. Most of the gain results from the normalization algorithm but the morphological processing contributes also to enhance the GA growth area (about +1.5dB). Setting $2K + 1 = 11$ leads to the highest contrasts (b,c) whatever M . But this experiment also shows that K is not a sensitive parameter since the performances remain similar for $(2K + 1) \in [5, 21]$.

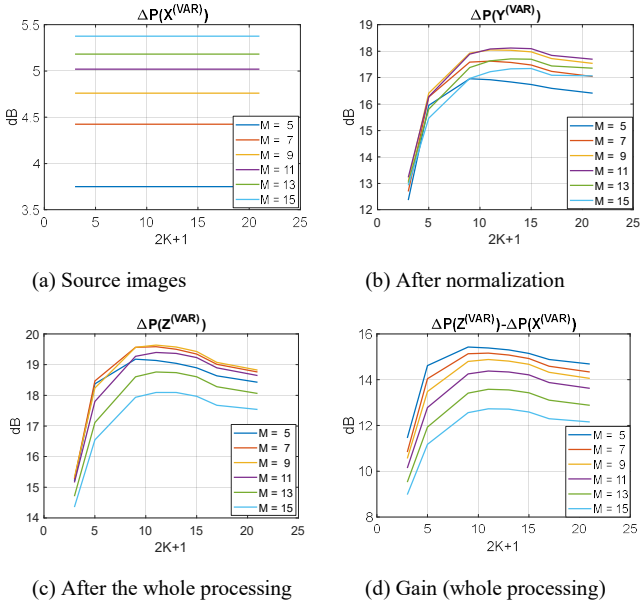


Fig. 5. Contrast enhancement between the GA growth area and the healthy area for several parameters K (temporal window in the normalization algorithm) and M (temporal window to compute the variance).

5.3. Evaluation from an enhanced image of the GA total growth

Let us define the image $G_n^{(X)}$ enhancing the area of growth between the first and the n^{th} image:

$$G_n^{(X)} = \max_{k \in [2, n], l < k} (X_k - X_l), n \in [2, N] \quad (17)$$

Note that maximizing over all differences, $(X_k - X_l)$ with $l < k$, makes sense only because all images have been jointly normalized. We compare the signal power in the total growth area GTh_N (13) to the power outside the GA, $L_N^{(out)}$ (12), by calculating $\Delta P(G_N^{(X)}, GTh_N, L_N^{(out)})$ (15)(16) at the different steps of the algorithm. Table 2 summarizes the results averaged on the 18 series of the experimentation sub-database. The main gain obviously results from the normalization step, but the morphological treatment improves generally slightly the contrast, (+0.6dB). The results are again stable over K , with slightly better contrasts for $2K + 1 = 7$.

$\Delta P(G_N^{(X)})$	K	$\Delta P(G_N^{(Y)})$	Gain	$\Delta P(G_N^{(Z)})$	Gain
2.76 ± 3.47	2	8.44 ± 2.35	5.68 ± 2.37	9.52 ± 2.04	6.77 ± 2.91
	3	9.08 ± 2.01	6.32 ± 3.24	9.90 ± 1.84	7.14 ± 3.32
	4	9.25 ± 2.07	6.49 ± 3.57	9.70 ± 1.71	6.94 ± 3.52
	5	9.06 ± 2.06	6.30 ± 3.79	9.52 ± 1.81	6.75 ± 3.50

Table 2. Contrast between the area of GA growth GTh_N and the outside $L_N^{(out)}$ in the enhanced image (17) calculated from X_n , Y_n or Z_n . The gains are calculated with respect to the initial contrast $\Delta P(G_N^{(X)})$. Values are in dB.

In Table 3, we evaluate the pairwise processing described in Section 4.3, where a fourth-order polynomial (15 parameters) is estimated to model and correct the illumination distortion between two consecutive images (10,11). This treatment extends the pairwise normalization proposed in [10] when iterated on the source images X_n (10). It is used as refinement of our joint processing when applied to $Y_n^{(morph)}$ (11).

The highest gains are achieved with our joint processing methods, with a slightly higher contrast when the last step is

the pairwise polynomial fitting (\tilde{Z}_n) instead of a second joint normalization (Z_n). The contrast between the total GA and the outside is improved by 7.4dB. In contrast, the polynomial fitting alone [10] leads to less good results (-2.2 dB), showing the interest of first processing the whole series jointly. It is even more evident when monitoring the images, as illustrated in Fig. 6: despite the noise, a good enhancement of the true structural changes between consecutive images is obtained with our method (\tilde{Z}_n); in contrast, the polynomial fitting applied directly on the source images (\tilde{X}_n , [10]) does not completely compensate for illumination variation, and white areas are still present in the background of the difference images.

Method (IP)	K	$\Delta P(G_N^{(IP)})$	Gain
Source images (X)	-	2.76 ± 3.47	
Polynomial fitting [10] (\tilde{X})	-	7.93 ± 2.81	5.17 ± 2.74
Our method 1 (Z)	3	9.90 ± 1.84	7.14 ± 3.32
Our method 2 (\tilde{Z})	2	10.12 ± 2.37	7.37 ± 2.99

Table 3. Contrast of the GA growth area GTh_N : comparison with the illumination correction method of [10]. Values are in dB.

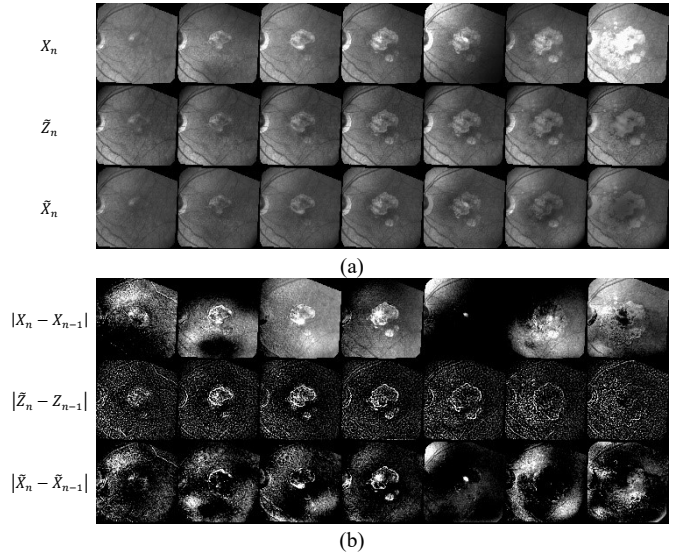


Fig. 6. Comparison between our joint processing method and the pairwise correction of [10]: (a) source images (first line), series processed with our joint normalization method (\tilde{Z}_n , second line), and series processed pairwise [10] (\tilde{X}_n , third line); (b) the difference images after contrast stretching.

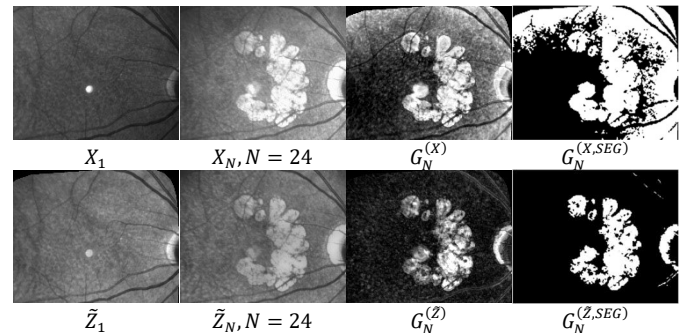


Fig. 7. GA growth enhancement and segmentation, applied on a series of 24 images with no lesion in the first image (001-R). The first row shows the images before normalization (X_n) and the second row after processing (\tilde{Z}_n). The enhanced GA growth is shown in the third column and the corresponding segmentation in the fourth column (F1 score: 0.82 for $G_N^{(Z,SEG)}$).

Fig. 7 shows the enhanced GA growth (17) before and after processing. In this example, the set of images starts from the very

early stages of the disease to advanced ones and the whole GA is enhanced with a very dark background ($G_N^{(Z)}$) compared to the result without processing ($G_N^{(X)}$). The growth of diseased tissues around the optical disk is also clearly visible.

A simple thresholding with Otsu's method enables us to segment the total GA growth, thanks to the proposed processing, while it was impossible with the source images (Fig. 7). We denote by $G_N^{(X,SEG)}$ the binarized $G_N^{(X)}$ image. The segmentation is evaluated in Table 4 via classical metrics: precision, recall and F1-score, which is the harmonic mean of precision and recall, given the ground truth Gth_N (13).

	T_{Otsu}	Lesion (5 series)			Growth (13 series)		
		Prec.	Rec.	F1-sc.	Prec.	Rec.	F1-sc.
$G_N^{(X,SEG)}$	0.44 ± 0.04	0,29 $\pm 0,10$	0,86 $\pm 0,15$	0,42 $\pm 0,13$	0,27 $\pm 0,23$	0,63 $\pm 0,16$	0,33 $\pm 0,17$
$G_N^{(Y,SEG)}$	0.37 ± 0.04	0,79 $\pm 0,11$	0,79 $\pm 0,08$	0,78 $\pm 0,07$	0,45 $\pm 0,27$	0,70 $\pm 0,12$	0,49 $\pm 0,21$
$G_N^{(Z,SEG)}$	0,37 $\pm 0,03$	0,84 $\pm 0,06$	0,78 $\pm 0,10$	0,81 $\pm 0,06$	0,51 $\pm 0,25$	0,68 $\pm 0,13$	0,54 $\pm 0,19$
$G_N^{(Z,SEG)}$	0,36 $\pm 0,02$	0,80 $\pm 0,18$	0,72 $\pm 0,11$	0,74 $\pm 0,08$	0,56 $\pm 0,26$	0,66 $\pm 0,13$	0,56 $\pm 0,19$
$G_N^{(\tilde{X},SEG)}$	0,37 $\pm 0,03$	0,66 $\pm 0,19$	0,76 $\pm 0,15$	0,70 $\pm 0,16$	0,49 $\pm 0,25$	0,68 $\pm 0,14$	0,52 $\pm 0,20$

Table 4. Evaluation of the GA growth segmentation by Otsu's method at the different steps of the algorithm ($G_N^{(X,SEG)}$, $G_N^{(Y,SEG)}$ with $K = 3$, $G_N^{(Z,SEG)}$ or $G_N^{(\tilde{X},SEG)}$ with $K = 2$) and comparison with the results obtained with the pairwise polynomial correction [10] ($G_N^{(\tilde{X},SEG)}$).

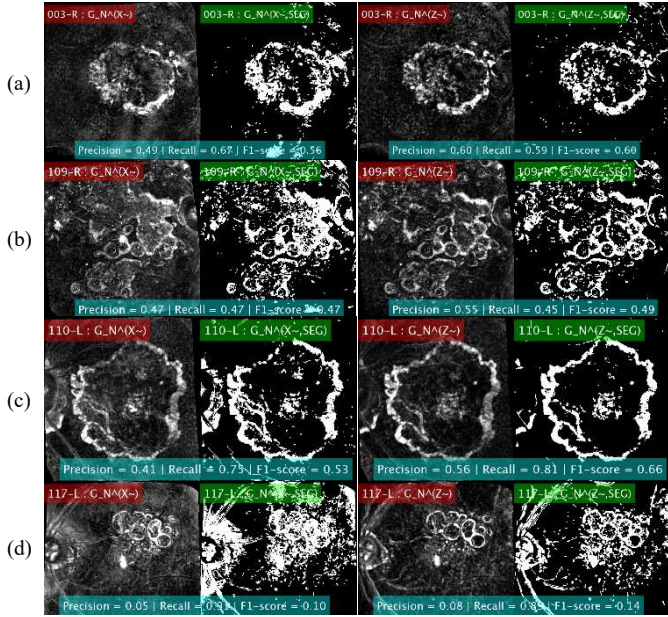
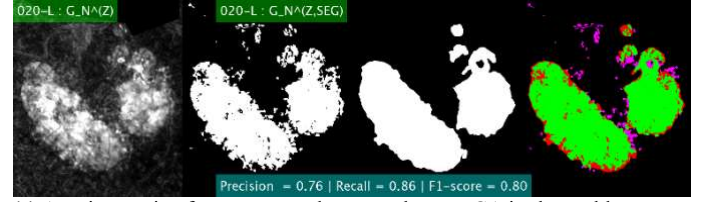


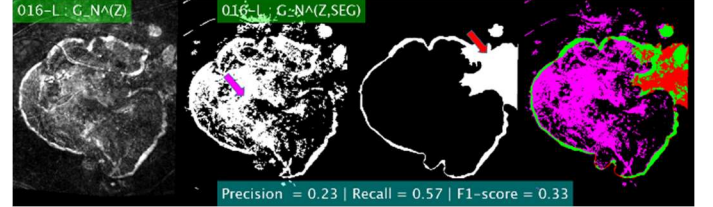
Fig. 8. Segmentation of the GA growth for series starting at an advanced stage: results obtained from \tilde{X}_n (pairwise method of [10], left) and from \tilde{Z}_n (our joint processing, right). The enhanced images $G_N^{(Z)}$ are better contrast than $G_N^{(X)}$, leading to a more accurate segmentation of the total GA growth.

The best scores are achieved with the joint normalization (Y_n , Z_n or \tilde{Z}_n), and the highest F1-scores are obtained with the whole processing including morphological treatment (Z_n or \tilde{Z}_n). The polynomial fitting alone (\tilde{X}_n) [10] performs less, with smaller mean scores and higher standard deviations; moreover, the precision scores are significantly worse, revealing over segmentation issues (Fig. 8, left column). We notice that Otsu's threshold (T_{Otsu}) is remarkably stable over the series (0.36 ± 0.02 for \tilde{Z}_n). The results for the detection of the whole growth are especially good with a F1-score equal to 0.81 ± 0.06 (Z_n) or

0.74 ± 0.08 (\tilde{Z}_n) (Fig. 7, Fig. 9a). For the other series, the F1-score is much lower, equal to 0.56 ± 0.19 (\tilde{Z}_n , Fig. 8, right column). The main problem is that the texture changes inside the GA over time, resulting in a lot of false positives and lower precision (Fig. 9b).



(a) A series starting from a very early stage where no GA is observable.



(b) A series starting from an advanced stage. Texture changes inside the lesion are classified as lesion growth areas (magenta), resulting in low precision. The recall is low since the expert segmented the right part as lesion in the last image (red), contrary to the algorithm.

Fig. 9. Segmentation of the GA growth. Left to right: enhanced image $G_N^{(Z)}$, proposed segmentation $G_N^{(Z,SEG)}$, ground truth Gth_N , and comparison (true positives in green, false negatives in red and false positives in magenta).

5.4. GA progress between consecutive images

Finally, still with the aim of evaluating our image processing algorithms, we tried to segment the GA growth between consecutive images. This process is a prerequisite to elaborate predictive models of the disease. For that, we consider the pixels belonging to the GA growth area, so satisfying $G_N^{(Z,SEG)}(x, y) = 1$ (Section 5.3), and we study their gray levels in the whole series. Let us denote by $V^{(Z,(x,y))}(n)$, $n \in [1, N]$, the gray levels of such a pixel (x, y) in the processed series Z_n . We assume that this pixel is not part of the GA in the first image and that it will be included in the GA at a given time t . This means that the gray levels in $V^{(Z,(x,y))}$ may be low for $1 \leq n < t$ and much higher for $t \leq n \leq N$ and roughly stable in the two states. For each index $k \in [2, N]$, we classify the elements of $V^{(Z,(x,y))}$ in 2 classes: $V^{(Z,(x,y))}(n) \in C_0(k)$ if $n < k$, $V^{(Z,(x,y))}(n) \in C_1(k)$ otherwise. We calculate the inter-class variance according to Otsu's criterion for each index k :

$$\sigma_{Otsu}^2(Z,(x,y))(k) = \frac{k-1}{N} \left(\overline{V_{1:k-1}^{(Z,(x,y))}} - \overline{V_{1:N}^{(Z,(x,y))}} \right)^2 + \frac{N-k+1}{N} \left(\overline{V_{k:N}^{(Z,(x,y))}} - \overline{V_{1:N}^{(Z,(x,y))}} \right)^2 \quad (18)$$

where $\overline{V_{k:l}^{(Z,(x,y))}}$ denotes the mean value of the elements k to l of vector $V^{(Z,(x,y))}$. We calculate a segmentation map by maximizing the interclass variance $\sigma_{Otsu}^2(Z,(x,y))$:

$$MAP^{(Z)}(x, y) = \arg \max_k \sigma_{Otsu}^2(Z,(x,y))(k) \quad (19)$$

This method is not relevant for pixels that do not switch from a healthy to an unhealthy state. Consequently, the method is only applied in the GA growth area, so for pixels satisfying $G_N^{(Z,SEG)}(x, y) = 1$ (Section 5.3). So, we calculate the series

$SEG_n^{(Z)}, n \in [1, N]$, of binary images of the GA from image 1 to image N by

$$SEG_n^{(Z)}(x, y) = \begin{cases} 1 & \text{if } n \geq MAP^{(Z)}(x, y) \text{ and } G_N^{(Z, SEG)}(x, y) = 1 \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

Finally, we calculate the growth between two consecutive images by

$$SEG_{n-1:n}^{(Z)} = SEG_n^{(Z)} \& \overline{SEG_{n-1}^{(Z)}} \quad (21)$$

which reveals the progression of the GA between the $(n-1)^{\text{th}}$ and the n^{th} images.

Table 5 shows the quantitative evaluation of this segmentation method, applied on the processed series \tilde{X}_n , Z_n and \tilde{Z}_n . The method is obviously not applicable before normalization. The automatic segmentations ($SEG_{n-1:n}^{(X)}$, $SEG_{n-1:n}^{(Z)}$, $SEG_{n-1:n}^{(Z)}$) are compared to the manual segmentations $gth_{n-1:n}$ (14) using the classical metrics, which are averaged over the images n and then over the series. The best results are obtained on the series \tilde{Z}_n , while the segmentations deduced from the pairwise treatment (\tilde{X}_n) [10] reach lower mean scores with higher standard deviations. This again demonstrates the interest of jointly normalizing all the images in a series before attempting to detect growth.

	Precision	Recall	F1-score
$SEG_{n-1:n}^{(Z)}$	0.29 ± 0.12	0.29 ± 0.14	0.25 ± 0.12
$SEG_{n-1:n}^{(Z)}$	0.30 ± 0.13	0.31 ± 0.14	0.28 ± 0.12
$SEG_{n-1:n}^{(X)}$ [10]	0.27 ± 0.16	0.28 ± 0.14	0.25 ± 0.14

Table 5. Evaluation of the segmentation of the GA progress between consecutive images. The results are given for the series Z_n and \tilde{Z}_n ($K = 2$) processed with our approach and for the series \tilde{X}_n processed with [10].

Fig. 10 shows samples of growth segmentation between consecutive images ($SEG_{n-1:n}^{(Z)}$, (21)) in one series, and Fig. 11 colored representations of the GA growth ($MAP^{(Z)}$, (19)) for two different series, along with the corresponding ground truth.

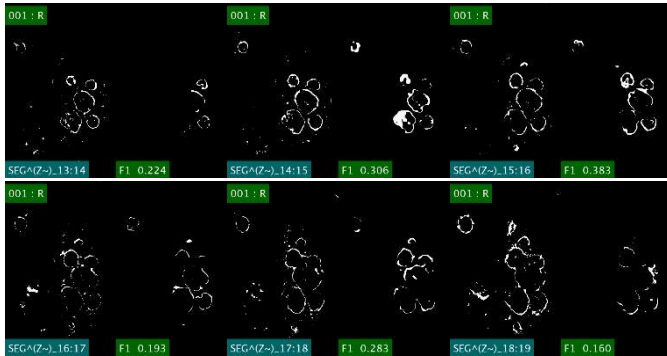


Fig. 10. Examples of very consistent change detections between consecutive images, with F1 scores. Every time, the automatic segmentation $SEG_{n-1:n}^{(Z)}$ (19) is on the left and the ground truth $gth_{n-1:n}$ (12) on the right.

Overall, the segmentation images $SEG_{n-1:n}^{(Z)}$ (Fig. 10) look very consistent, and the colored maps (Fig. 11) look similar to their equivalent deduced from the ground truth. However, the average metrics are quite low (Table 5). There are several explanations for

this. First, it should be remembered that the manual segmentation is very difficult to achieve and is not accurate either. Second, the GA progress between two consecutive images is structurally very thin, which explains that a small difference between the segmented area and the ground truth results immediately in low metrics. Especially, small shifts in the change detection maps lead to an increase of false positives and false negatives, therefore to weaker metrics, although the automatic GA growth detection looks very nice (Fig. 10, Fig. 12b).

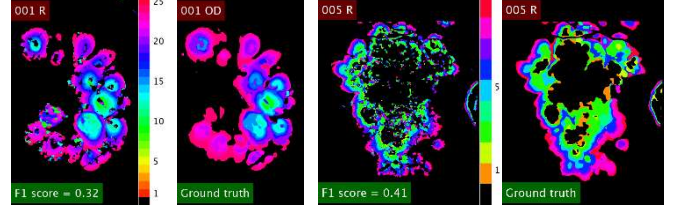


Fig. 11. Colored map of the GA progress in the series (17). The n^{th} color in the colormap corresponds to a GA progress between image $n-1$ and image n . The automatic result is on the left and the ground truth on the right. The F1 score is the average of all F1 scores obtained for each pair of consecutive images (21).

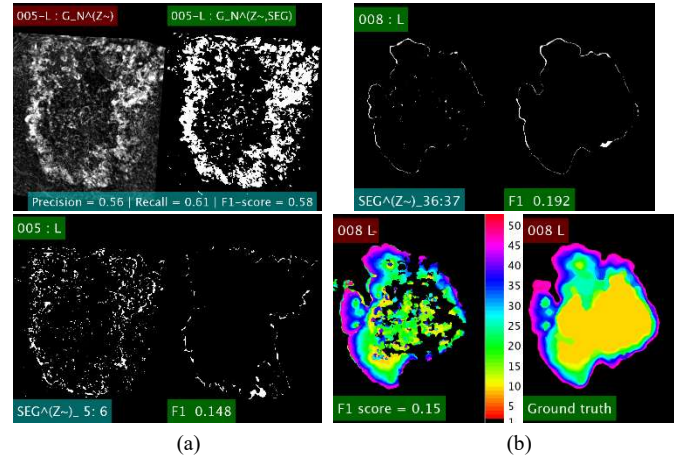


Fig. 12. Limits of the segmentation method. In each illustration, the automatic result is on the left and the ground truth on the right. (a) Case of a series of 9 images corresponding to late stages of ARMD. The area of total GA growth $G_N^{(Z, SEG)}$ is over segmented (first row) leading to many false positives in the detection of changes between consecutive images, inside and outside the GA (second row). (b) Case of a series starting from the beginning of the disease. The black areas inside the colored map $MAP^{(Z)}$ come from the under-segmentation of the total GA growth $G_N^{(Z, SEG)}$ (second row). Apart from that, the segmentation between two consecutive images looks very consistent (first row), and the low metrics result also from their sensitivity to manual segmentation imprecision, as the structures to be detected are very thin.

Finally, the main limits of the proposed segmentation method come from the texture changes inside the lesion during the growth, disturbing the decision criterion (high variance in the class C_1). In addition, an over-segmentation of the total growth ($G_N^{(Z, SEG)}$) allows false positives inside and outside the GA (20), and this has a strong impact on the precision and on the F1 score, especially for series starting with advanced GA for which we have a low precision in the $G_N^{(Z, SEG)}$ segmentations (0.56 ± 0.26 , Table 4). Under-segmentation of the total GA growth has also a negative impact. Fig. 12 illustrates all of these problems, responsible for the rather low metrics presented in Table 5, despite the detections of changes appear to be very consistent. The best results are obtained on average with the normalized images processed by mathematical morphology and refined pairwise (\tilde{Z}_n), for $K = 2$ (5 images in the sliding window), with a F1 score around 0.28 ± 0.12 . The

morphological treatment improves the joint normalization by introducing a better consistency along the temporal axis, and the final pairwise refinement further enhance the GA growth between two consecutive images.

Fig. 13 shows additional examples of colored maps obtained with our joint processing method (center column) or with the pairwise treatment [10] (right column), along with the ground truth (left). Despite the metrics are not much higher with our method (Table 5), we can observe that the colored maps derived from \tilde{Z}_n are more consistent than the ones derived from \tilde{X}_n . We can assume that the \tilde{Z}_n series are likely to give better segmentation results than the \tilde{X}_n series once a more advanced segmentation algorithm involving spatial regularization and/or morphological processing is applied. In some cases, our colored map looks also more consistent than the one derived from the manual segmentation, especially when the GA has a very complex shape (Fig. 13b,d).

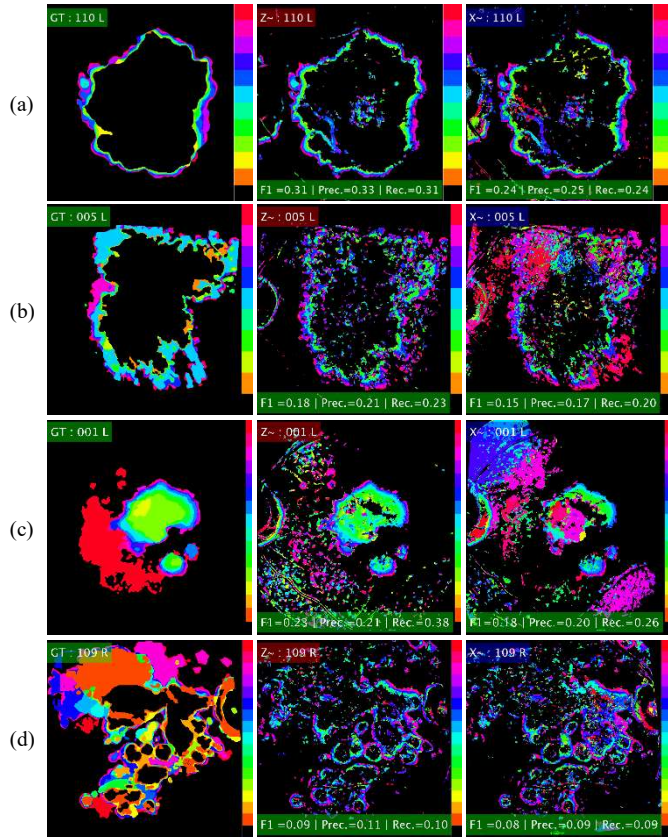


Fig. 13. Colored maps of the GA progress in the series: comparison between the ground truth (left), the maps derived from our joint normalization method (center, \tilde{Z}) and from the pairwise processing of [10] (right, \tilde{X}).

6. Discussion

We have presented a set of experiments evaluating the advantages of our normalization algorithm for the direct comparison of gray levels in a series of fundus images, and, ultimately, the detection of the GA growth by differential analysis. We first calculated the variance at each pixel, which must be low in the areas that remain healthy and, if applicable, within the GA present in the first image, and high in the region of GA growth (Section 5.2). We found that our processing increases the contrast between the growth area and the rest of the image by 11 to 15dB (Fig. 5), which is very important. It means that the visual inspection of the series, by scrolling through the images, is greatly facilitated, since illumination distortions no longer disturb the interpretation. The GA expansion becomes much more visible. The main gain results from the illumination normalization, but the

morphological treatment attenuates local artefacts that could be mistaken for new GA zones. This experiment also demonstrates the advantage of normalizing on a sliding window and helps us to set its optimal size (around 9 consecutive images).

In addition, the proposed preprocessing allows to perform a differential analysis, in order to detect the GA growth between any two examinations.

Thus, in Section 5.3, we calculated an enhanced image of the total GA growth and we demonstrated that the contrast between the GA growth area and the outside (still healthy region) is around 10dB while it is only around 3dB without treatment (Table 2) and less than 8dB with the pairwise processing proposed in [10] (Table 3). We then experimented standard segmentation methods to further evaluate the benefit of our image preprocessing. We considered two processing flows, including both the joint normalization followed by the morphological processing, but lastly refined by a second joint normalization (Z_n) or a pairwise polynomial fitting (\tilde{Z}_n). The gain in contrast enables a segmentation of the total GA growth area with a F1-score equal to 0.81 ± 0.06 (Z_n) for series starting from the earliest stages of the disease (0.42 ± 0.13 without treatment). The F1-score is equal to 0.70 ± 0.16 with the pairwise processing [10], which is significantly lower (Table 4). The segmentation methods is basic, based on Otsu's thresholding, but the results are already good despite this. The segmentation of the GA growth from an already advanced stage is less good, with an average F1-score equal to 0.56 ± 0.19 (\tilde{Z}_n), against 0.33 ± 0.17 without treatment and 0.52 ± 0.20 with the pairwise processing [10]. The smaller metrics are due to texture changes inside the lesion during time, resulting in higher values in the difference images, and therefore lower precision.

Finally, we explored the possibility of calculating colored maps of the GA progress, in order to obtain a visualization of the growth and speed of growth in a single image (Section 5.4). The proposed method is based on the idea that the grey-levels are stable in each state, healthy (low values) or unhealthy (high value). We faced again the problem of texture variations inside the GA that disturb the detection of the time at which a pixel has switched from a healthy state to an unhealthy one. However, we got nice maps in many cases (Fig. 11, Fig. 13), close to the ground truth, and also very good sets of binary images showing the GA growth between two consecutive images (Fig. 10). The best F1-scores evaluating the segmentation of the GA growth between consecutive images are obtained on \tilde{Z}_n (joint normalization with a window of size 5 followed by morphological processing and pairwise refinement). The F1-scores are not high (around 0.28 ± 0.12 on the whole database, Table 5), but these poor results should be put into perspective, as the structures to be detected are very thin and the manual segmentation prone to inaccuracies. That is why we measured weak metrics in many cases despite the binary images of change looked very consistent (Fig. 10), as did the colored maps summarizing the GA progression (Fig. 11, Fig. 13).

All these experiments demonstrate that the proposed image processing enables the detection of changes between any two images of a series of fundus images, regardless the number of images, the size and shape of the GA. This is very new compared to the literature, which has focused much more on the segmentation of the GA in single images (Section 2), and much less on change detection in series [10][13]. Contrary to [10], we normalize all images jointly and not just in pairs, and the visual inspection of the GA progression maps shows clearly a better consistency with the joint processing (Fig. 13), confirmed by the metrics (Table 5). The differential analysis proposed in [13], where the images are preprocessed separately by a homomorphic filter, could also benefit from our global normalization.

7. Conclusion and perspectives

Very large databases of patients available through routine imaging systems offer an opportunity to study a vast amount of data on dry ARMD. Especially, deriving models of GA progression from fundus images is of high medical interest, in order to better understand the intimate mechanisms of lesion growth or to assess the benefit of medical treatments. However, there is a real lack of automatic algorithms to process all these images, and especially to detect and quantify the GA progression between two medical examinations. Many research works deal with the GA segmentation in single images, but there is not yet an automatic and reliable solution. A promising alternative is to carry out a differential analysis between consecutive images, but the first obstacle is the great illumination inhomogeneity inside the images and especially between the images. This article deals with the correction of illumination inhomogeneity between all the images in a series in order to allow the differential analysis between any two images. For that, we have proposed a method which recursively estimates the illumination ratio between any two images in the series and correct the illumination of every image to the geometric mean of all illuminations. Then, we have presented a morphological processing that enhances the consistent changes between consecutive images and attenuate the others, based on the fact that the GA increases continuously. The resulting images are refined either by another joint normalization or a pairwise polynomial fitting. Finally, we have presented many experiments realized on a database of 18 series (328 images) acquired from 15 patients and segmented manually by a medical expert. The experiments show the benefit of the normalization algorithm and of the morphological processing, in terms of contrast between the healthy area and the GA growth area, and in terms of ability to detect cumulative changes in the series or between consecutive images. The experimental results demonstrate that the normalization algorithm makes possible the differential analysis, and that the morphological processing brings some improvements by attenuating light artefacts. Our first attempt to detect the GA growth in consecutive images shows that it is possible to process all the images of the series jointly, since the gray levels are now comparable at each pixel. We got very nice colored maps of GA progression, offering an intuitive understanding for clinicians, in terms of growth and speed of growth. Further research is needed to determine which refinement is the most suitable after the morphological processing, a second joint normalization or the pairwise polynomial fitting.

It is worth noting that the joint normalization process (Section 4.1) can apply to other types of images, for example color or autofluorescence fundus images, which present the same type of illumination defects. One of the strengths of our algorithm is that there is no precondition, except that the structural changes from one image to the next one must not be too important. There is no sensitive parameter either: the only parameter is the length of the sliding window and it does not require fine tuning. Furthermore, the joint normalization process can be useful to monitor clinical signs of other diseases than dry ARMD. Only the morphological processing (Section 4.2) exploits the fact that the GA continuously grows in dry ARMD, which does not apply medically treated wet ARMD, for example.

More robust segmentation algorithms are required to segment the GA growth in dry ARMD, and we continue investigating several approaches. We have proposed an unsupervised deep learning architecture, based on joint autoencoders, to detect structural changes between two consecutive images [14]. The neural network takes as input the normalized images, as described in Section 4, to help it learn a meaningful representation of the

input data and get rid of insignificant residual illumination variations. Ongoing works follow on from what was presented in section 5.4, but with spatial and temporal regularizations. In parallel, we carry on deep learning approaches, but based on LSTM architectures to process entire time series.

8. References

- [1] Deckert, A., Schmitz-Valckenberg, S., Jorzik, J., Bindewald, A., Holz, F. G., & Mansmann, U. (2005). Automated analysis of digital fundus autofluorescence images of geographic atrophy in advanced age-related macular degeneration using confocal scanning laser ophthalmoscopy (cSLO). *BMC ophthalmology*, 5(1), 8.
- [2] Feeny, A. K., Tadarati, M., Freund, D. E., Bressler, N. M., & Burlina, P. (2015). Automated segmentation of geographic atrophy of the retinal epithelium via random forests in AREDS color fundus images. *Computers in biology and medicine*, 65, 124-136.
- [3] Hu, Z., Medioni, G. G., Hernandez, M., Hariri, A., Wu, X., & Sadda, S. R. (2013). Segmentation of the geographic atrophy in spectral-domain optical coherence tomography and fundus autofluorescence images. *Investigative ophthalmology & visual science*, 54(13), 8375-8383.
- [4] Hu, Z., Medioni, G. G., Hernandez, M., & Sadda, S. R. (2015). Automated segmentation of geographic atrophy in fundus autofluorescence images using supervised pixel classification. *Journal of Medical Imaging*, 2(1), 014501.
- [5] I2k Retina Technology: <https://www.dualalign.com/retinal/image-registration-montage-software-overview.php>
- [6] Köse, C., Şevik, U., & Gençalioglu, O. (2008). Automatic segmentation of age-related macular degeneration in retinal fundus images. *Computers in Biology and Medicine*, 38(5), 611-619.
- [7] Köse, C., Şevik, U., Gençalioglu, O., İkibaş, C., & Kayıkcioglu, T. (2010). A statistical segmentation method for measuring age-related macular degeneration in retinal fundus images. *Journal of medical systems*, 34(1), 1-13.
- [8] Lee, N., Laine, A. F., & Smith, R. T. (2007, August). A hybrid segmentation approach for geographic atrophy in fundus autofluorescence images for diagnosis of age-related macular degeneration. In *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (pp. 4965-4968). IEEE.
- [9] Lee, N., Smith, R. T., & Laine, A. F. (2008, October). Interactive segmentation for geographic atrophy in retinal fundus images. In *2008 42nd Asilomar Conference on Signals, Systems and Computers* (pp. 655-658). IEEE.
- [10] Marrugo, A. G., Millan, M. S., Sorel, M., & Sroubek, F. (2011). Retinal image restoration by means of blind deconvolution. *Journal of Biomedical Optics*, 16(11), 116016.
- [11] Priya, R., & Aruna, P. (2011, April). Automated diagnosis of Age-related macular degeneration from color retinal fundus images. In *2011 3rd International Conference on Electronics Computer Technology* (Vol. 2, pp. 227-230). IEEE.
- [12] Ramsey, D. J., Sunness, J. S., Malviya, P., Applegate, C., Hager, G. D., & Handa, J. T. (2014). Automated image alignment and segmentation to follow progression of geographic atrophy in age-related macular degeneration. *Retina*, 34(7), 1296-1307.
- [13] Troglio, G., Alberti, M., Benediksson, J. A., Moser, G., Serpico, S. B., & Stefánsson, E. (2010, April). Unsupervised change-detection in retinal images by a multiple-classifier approach. In *International Workshop on Multiple Classifier Systems* (pp. 94-103). Springer, Berlin, Heidelberg.
- [14] Dupont, G., Kalinicheva, E., Sublime, J., Rossant, F., & Paques, M. (2020). Analyzing Age-Related Macular Degeneration Progression in Patients with Geographic Atrophy Using Joint Autoencoders for Unsupervised Change Detection. *Journal of Imaging*, 6(7), 57.
- [15] Sadda, S. R., Guymer, R., Holz, F. G., Schmitz-Valckenberg, S., Curcio, C. A., Bird, A. C., ... & Staurengi, G. (2018). Consensus definition for atrophy associated with age-related macular degeneration on OCT: classification of atrophy report 3. *Ophthalmology*, 125(4), 537-548.

Conflict of interest statement

The authors have no affiliation with any organization with a direct or indirect financial interest in the subject matter discussed in the manuscript.