ÓPTICA PURA Y APLICADA
www.sedoptica.es

Type: Research paper
Section: Imaging Techniques

Detection and removal of dust artifacts in retinal images via sparse-based inpainting

Erik Barrios[1,2*], Enrique Sierra[2], Lenny A. Romero[3], María S. Millán[4,S], Andres G. Marrugo[2],

1. Esc. de Ciencias Básicas, Tecnología e Ingeniería, Universidad Nacional Abierta y a Distancia, Corozal, Colombia.

2. Facultad de Ingeniería, Universidad Tecnológica de Bolívar, Cartagena, Colombia.

3. Facultad de Ciencias Básicas, Universidad Tecnológica de Bolívar, Cartagena, Colombia.

4. Dept. de Óptica y Optometría, Universitat Politècnica de Catalunya - BARCELONATECH, Terrassa, Spain.

(*) E-mail: embm32@gmail.com S: miembro de SEDOPTICA / SEDOPTICA member

Received: 13/4/2021 Accepted: 6/9/2021
DOI: 10.7149/OPA.54.3.51060

ABSTRACT:

Dust particle artifacts are present in all imaging modalities but have more adverse consequences in medical images like retinal images. They could be mistaken as small lesions, such as microaneurysms. We propose a method for detecting and accurately segmenting dust artifacts in retinal images based on multi-scale template-matching on several input images and an iterative segmentation via an inpainting approach. The inpainting is done through dictionary learning and sparse-based representation. The artifact segmentation is refined by comparing the original image to the initial restoration. On average, 90% of the dust artifacts were detected in the test images, with state-of-the-art restoration results. All detected artifacts were accurately segmented and removed. Even the most challenging artifacts located on top of blood vessels were removed. Thus, ensuring the continuity of the retinal structures. The proposed method successfully detects and removes dust artifacts in retinal images, which could be used to avoid false-positive lesion detections or as an image quality criterion. An implementation of the proposed algorithm can be accessed and executed through a Code Ocean compute capsule.

Key words: Artifact detection, dust particle, retinal image, fundus image, image restoration, dictionary learning, inpainting, sensor artifact, sparse representation.

REFERENCES AND LINKS / REFERENCIAS Y ENLACES

[1] Abràmoff, M. D., Garvin, M. K., & Sonka, M. (2010). Retinal imaging and image analysis. IEEE reviews in biomedical engineering, 3, 169-208.

[2] B. J. Fenner, R. L. M. Wong, W.-C. Lam, et al., "Advances in Retinal Imaging and Applications in Diabetic Retinopathy Screening: A Review," Ophthalmology and Therapy 7, 333–346 (2018).

[3] A. G. Marrugo and M. S. Millan, "Retinal image analysis: Image processing and feature extraction oriented to the clinical task," Opt. Pura Apl 50(1), 49–62 (2017).

[4] E. Sierra, A. G. Marrugo, and M. S. Millán, "Dust particle artifact detection and removal in retinal images," Opt. Pura Apl 50(4), 379–387 (2017).

[5] A. G. Marrugo, M. S. Millan, M. Sorel, et al., "Restoration of retinal images with spacevariant blur," Journal of Biomedical Optics 19(1), 016023 (2014).

[6] H. Narasimha-Iyer, A. Can, B. Roysam, et al., "Robust detection and classification of longitudinal changes in color retinal fundus images for monitoring diabetic retinopathy," IEEE Trans on Biomed Eng 53(6), 1084–1098 (2006).

[7] P. F. Ordóñez, C. M. Cepeda, J. Garrido, et al., "Classification of images based on small local features: a case applied to microaneurysms in fundus retina images," Journal of Medical Imaging 4(4), 041309 (2017).

[8] A. Manjaramkar and M. Kokare, "Statistical geometrical features for microaneurysm detection," Journal of digital imaging 31(2), 224–234 (2018).

[9] M. Zamfir, E. Steinberg, Y. Prilutsky, et al., "Image defect map creation using batches of digital images," (2010). Patent US 2010/0321537 A1.

[10] R. G. Willson, M. W. Maimone, A. E. Johnson, et al., "An optical model for image artifacts produced by dust particles on lenses," 8th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS) (2005).

[11] C. Li, K. Zhou, and S. Lin, "Removal of dust artifacts in focal stack image sequences. ," in Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012), 2602–2605, IEEE (2012).

[12] H. Altamar-Mercado, A. Patino-Vanegas, and A. G. Marrugo, "Extended Focused Image in White Light Scanning Interference Microscopy," in Imaging and Applied Optics 2019, ITh1C.3, Optical Society of America, (Munich) (2019).

[13] C. Zhou and S. Lin, "Removal of image artifacts due to sensor dust," in Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, 1–8, IEEE (2007).

[14] A. D. Mora, J. Soares, and J. M. Fonseca, "A template matching technique for artifacts detection in retinal images," in 2013 8th International Symposium on Image and Signal Processing and Analysis (ISPA), 717–722, IEEE (2013). [doi:10.1109/ispa.2013.6703831].

[15] M. Niemeijer, M. D. Abramoff, and B. van Ginneken, "Image structure clustering for image quality verification of color retina images in diabetic retinopathy screening," Medical image analysis 10(6), 888–898 (2006). [doi:10.1016/j.media.2006.09.006].

[16] M. S. Millan, A. G. Marrugo, and F. Alba-Bueno, "Quality Changes in Fundus Images of Pseudophakic Eyes," Opt. Pura Apl. 51(4), 50015:1–8 (2018).

[17] T. Köhler, A. Budai, M. F. Kraus, et al., "Automatic no-reference quality assessment for retinal fundus images using vessel segmentation," in Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems, 95–100, IEEE (2013).

[18] D. Veiga, C. Pereira, M. Ferreira, et al., "Quality evaluation of digital fundus images through combined measures.," Journal of Medical Imaging 1(1), 014001 (2014).

[19] S. A. A. Shah, A. Laude, I. Faye, et al., "Automated microaneurysm detection in diabetic retinopathy using curvelet transform," Journal of Biomedical Optics 21(10), 101404 (2016).

[20] P. Yang, L. Chen, J. Tian, et al., "Dust particle detection in surveillance video using salient visual descriptors," Computers & Electrical Engineering 62, 224–231 (2017).

[21] L. Chen, D. Zhu, J. Tian, et al., "Dust particle detection in traffic surveillance video using motion singularity analysis," Digital Signal Processing 58, 127–133 (2016).

[22] L. Hu, L. Chen, and J. Cheng, "Gray spot detection in surveillance video using convolutional neural network," in 2018 13th IEEE Conference on Industrial Electronics and Applications (ICIEA), 2806–2810, IEEE (2018). [doi:10.1109/ICIEA.2018.8398187].

[23] D. A. Forsyth and J. Ponce, Computer vision: A modern approach, Prentice Hall (2011).

[24] K. Zeng, G. Erus, A. Sotiras, et al., "Abnormality Detection via Iterative Deformable Registration and Basis-Pursuit Decomposition," Medical Imaging, IEEE Transactions on 35(8), 1937–1951 (2016).

[25] F. Girard, C. Kavalec, and F. Cheriet, "Statistical atlas-based descriptor for an early detection of optic disc abnormalities," Journal of Medical Imaging 5(01), 1–15 (2019).

[26] C. Kou, W. Li, W. Liang, et al., "Microaneurysms segmentation with a U-Net based on recurrent residual convolutional neural network," Journal of Medical Imaging 6(02), 1–12 (2019).

[27] Enrique Sierra, Andres G Marrugo, Erik Barrios (2021) Dust particle artifact detection in retinal images [Source Code].

[28] R. C. Gonzalez, R. E. Woods, and S. L. Eddins, "Digital image processing using matlab," Gatesmark Publishing (2009).

[29] J. Lewis, "Fast normalized cross-correlation," in Vision interface, 10(1), 120–123 (1995).

[30] J. Lin, L. Yu, Q. Weng, et al., "Retinal image quality assessment for diabetic retinopathy screening: A survey," Multimedia Tools and Applications, 1–27 (2019).

[31] A. Awati, H. C. Rao, and M. R. Patil, "Image Inpainting for Hemorrhage Detection in Mass Screening of Diabetic Retinopathy," in Computing, Communication and Signal Processing, 1011–1019, Springer Singapore, Singapore (2018).

[32] M. Elad, "From exact to approximate solutions," in Sparse and Redundant Representations, 79–109, Springer (2010).

[33] C. Guillemot and O. Le Meur, "Image inpainting: Overview and recent advances," IEEE signal processing magazine 31(1), 127–144 (2014).

[34] E. M. Barrios, A. G. Marrugo, and M. S. Millán, "Lremoving dust artifacts in retinal images via dictionary learning and sparse-based inpainting," in 2019 XXII Symposium on Image, Signal Processing and Artificial Vision (STSIVA), 1–5, IEEE (2019).

[35] M. Aharon, M. Elad, A. Bruckstein, et al., "K-svd: An algorithm for designing overcomplete dictionaries for sparse representation," IEEE Transactions on signal processing 54(11), 4311 (2006).

[36] K. Engan, S. O. Aase, and J. H. Husoy, "Method of optimal directions for frame design," in Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on, 5, 2443–2446, IEEE (1999).

[37] S. Manat and Z. Zhang, "Matching pursuit in a time-frequency dictionary," IEEE Trans Signal Processing 12, 3397–3451 (1993).

[38] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," Systems, Man and Cybernetics, IEEE Transactions on 9(1), 62–66 (1979).

[39] E. Sierra, E. Barrios, A. G. Marrugo, et al., "Robust detection and removal of dust artifacts in retinal images via dictionary learning and sparse-based inpainting," in Pattern Recognition and Tracking XXX, M. S. Alam, Ed., 109950L, SPIE (2019).

1. Introduction

Retinal images are photographs of the eye fundus in which different structures of the retina can be distinguished, e.g., optic disc, fovea, and blood vessels. Since a wide variety of diseases manifest themselves in the retina by signs or anomalies [1], retinal images are a powerful tool for medical diagnosis and screening [2,3]. These images are acquired with eye fundus cameras, which are composed of a special camera attached to a low magnification microscope. Fundus cameras, like any other imaging device, may suffer from dust particles attached to the sensor and lenses. These particles impede light from reaching small regions of the sensor. As shown in Fig. 1, this phenomenon generates small dark spots in the obtained images that are not part of the observed scene's actual signal. Also, dead or stuck pixels may produce artifacts with a similar appearance. The three different retinal images show in Fig. 1 were acquired with the same fundus camera. However, they have small dark spots in approximately the exact location on the images marked with arrows. If it were not for a close side-by-side comparison, the dark spots could be mistaken as small retinal lesions such as microaneurysms [4].

Dust particles can also produce unreliable counts in automated microaneurysm detection tasks, image restoration [5], or longitudinal change detection in retinal images [6,7] because microaneurysm sizes range from 10 to 125 microns in diameter [8]. The dust particle artifact problem is unavoidable and accumulative in clinical practice. It is possible to clean the lenses and the camera sensor, but it requires professional assistance and involves a lapse of inactivity. In practice, sensor cleaning is not done very often, and many images can be acquired before the sensor is cleaned or the technician notices the artifacts. Some cameras incorporate automatic defect map creation, but this does not solve the problem of contaminated images. Instead, it is oriented toward determining the need for maintenance [9].

The effect of dust particles in imaging has been described by Wilson et al. [10]. However, their work is focused on modeling the dust artifacts and not on their detection or removal. Li et al. [11] propose removing dust artifacts in focal stack image sequences by taking advantage of the slight shift in image position of dust artifacts relative to the focus setting. However, this detection approach requires careful scanning through many possible focus positions. Also, the removal stage is based on a fusion approach from the stack of images typically used in microscopy [12]. Alternatively, Zhou and Lin [13] propose the detection of dust

candidate regions via an optimization approach to estimate a dust-free image. The optimization scheme is dependent on several parameters, such as the camera settings. However, these parameters are not always known, and their approach may mistake retinal lesions, such as hemorrhages, for artifacts.

Artifact detection in retinal images has been mostly focused on artifacts such as flash flares [14], image quality assessment [15–18], or as a previous stage for automatic retinopathy screening [19]. However, the detection and removal of dust particle artifacts in retinal images have not been sufficiently studied. In recent years, several works have been proposed to detect dust artifacts in video sequences [20–22], where the common assumption is that the scene may change in time, but the artifacts remain in the same position. Dust particles can occlude the moving objects in a scene. Therefore, by analyzing the motion flow field, regions with no motion are candidate artifacts. However, this approach is not directly applicable to retinal images since optical flow algorithms require small displacements between consecutive frames [23]. Here we consider the general case in which several different retinal images acquired with the same fundus camera have dust artifacts.

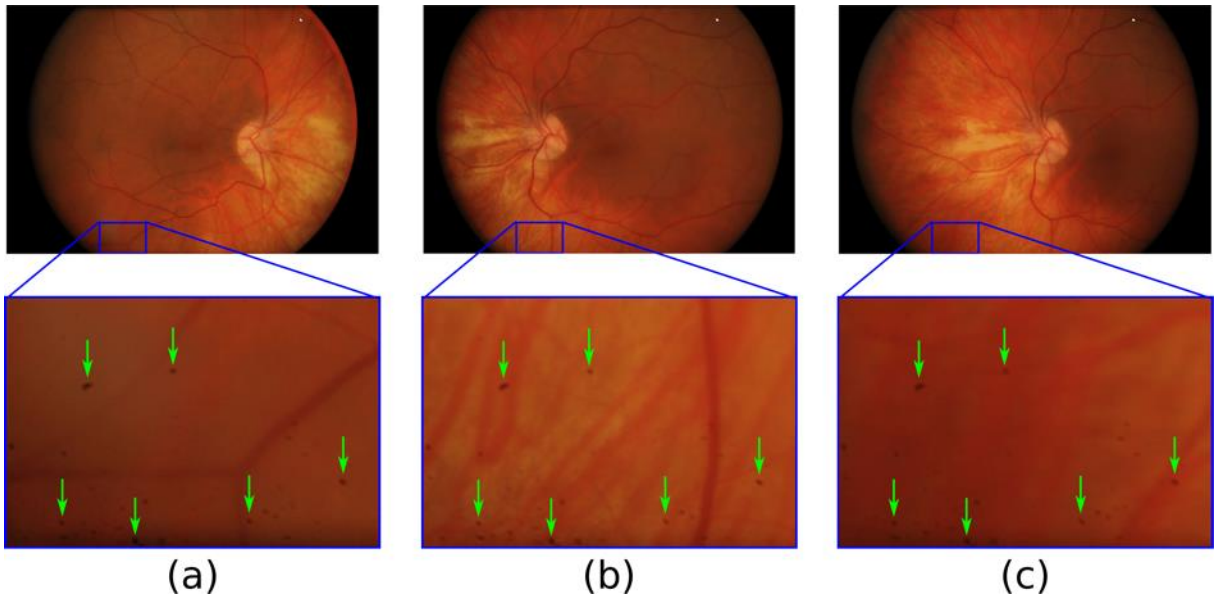

Fig 1: Close-up from three different retinal images acquired with the same fundus camera (real sample images from clinical practice). The green arrows indicate the dust particle artifacts (dark spots) in approximately the same pixel locations.

In this paper, we propose a method for detecting and removing dust particle artifacts in retinal images. Our algorithm compares detections on two or more input images to avoid false detections. This paper is a continuation of our previous work on dust artifact detection and removal [4]. Our new method detects more artifacts than the previous approach, and we now use a sparse-based inpainting approach for their removal, which achieves better image restoration. The algorithm proposed here can detect dust artifacts of a broader range of sizes and different shapes.

Additionally, we present a methodology of iterative segmentation based on inpainting. Precise segmentation of artifacts is challenging, mainly due to a blurred boundary, and over-segmentation is not acceptable. Our new segmentation approach is based on comparing an initial restored image and the original input image to achieve precise artifact segmentation. The restoration yields plausible results since small details are preserved after artifact removal, even in fine fundus structures such as blood vessels. This segmentation via inpainting approach may prove useful in abnormality detection [24, 25], including microaneurysms [26]. Finally, to comply with modern open and reproducible research practices, the archived version of the code can be freely accessed and executed through Code Ocean [27].

2. Method

We show an overview of the proposed method in Fig. 2. The input of the algorithm is a set of retinal images $(I_1, I_2, \dots, I_n)$ acquired with the same fundus camera. The images have artifacts caused by dust particles in some components of the camera.

The first stage consists in generating a mask M_0 that contains the locations and the estimated size of the artifacts. The detections are obtained with a correlation-based approach. The second stage consists of artifact segmentation and image restoration by inpainting. The algorithm predicts the values of the missing regions (artifacts) with sparse-based representation. The predicted values are replaced in the original image to obtain restored images $(I'_1, I'_2, \dots, I'_n)$. In the following sections, we describe the details of each stage.

2.a. Dust artifact detection algorithm

In Fig. 3, we show the block diagram of the artifact detection algorithm. It consists of two steps. First, the RGB input images are converted to grayscale, and the algorithm applies morphological processing to remove all the information that is not relevant to the artifacts. This step produces an image that contains only the artifacts and other small details, as shown in Fig. 4(c). In the second step, the algorithm searches for the dust artifacts via normalized cross-correlation, nonmaximal suppression, and discards false detections through spatial matching, i.e., we consider true detections those present at the same location in all input images. The output of this stage is a matrix M_0 that contains the location and estimated size of the detected artifacts.

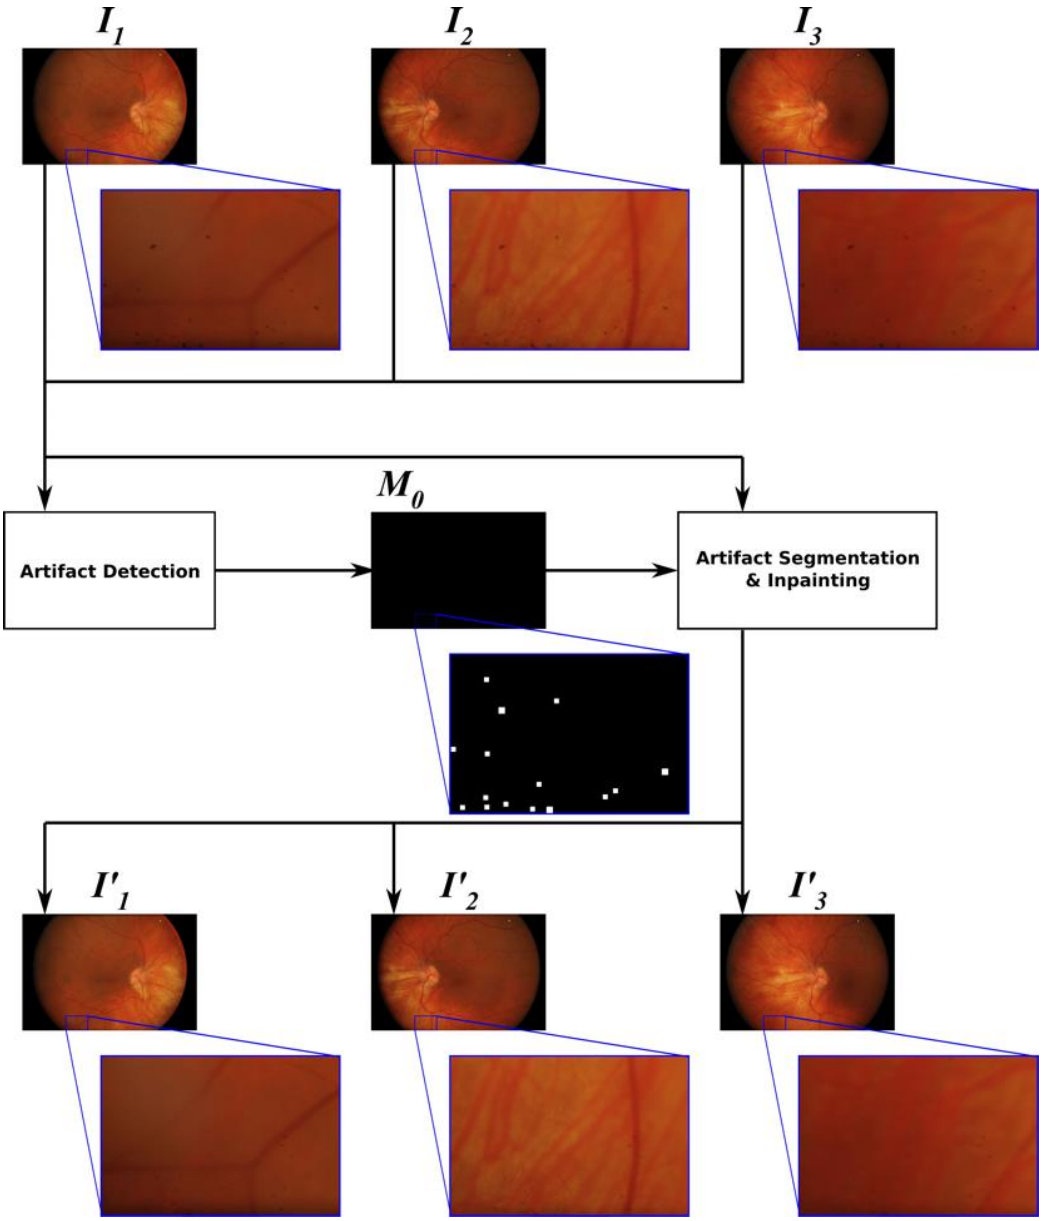

Fig 2: Block diagram illustrating the proposed method. I_i are the input images, I'_i are the restored versions. M_0 is the initial segmentation mask resulting from the artifact detection stage.

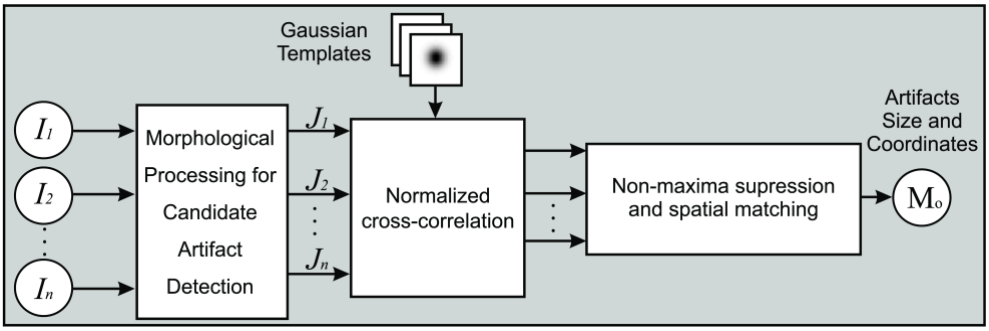

Fig 3: Artifact detection algorithm diagram, where I_i are the input images and M_0 is the output mask with the location and size of the detected dust artifacts.

Morphological processing for candidate artifact detection

The algorithm performs a morphological closing (a dilation followed by erosion) to every grayscale image with a circular structuring element. This operation roughly removes dark details that are smaller than the structuring element. For the resolution of our images (2784x1846), we empirically found that a circular structural element of 20 pixels in diameter is large enough to cover all possible dust artifact sizes. This setting parameter needs to be established from sample images at the initialization stage. The dilation $\oplus$ of a gray-scale image I by a structuring element b is defined as [28],

$$(I \oplus b)(x,y) = \max\{I(x-x', y-y') | (x',y') \in D_b\}, \quad (1)$$

and the erosion $\ominus$ of a gray-scale image I by a structuring element b is defined as,

$$(I \ominus b)(x,y) = \min\{I(x-x', y-y') | (x',y') \in D_b\}, \quad (2)$$

where D_b is the domain of b. Therefore, the closing of an image I by a structuring element b, denoted as $I \cdot b$, is given by

$$I \cdot b = (I \oplus b) \ominus b. \quad (3)$$

To correctly remove the artifacts at the borders of the retinal image, we pad around the image before applying the closing operation. The pad must be similar to the content of the image. We use symmetric replication which pads the array with a mirror reflection of itself. After morphological processing, the image is cropped to the original dimensions.

For each input image I, the algorithm calculates the difference between the original grayscale image and the image after applying the closing operation as

$$J = I - (I \cdot b), \quad (4)$$

where J is called the difference image and only contains the small details that were removed in the closing step, including the dust artifacts. In Fig. 4, we illustrate the morphological processing procedure. Note that the artifacts are better contrasted against the background and are more easily detectable in the resulting difference image J, as shown in Fig. 4(c).

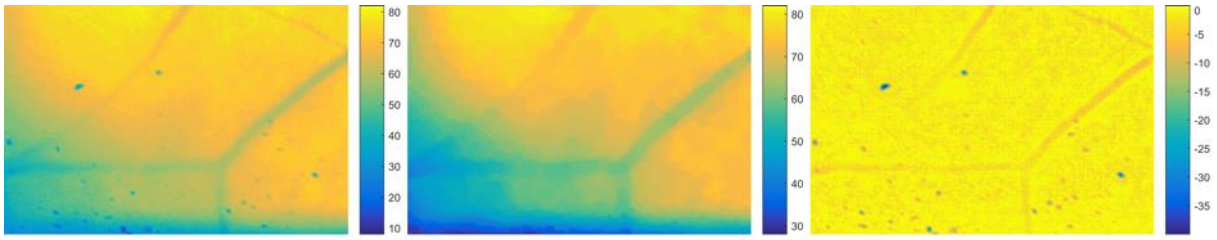

Fig 4: Morphological processing for candidate artifact detection. (a) Pseudocolor original grayscale image. (b) Image after morphological closing. (c) Difference image *J* in which the artifacts are easier to identify.

Artifact localization via cross-correlation, non-maxima suppression, and spatial matching

To obtain the coordinates of the dust artifacts, the algorithm computes the normalized cross correlation $\gamma(u,v)$ of each difference image with a set of templates as given by

$$\gamma(u,v)=\frac{\sum_{x,y}\left[J(x,y)-\bar{J}_{u,v}\right]\left[t(x-u,y-v)-\bar{t}\,\right]}{\left\{\sum_{x,y}\left[J(x,y)-\bar{J}_{u,v}\right]^2\sum_{x,y}\left[t(x-u,y-v)-\bar{t}\,\right]^2\right\}^{0.5}}, \tag{5}$$

where J is the difference image, $\bar{J}_{u,v}$ is the mean of $J(x,y)$ values in the region under the template placed at (u,v), t is the template, and $\bar{t}$ the mean of the template [29].

To find the location of artifacts, we use templates that resemble their typical shape. Among the properties of these artifacts, we find that dust particles are made of a single material, and their projections are approximately monochromatic [12]. Dust artifacts are darker in the center, and their opacity is attenuated far from the center, and most of them are approximately circular in shape [9]. Moreover, since the search image J is a difference image, the artifact should have the minimum value in the center and increase smoothly from the center. Therefore, we use inverted Gaussian filters, defined as

$$t_\sigma(x,y)=-\frac{1}{2\pi\sigma^2}\exp\left(-\frac{x^2+y^2}{2\sigma^2}\right), \tag{6}$$

where σ is the parameter that controls the width of the Gaussian. Here, we use a set of Gaussian filters with different σ values for detecting artifacts of different sizes.

Every difference image J undergoes the same process. The algorithm performs one normalized cross-correlation for each Gaussian filter with the difference image. In Fig. 5, we illustrate the multi-scale cross-correlation with four Gaussian filters of different size. From the multi-scale correlation, the algorithm searches peak values that are higher than a predefined threshold via nonmaximal suppression. These peaks account for the artifact detection locations and the filter σ values for the artifact approximate size. If an artifact is detected with more than one filter, the σ value of the Gaussian filter with the highest cross-correlation value is chosen.

As a result, there are usually many spots detected for each input image. To distinguish between true and false artifacts, we assume that the true dust artifacts are in approximately the same locations on every image acquired with the contaminated camera. In other words, we consider true artifacts the detections common to all input images within a 5-pixel radius. Otherwise, detections not common to all input images are discarded. This condition is a cautious approach considering the consequences of possible false positives, e.g., detecting a microaneurysm as an artifact.

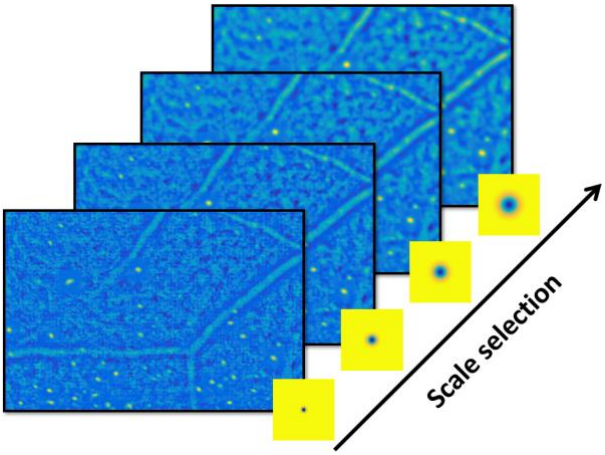

Fig 5: Normalized cross-correlation between the difference image J and four Gaussian filters of different size.

In Figures 6(a)-(c) we show an example of dust artifact detection on the three close-ups shown in Fig. 1. The green lines connect the detection common to all images. The remaining detections are discarded. In Figures 6(d)-(f) we show the detected artifacts (marked with green squares) after discarding false detections. Most artifacts were successfully detected, while several small artifacts were not classified as true detections because they were not consistently detected in all input images. Although this is a limitation of the method, this condition limits false positives which may hinder the overall performance of the method. The result of the detection stage is a matrix M_0 that contains the location and estimated size of the detected artifacts. The matrix M_0 is the same size as the input images. This matrix is one of the inputs of the segmentation and inpainting stage.

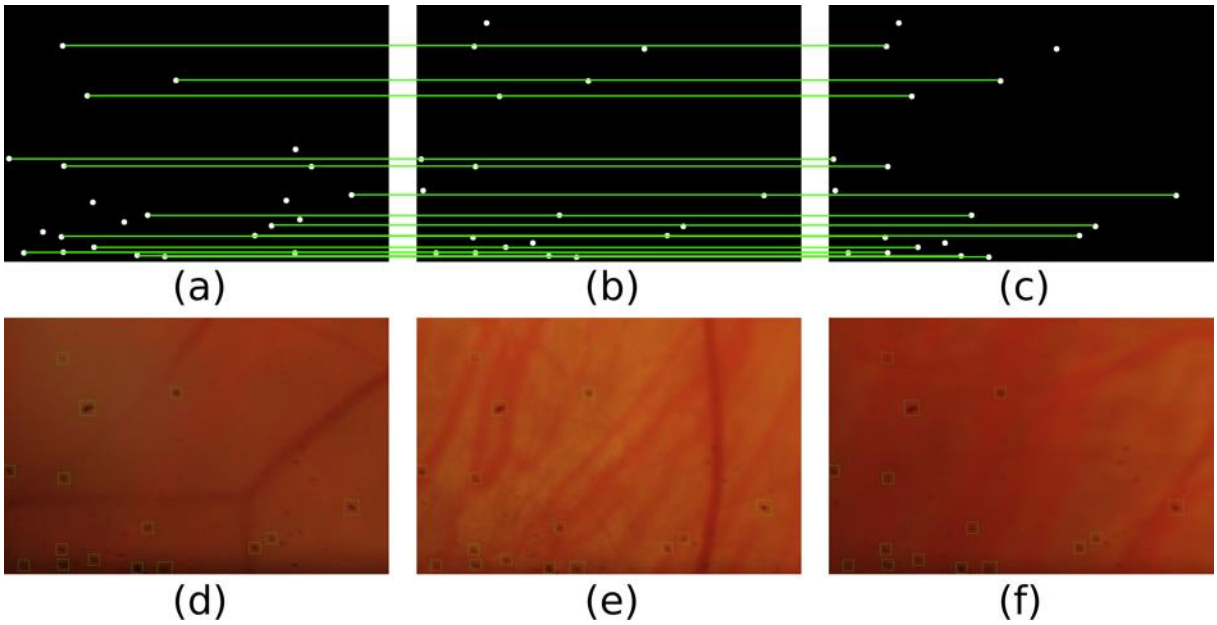

Fig 6: (a)-(c) Candidate artifact detections on three input images. (d)-(f) Detected artifacts (marked with green squares) after discarding false detections. We consider true artifacts the detections common to all input images. This figure is best viewed by zooming in the electronic version.

2.b. Artifact segmentation and inpainting

Precise artifact segmentation is useful as a preprocessing stage or as an image quality estimation index [30]. The digital removal of the artifacts implies an estimation of the missing pixels of the image while maintaining the structures and other characteristics of the neighboring regions. If a digital restoration is performed on a medical image, this modification should be informed to the specialist assessing the image. However, we believe that the restored image should help the specialist or the computer-aided diagnosis system in reducing false-positive detections, e.g., in the detection of microaneurysms [31].

Inpainting

In medical images, it is important to have well-defined inpainting criteria. In general, we can classify inpainting methods into three categories: (i) diffusion-based methods, (ii) exemplar-based methods, and (iii) hybrid methods [32]. The first category consists of methods based on parametric models or partial differential equations used to propagate (or diffuse) the local structures of the exterior to the interior of the target region. The biggest drawback of these methods is the failure to reconstruct texture. As a result, the reconstructed regions are smooth and lacking the surrounding texture. The second category consists of exemplar-based methods. These methods are defined from image patches. They assume that the statistics of the texture in the image are stationary (in the case of random texture) or homogeneous (in the case of a regular pattern). The texture to be synthesized is learned from similar regions in a texture sample or a known part of the image. Finally, several authors have proposed hybrid methods as approximations that attempt to combine the two previous categories [33]. These methods allow inpainting with texture information while maintaining the linear structure simultaneously. However, such methods rely on the statistics and structure often found in natural and human-made scenes.

Since we deal with medical images, we have chosen an exemplar-based approach with sparse representation and dictionary learning to remove the dust artifacts found in retinal images, which has been found appropriate in a previous work [34]. The initial inpainting mask is produced from the σ values from the detected artifacts. We found that a value of 3σ provides an initial mask that covers the whole artifact. The mask is also used to avoid including the artifact pixels in the dictionary learning process.

The first stage of the inpainting method relies on patch-based processing of the image. Let vector $I \in \mathbb{R}^N$ represent the original image of size N, and $n \leq N$ be the number of patches in the image. Let us consider $I_k = R_k(I)$, in which $R_k(.)$ is an operator that extracts the k^{th} patch from the image I, and its transpose, denoted by $R_k^T(.)$, is the operator that allows the reconstruction of the image I from the set of patches $\{I_k\}$ through the following equation,

$$I = \left(\sum R_k^T(R_k)\right)^{-1}\left(\sum R_k^T(I_k)\right). \tag{7}$$

A sparse representation seeks to determine the subspaces where the most relevant data of a signal are found. Resulting in a reduction of the dimensionality of the signal. The principle of a sparse representation relies on approximating a signal I of dimension L as a linear combination of a small number of signals taken from a base called the dictionary D of size $N \times L$, which contains L elements of length N, often called atoms ($D = [d_1, d_2, \ldots, d_L]$). The sparse vector $x \in \mathbb{R}^L$ contains the signal representation coefficients. In our case, $I_k \approx Dx_k$, we are interested in generating a patch-based sparse representation of the image given by

$$I \approx \left(\sum R_k^T(R_k)\right)^{-1}\left(\sum R_k^T(Dx_k)\right), \tag{8}$$

where I_k has been replaced by Dx_k.

The second stage consists in generating a dictionary from the image patches that do not contain artifacts. The dictionary should be compact and satisfy the condition for sparsity [35]. The dictionary learning is carried out by solving the following optimization problem

$$\min_{D,x_k} \sum_{k=1}^{N} \|I_k - Dx_k\|_2^2, \text{s.t. } \|x\|_0 \leq T_0, 1 \leq k \leq N\,, \tag{9}$$

where $\|.\|_2$is the ℓ_2-norm, $\|.\|_0$is the ℓ_0-norm, and T_0 limits the sparsity of the representation coefficients. The dictionary learning problem defined in Eq. (9) can be solved using the Method of Optimal Directions (MOD) [36], K-SVD [35], among others. For this problem, we used the K-SVD algorithm and Orthogonal Matching Pursuit (OMP) [37] algorithm to obtain the x_k sparse vectors. We have found that the K-SVD

algorithm is simple, flexible, and works in conjunction with OMP. It is highly efficient due to effective sparse coding and an accelerated dictionary update method [35].

In the third stage, the algorithm predicts the values of the missing regions (artifacts) by constructing an approximate image without artifacts. Afterward, these values are replaced in the original image I to obtain the restored image I'. The dictionary atoms and the sparse coefficients given by $\hat{d}_k$, $\hat{x}_k$, respectively, are the optimal solution. Once the sparse coefficients $\hat{x}_k$ for all patches and the learned dictionary atoms $[\hat{d}_1, \hat{d}_2, \ldots, \hat{d}_L]$ have been estimated, we obtain the best approximation to the original image without artifacts $\hat{I}$. This approximation is based on Eq. (8), so that

$$\hat{I} \approx \left(\sum R_k^T(R_k)\right)^{-1} \left(\sum R_k^T(\hat{D}\hat{x}_k)\right). \tag{10}$$

Taking into account that the patches generally overlap, Eq. (10) describes a weighted average as a reconstruction strategy. Subsequently, an indexing approach is carried out which consists in replacing the missing pixels in I (original image) by the values of $\hat{I}$ in the positions of the artifacts, thus obtaining the restored image I'.

Artifact Segmentation Refinement

Recent advances in inpainting techniques are mostly focused on different ways of filling the target region. However, they assume that the inpainting mask is correct and does not need to be updated during the inpainting process. A successful image restoration depends to a great extent on a precise mask. If the mask is too large, it will end up removing useful information. However, if the mask is too small, the algorithm will fail in removing the desired pixels. This issue is particularly relevant for the removal of artifacts which do not have a well-defined boundary. Therefore, we propose a new stage to further improve the method by refining the segmentation mask.

In Fig. 7, we present an approach for updating the inpainting mask, which we call segmentation via inpainting (SVI). We highlight in red the new mask refinement approach. It resembles a feedback control loop, in which the error signal is obtained from the difference between the input signal, i.e., the original image I, and the corresponding restored image I output. The idea is to threshold the image difference to accurately segment the artifacts. The difference image is expected to have a maximum value in the center of the artifact and decrease rapidly from the center. For this segmentation, we found that Otsu's method [38] yielded reliable results to produce a new mask M_i. This mask iteratively shrinks towards a more accurate segmentation of the artifact.

In Fig. 8, we show a close-up of the image shown in Fig. 1(a). Although the locations of the artifacts in the initial mask M_0 are correct, the size and shape of the artifacts are overestimated (Fig. 8(b)). In Fig. 8(c) we present the segmentation by the proposed mask refinement method. Note the improvement in the segmentation of the artifacts.

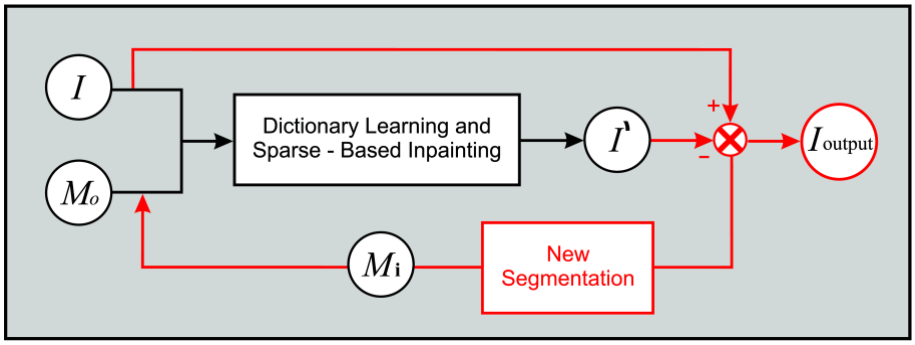

Fig 7: Segmentation refinement based on inpainting. The new mask M_i is computed from the difference of the restored image and the original image. The final restored image is obtained from a refined inpainting mask.

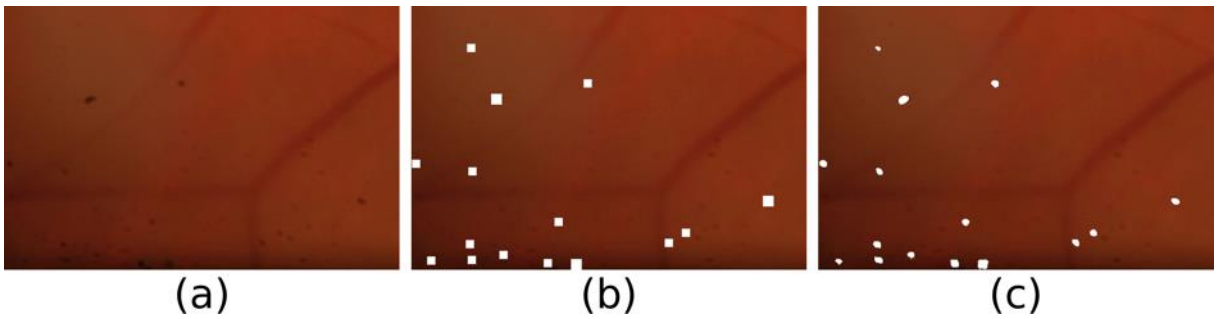

Fig 8: Segmentation of artifacts via inpainting. (a) Close-up of image shown in Fig. 1(a). (b) Initial detection and segmentation mask M_0. (c) Refined segmentation (M_i).

3. Experiments and Results

Two experiments were carried out. The objective of the first experiment was to determine a quantitative measurement of dust artifact detection performance and to evaluate the image restoration qualitatively. The second experiment consisted of comparing the detection and restoration results of the proposed method with the method described in Ref. [4]. An implementation of the proposed algorithm can be accessed and executed through a Code Ocean compute capsule [27].

3.a. Image corpus

A set of 8 fundus images from the clinical practice were used. They were all grouped in pairs acquired in the same session with the same contaminated fundus camera. They are color RGB 24 bit-depth fundus images of size 2784×1846 digitized in JPEG format.

3.b. Experiment 1

All pair of retinal images were processed for artifact detection and removal with the proposed method. To illustrate a complete realization of the algorithm, in Fig. 9, we show a pair of retinal images from the same eye but with different framing. The green squares mark the detected artifacts. There are many artifacts located throughout the retinal images. Many are on top of retinal structures like blood vessels, the optic disc, and even near the macula. The artifacts were detected with four Gaussian templates with $\sigma = \{1, 2, 3, 4\}$. The number of detected artifacts (> 100) and their distribution throughout the fundus image reveals the severity of the problem.

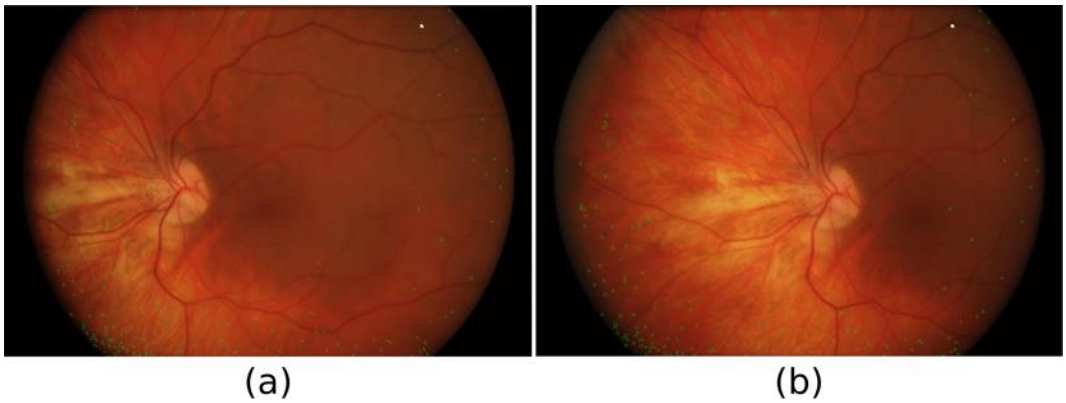

Fig 9: (a) and (b) Detected artifacts (green small squares) on a pair of retinal images acquired with the same fundus camera. Both pictures correspond to the same eye but with different framing. The figure is best viewed by zooming in to see the artifact detections.

Due to the lack of ground truth detections and a large number of artifacts present in most images, we cropped sample regions from the test images. These were selected to have varying amounts of artifacts and

located in different parts of the images. Each region was assessed by one individual by manually selecting the artifacts, thus obtaining the ground-truth. In Fig. 10, we show six different cases of image regions with artifacts. Each case has two corresponding regions from the two input images. From left to right, the first (image A) and third (image B) columns are the regions of the input images with the detected artifacts marked with green squares. The second and fourth columns are the restored versions of images A and B, respectively.

In case 1 of Fig. 10, 3 out of 3 dust artifacts were successfully detected. In this case, image B has 2 dust artifacts located over blood vessels (pointed out by green arrows), which were correctly detected even though one of them is small. The dust artifact located on the blood vessel on the right side of the region was satisfactorily removed while keeping the vessel structure intact. Similarly, in case 2, the algorithm detected all 4 dust artifacts even though they do not have high contrast, and one of them occludes a blood vessel in image A (see the green arrow). In image A, a single artifact located on top of a blood vessel was successfully detected. The inpainting result was satisfactory. The other dust artifacts were located in uniform regions with no texture, which facilitates their detection and removal.

Cases 3 and 4 show the detection of artifacts of different sizes, which are also occluding vessels, including barely noticeable ones. All artifacts were correctly detected, segmented, and removed. In case 3, 12 out of 12 dust artifacts were detected. In case 4, 7 out of 7 artifacts were detected, including two small artifacts located in a dark region. In both cases, there are also artifacts located on top of blood vessels, which were satisfactorily restored.

In case 5, 20 out of 30 dust artifacts of different sizes were detected. The undetected 10 artifacts have low contrast. The algorithm was able to remove the detected artifacts regardless of their size and appearance. However, the inpainting area is slightly noticeable in two artifacts at the bottom edge of the image where the surrounding area was darker (see the two green arrows in the restored version of image A). Probably because this edge region is not well represented in the learned dictionary. An extreme situation is case 6, in which we show a widely contaminated region. There are many dust artifacts of different sizes. The algorithm successfully detected 29 out of 41 dust artifacts. The undetected artifacts are mostly small and not well contrasted. The inpainting process removed the detected dust artifacts without leaving noticeable marks.

In the six cases presented in Fig. 10, there are a total of 97 dust particle artifacts, of which 75 were detected, i.e., the results present a specificity of 77.32% with an average artifact detection rate of 90%. No false positives were obtained in the results of the dust artifacts detection phase. The false positives that might have resulted from the normalized cross-correlation were discarded by comparing several images taken with the same camera (spatial matching). Due to the complexity of the problem, this detection rate is sufficiently good either for image quality estimation or for the complete process of artifact detection and removal. However, further research with a larger dataset may be required to provide a more accurate detection rate. Nevertheless, from our tests, the algorithm often detects most artifacts. If used as a binary classifier for determining the presence of artifacts, the classification rate would be much higher.

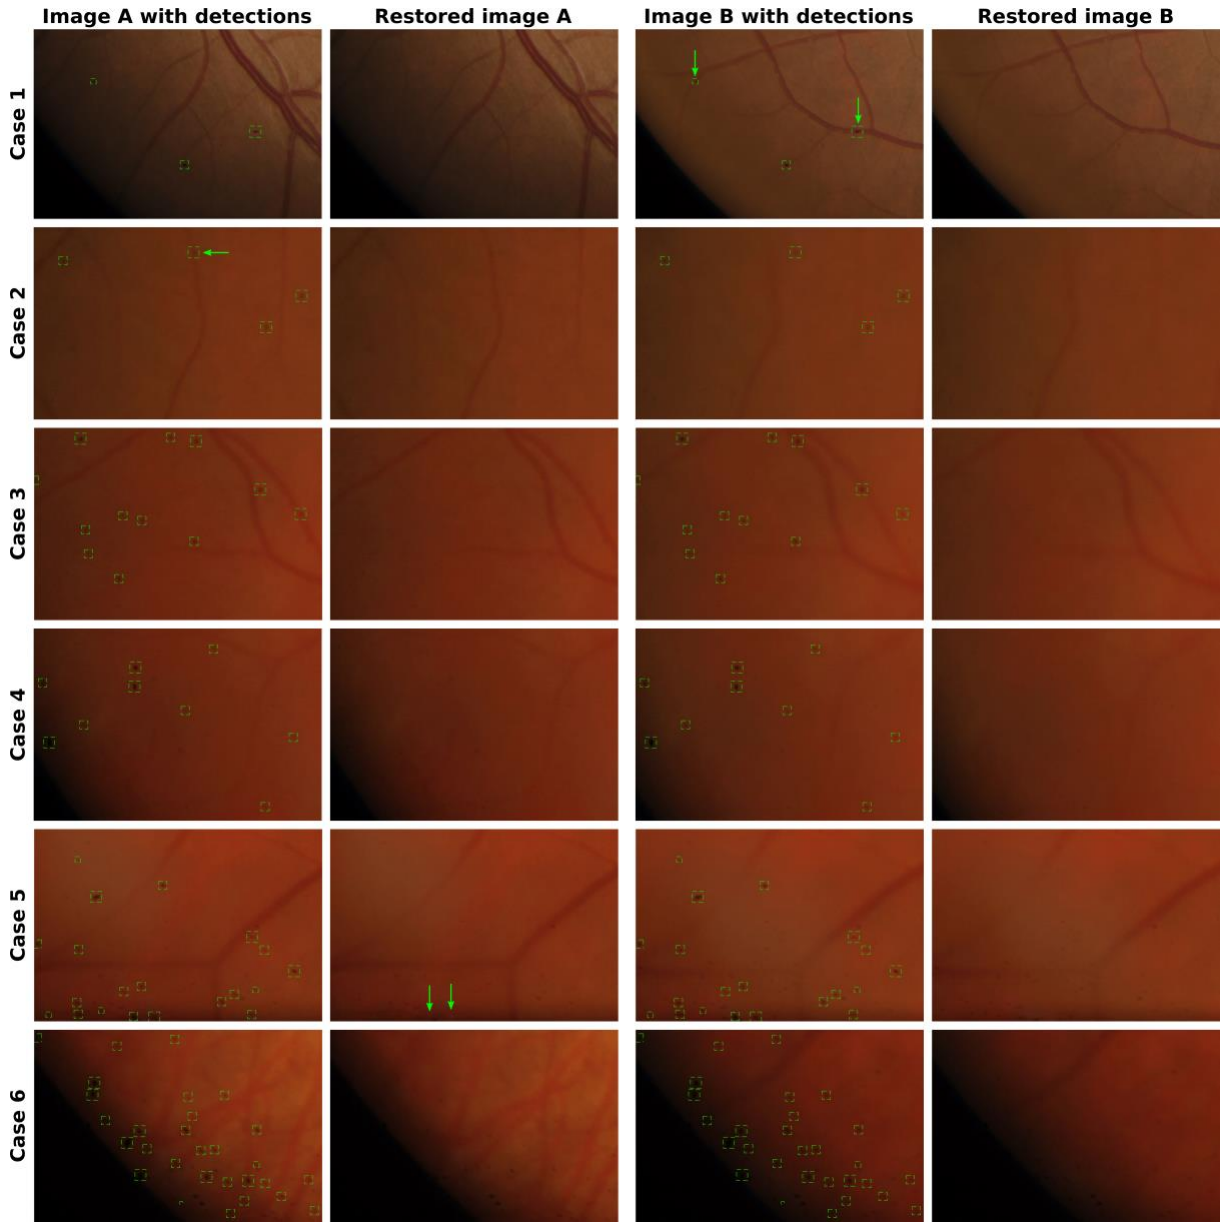

Fig 10: Six cases, one for each row, of contaminated regions. For each case, the first and third image, left to right, are the input images with green squares that mark the detected artifacts. The second and fourth image correspond to the result after inpainting of the first and third image, respectively. Dust artifact count ground truth vs detected: Case 1, 3/3; case 2, 4/4; case 3, 12/12; case 4, 7/7; case 5, 20/30; case 6, 29/41. Green arrows point out details described in the text.

3.c. Experiment 2

In Fig. 11, we show a comparison between the results of our previous dust artifact detection method [4] (Fig. 11(a)) and the results of the method proposed in this paper (Fig. 11(b)). The previous method detected 16 out of 41 artifacts, while the new method detected 29. The proposed method is better at detecting dust artifacts in darker regions, and at estimating the correct artifact size. The previous method used an average template obtained from artifacts, whereas a filter bank of templates describes better the different artifacts and allows for size estimation.

To assess the restoration quality of the proposed method, in Fig. 12 we show a comparison between the results of the inpainting by a diffusion method [4] (Fig. 12(b)) and by the proposed sparse-based inpainting method (Fig. 12(c)). Note that with the diffusion method, the restoration of artifacts located on top of blood vessels lead to an image where vessels are smoothed out. Whereas with the proposed method, the blood vessel structure remains continuous without introducing new artifacts.

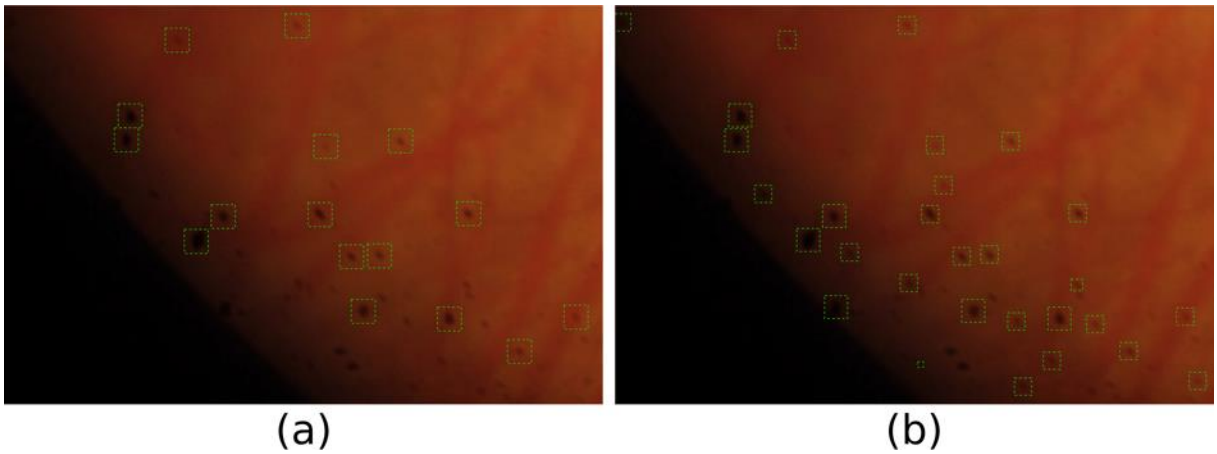

Fig 11: Dust artifact detection method comparison on case 6 from Fig. 10. Green squares mark the detected artifacts and their size. (a) Detected dust artifacts with previous method [4] 16/41. (b) Detected dust artifacts with the proposed method 29/41.

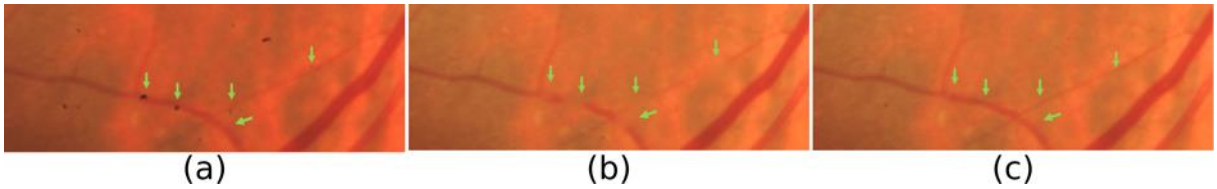

Fig 12: Inpainting results. (a) Original image with artifacts. (b) Inpainting result by diffusion method from Ref. [4]. (c) Inpainting result by the proposed method in which the blood vessel structure remains continuous.

4. Conclusions

Dust artifact detection in retinal images is needed for preliminary image analysis to avoid false-positive detections or to be used as an image quality criterion. They may appear in any image region and are generally difficult to detect as they are small and of low contrast. The method proposed in this paper achieves high detection rates, precise artifact segmentation, and state-of-the-art image restoration through sparse-based inpainting. Using two images to confirm the presence of artifacts allows us to rule out a possible artifact with a microaneurysm since the artifacts maintain their positions despite the fact that the images are different, not being the case of the microaneurysms. The proposed iterative segmentation via inpainting approach enables well-defined artifact segmentation, improving the restoration results even under challenging regions like blood vessels. Future work involves using this method as a preprocessing stage in other retinal image processing algorithms. An implementation of the proposed algorithm can be accessed and executed through Code Ocean.

Acknowledgements

The authors acknowledge the financial support from the Centre de Cooperació i Desenvolupament (CCD) at the Universitat Politècnica de Catalunya under project ref. CCD 2019-B004, and from the Universidad Tecnológica de Bolívar. Authors are grateful to Juan Luís Fuentes from the Miguel Servet University Hospital (Zaragoza, Spain) for providing the real images from clinical practice. E. Barrios thanks Minciencias and Sistema General de Regalías (Programa de Becas de Excelencia) for a PhD scholarship. E. Sierra thanks the Universidad Tecnológica de Bolívar for a post-graduate scholarship. Parts of this work were presented at the Pattern Recognition and Tracking XXX - SPIE DCS, 2019 [39]. L. Romero, A. Marrugo, and M.S. Millán thank the funds provided by the Spanish Ministerio de Ciencia e Innovación under the project reference PID2020-114582RB-I00.

