Restoration of retinal images with space-variant blur

Andrés G. Marrugo
María S. Millán
Michal Šorel
Filip Šroubek

Restoration of retinal images with space-variant blur

Andrés G. Marrugo,[a,b,]* **María S. Millán,**[a] **Michal Šorel,**[c] **and Filip Šroubek**[c]
[a]Universitat Politècnica de Catalunya, Department of Optics and Optometry, Group of Applied Optics and Image Processing, Violinista Vellsolà 37, 08222 Terrassa, Spain
[b]Universidad Tecnológica de Bolívar, Facultad de Ciencias Básicas, Km 1 vía Turbaco, Cartagena, Colombia
[c]Academy of Sciences of the Czech Republic, Institute of Information Theory and Automation, Pod Vodárenskou věží 4, 18208 Prague 8, Czech Republic

Abstract. Retinal images are essential clinical resources for the diagnosis of retinopathy and many other ocular diseases. Because of improper acquisition conditions or inherent optical aberrations in the eye, the images are often degraded with blur. In many common cases, the blur varies across the field of view. Most image deblurring algorithms assume a space-invariant blur, which fails in the presence of space-variant (SV) blur. In this work, we propose an innovative strategy for the restoration of retinal images in which we consider the blur to be both unknown and SV. We model the blur by a linear operation interpreted as a convolution with a point-spread function (PSF) that changes with the position in the image. To achieve an artifact-free restoration, we propose a framework for a robust estimation of the SV PSF based on an eye-domain knowledge strategy. The restoration method was tested on artificially and naturally degraded retinal images. The results show an important enhancement, significant enough to leverage the images' clinical use. © *2014 Society of Photo-Optical Instrumentation Engineers (SPIE)*
[DOI: 10.1117/1.JBO.19.1.016023]

Keywords: blind deconvolution; space-variant restoration; image restoration; deblurring; retinal image.

Paper 130016RR received Jan. 11, 2013; revised manuscript received Dec. 23, 2013; accepted for publication Dec. 30, 2013; published online Jan. 27, 2014.

1 Introduction

Blur is one of the main image quality degradations in eye fundus imaging, which along with other factors such as nonuniform illumination or scattering, hinder the clinical use of the images. Its main causes are inherent optical aberrations in the eye, relative camera-eye motion, and improper focusing. Eye motion is related to the patient inability to steady fixate a target in the fundus camera. Many patients have difficulty in fixating, like many elderly patients or those that suffer from amblyopia.[1] Because the optics of the eye is part of the optical imaging system, the aberrations of the eye are a common source of image quality degradation. To overcome this limitation, adaptive optics techniques have been successfully applied to correct the aberrations, thus producing high resolution images.[2] However, most commercial fundus cameras compensate for spherical refractive errors, but not for astigmatism[3]—let alone higher-order aberrations. In general, the aberrations of the eye have a stronger impact in image degradation than the aberrations introduced by the rest of the optical system, i.e., the retinal camera. Besides, even though it is possible to measure the optical quality of the camera, it would be exceptional to have readily available additional information related to the optical quality of the patient's eye. The described scenario is commonplace in the clinical setting, for which we assume the same conditions here. This brings about the need for a restoration procedure that accounts for the lack of information related to the origin of the image degradation.

The technique for recovering an original or unblurred image from a single or a set of blurred images in the presence of a poorly determined or unknown point-spread function (PSF) is called blind deconvolution. Removing blur from a single blurred image is an ill-posed problem as there are more unknowns (image and blur) than equations. Having more than one image of the same scene better poses the problem. In retinal imaging, it is not difficult to obtain a second image from the same eye, with the convenience that acquisition conditions remain quite similar. In fact, in Ref. 4, we took advantage of this condition and proposed a blind deconvolution method to restore blurred retinal images acquired with a lapse of time, even in the case where structural changes had occurred in the images. In that work, we detected structural changes, which in turn have clinical relevance, and applied a masking operator so that the images would comply with the considered degradation model. This enabled the successful restoration of many degraded retinal images coming from patient follow-up visits. However, the method is limited to images blurred uniformly; in other words, we assumed the blur to be space-invariant. In Sec. 4.2, we show an attempt at restoring an image degraded with spatially variant blur with this approach. The space-invariant assumption is commonplace in most of the restoration methods reported in the literature,[5] but in reality it is a known fact that blur changes throughout the image.[6] In this work, we consider the blur to be both unknown and space-variant (SV). This in itself is a novel approach in retinal imaging; relevant to such extent that many common eye related conditions, such as astigmatism, keratoconus, corneal refractive surgery, or even tear break-up, may contribute significantly to a decline in image quality[7,8] typically in the form of an SV degradation. An example of such a condition is shown in Fig. 1(a). The image corresponds to an eye from a patient with corneal abnormalities that lead to a loss in visual acuity and a quality degradation of the retinal image [Fig. 1(b)].

Restoration of images with SV blur from optical aberrations has been reported in the literature,[9] although the main limitation

---

*Address all correspondence to: Andrés G. Marrugo, E-mail: agmarrugo@utbvirtual.edu.co
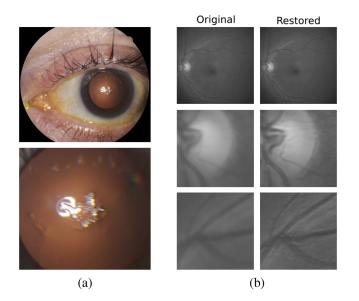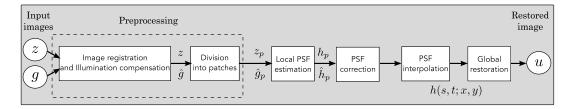
(a)   (b)

Fig. 1 (a) Top: Eye with corneal defects that induce retinal images with space-variant (SV) degradation. Bottom: zoomed region. (b) Left column: original image and details. Right column: restored image with proposed approach and details.
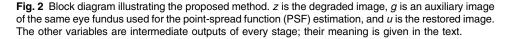
is that the blurred image is often restored in regions or patches, which are then stitched together. This inevitably leads to ringing artifacts associated with frequency-domain filtering like in Wiener filtering. Another clear disadvantage is a significant complexity for accurately estimating the SV PSF, for instance Bardsley et al.[10] use a phase-diversity based scheme to obtain the PSF associated with an image patch. This type of approach is common in atmospheric optics where the conditions and setup of the imaging apparatus (typically a telescope) are well known and calibrated. Unfortunately, this is not immediately applicable to retinal imaging, at least nonadaptive optics retinal imaging. Recently, there have been several works[11–13] that try to solve the SV blind deconvolution problem from a single image. The common ground in these works is that the authors assume that the blur is only due to camera motion. They do this in order to reduce the space in which to search for SV blurs. Despite their approach being more general, the strong assumption of camera motion is simply too restrictive to be applied in the retinal imaging scenario.

1.1 Contribution

In this work, we propose a method for removing blur from retinal images. We consider images degraded with SV blur, which may be due to factors like aberrations in the eye or relative camera-eye motion. Because restoring a single blurred image

is an ill-posed problem, we make use of two blurred retinal images from the same eye fundus to accurately estimate the SV PSF. Before the PSF estimation and restoration stages take place, we preprocess the images to accurately register them and compensate for illumination variations not caused by blur, but by the lighting system of the fundus camera. This is depicted in the block diagram shown in Fig. 2. The individual stages of the method are explained in Sec. 3.

We assume that in small image patches, the SV blur can be approximated by a spatially invariant PSF. In other words, that in a small region, the wavefront aberrations remain relatively constant; the so-called isoplanatic patch.[6] An important aspect of our approach is that instead of deblurring each patch with its corresponding space-invariant PSF—and later stitching together the results—we sew the individual PSFs by interpolation and restore the image globally. This is intended to reduce some artifacts that otherwise would likely appear at the seams of the restored patches. The estimation of the local space-invariant PSFs may fail in patches with hardly any structural information (e.g., such as blood vessels). These poorly estimated or nonvalid PSFs introduce artifacts in the restored image. Detecting such artifacts and inferring the nonvalid PSFs is a difficult problem. Recently, Tallón et al.[14] developed a strategy for detecting these patches in an SV deconvolution and denoising algorithm from a pair of images acquired with different exposures: a sharp noisy image with a short exposure and a blurry image with a long exposure. Because they had two distinct input images that were able to: (i) Identify patches where the blur estimates were poor based on a comparison (via a thresholding operation) of the deconvolved patches with the sharp noisy patches. (ii) In those patches, instead of correcting the local PSFs and deconvolving the patches again, they performed denoising in the noisy sharp image patch. The end result is a patchwork-image of deconvolved patches stitched together with denoised patches. Their method is mainly oriented at motion blur, this is the reason for a dual exposure strategy. This is not readily implementable in the retinal imaging scenario where the SV blur is generally caused by factors like aberrations, including those belonging to the patient's eye optical system, the eye fundus shape, and the retinal camera. In this paper, we address the question "how to identify PSF estimation failure to improve the SV deconvolution of retinal images?" Retinal imaging provides a constrained imaging scenario from which we can formulate a restoration approach that incorporates prior knowledge of blur through the optics of the eye. The novelty in our approach is in the strategy based on eye-domain knowledge for identifying the nonvalid local PSFs and replacing them with appropriate ones. Even though methods for processing retinal images in a space-dependent way (like locally adaptive filtering techniques[15,16])

Fig. 2 Block diagram illustrating the proposed method. z is the degraded image, g is an auxiliary image of the same eye fundus used for the point-spread function (PSF) estimation, and u is the restored image. The other variables are intermediate outputs of every stage; their meaning is given in the text.

have been proposed in the literature; to the best of our knowledge, this is the first time a method for SV deblurring of retinal images is proposed.

2 SV Model of Blur

In our previous work,[4] we modeled the blurring of a retinal image by convolution with a unique global PSF. This approximation is valid as long as the PSF changes little throughout the field of view (FOV). In other words, that the blurring is homogenous. In reality, we know that the PSF is indeed spatially variant,[6] to such an extent that in some cases the space-invariant approach completely fails, bringing forth the need for an SV approach. To address this limitation, in this work, we model the blurred retinal image z by the linear operation
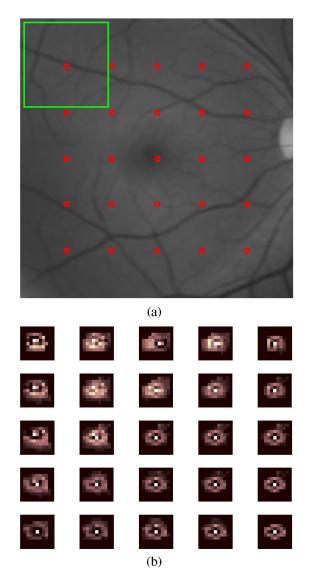
$$z(x, y) = [Hu](x, y) = \int u(s, t)h(s, t, x - s, y - t)\mathrm{d}s\mathrm{d}t,$$

(1)

where u is the unblurred retinal image and h is the SV PSF. The operator H is a generalization of standard convolution where h is now a function of four variables. We can think of this operation as a convolution with a PSF $h(s, t, x, y)$ that is now dependent on the position (x, y) in the image. Standard convolution is a special case of Eq. (1), where $h(s, t, x, y) = h(s, t)$ for an arbitrary position (x, y). Note that the PSF h is a general construct that can represent other complex image degradations which depend on spatial coordinates, such as motion blur, optical aberrations, lens distortions, and out-of-focus blur.

2.1 Representation of SV PSF

An obvious problem of spatially varying blur is that the PSF is now a function of four variables. Except trivial cases, it is hard to express it by an explicit formula. Even if the PSF is known, we must solve the problem of a computationally efficient representation.

In practice, we work with a discrete representation, where the same notation can be used but with the following differences: the PSF h is defined on a discrete set of coordinates, the integral sign in Eq. (1) becomes a sum, operator H corresponds to a sparse matrix and u to a vector obtained by stacking the columns of the image into one long vector. For example in the case of standard convolution, H is a block-Toeplitz matrix with Toeplitz blocks and each column of H corresponds to the same kernel $h(s, t)$.[17] In the SV case that we address here, as each column of H corresponds to a different position (x, y), it may contain a different kernel $h(s, t, x, y)$.

In retinal imaging, all typical causes of blur change in a continuous gradual way,[18] which is why we assume the blur to be locally constant. Therefore, we can make the approximation that locally the PSFs are space-invariant. By taking advantage of this property, we do not have to estimate local PSFs for every pixel. Instead, we divide the image into rectangular windows and estimate only a small set of local PSFs [see Fig. 3(a)] following the method described in Ref. 4 and outlined in Sec. 3. The estimated PSFs are assigned to the centers of the windows from where they were computed. In the rest of the image, the PSF h is approximated by bilinear interpolation from the four adjacent local PSFs. This procedure is explained in further detail in the following section.

(a)

(b)

Fig. 3 (a) Retinal image degraded with artificial SV blur given by (b) grid of PSFs. The grid and the image patches shown in (a) are used for local PSF estimation.

3 Description of the Method

In this section, we describe the different stages of the proposed restoration method shown in Fig. 2. This paper follows Ref. 4 but addresses a more general problem: restoration of retinal images in the presence of an SV PSF. In Ref. 4, we showed that the single image blind deconvolution for blurred retinal images does not provide a suitable restoration. Moreover, in images with SV blur, the restoration is even worse. Alternatively, by taking two images of the same retina we can use a multichannel blind deconvolution strategy that is mathematically better-posed.[19]

In this paper, the estimation of the SV PSF is carried out via local multichannel deconvolution. To illustrate the method and to study its dependence on its tunable parameters, we use an original real image of the retina and obtain two artificially degraded versions from it, denoted by z and g. Figure 3(a) contains image z. The degraded images z and g have been obtained by blurring the original image with an SV PSF represented by the grid of local PSFs shown in Fig. 3(b) and adding Gaussian

zero-mean noise ($\sigma = 10^{-6}$). The PSF grid was built with realistic PSFs estimated from real blurred retinal images using the method of Ref. 4.

3.1 Preprocessing

Because we use a multichannel scheme for the estimation of the local PSFs, the images are preprocessed so that they meet the requirements imposed by the space-invariant convolutional model given by Eq. (2). This consists in registering the images and adjusting their illumination distribution following Ref. 4. By carrying out this procedure, the remaining radiometric differences between the images are assumed to be caused by blur and noise. Unlike the case considered in Ref. 4, where a relatively long lapse of time between the two image acquisitions may possibly involve a potential structural change, in this study, both the images z and g are originated from the same image or, in practice, they are acquired one shortly after the other. Therefore, no structural change is expected. Since image g is registered and its illumination matched to z, we denote this transformed auxiliary image as $\hat{g}$.

3.2 Estimation of the Local PSFs

In Sec. 2, we described the model for a spatially varying blur in which we assume the PSF h to vary gradually, which means that within small regions the blur can be locally approximated by convolution with a space-invariant PSF. For this reason, we approximate the global function h from Eq. (1) by interpolating local PSFs estimated on a set of discrete positions. The main advantage of this approach is that the global PSF needs not be computed on a perpixel basis which is inherently time-consuming.

The procedure for estimating the local PSFs is the following. We divide the images z and $\hat{g}$ with a grid of $m \times m$ patches [Fig. 3(a)]. In each patch p, we assume a convolutional blurring model where an ideal sharp patch u_p originates from two degraded patches z_p and $\hat{g}_p$ (for $p = 1, \ldots, m \times m$). The local blurring model is

$$z_p = h_p * u_p + n,$$
$$\hat{g}_p = \hat{h}_p * u_p + \hat{n}, \qquad (2)$$

where $*$ is the standard convolution and h_p and $\hat{h}_p$ are the convolution kernels or local PSFs. The noise (n and $\hat{n}$) is assumed to have a constant spectral density and a zero-mean Gaussian distribution of amplitude. Despite the fact that this may not be the most accurate representation of the noise, because the retinal images considered here are acquired by illuminating with a flash, the resulting signal-to-noise ratio is high enough that in the estimation of the PSFs, the impact of noise is not significant.

From this model, we can estimate the local PSFs with an alternating minimization procedure as described in Ref. 4 but applied locally. The general guideline is that the patch size should be large enough to include retinal structures and much larger than the size of the local PSF. In Sec. 4, we show further analysis on the robustness of the method to these parameters. Every local PSF is computed on each patch p by minimizing the functional

$$\arg \min_{u_p, h_p, \hat{h}_p} \left(\frac{1}{2} \| u_p * h_p - z_p \|^2 + \frac{1}{2} \| u_p * \hat{h}_p - \hat{g}_p \|^2 \right.$$
$$\left. + \lambda_1 \int |\nabla u_p| \mathrm{d}x \mathrm{d}y + \lambda_2 \| z_p * \hat{h}_p - \hat{g}_p * h_p \|^2 \right), h_p, \hat{h}_p(s,t) \geq 0,$$
$$(3)$$

with respect to the ideal sharp patch u_p and the blur kernels h_p and $\hat{h}_p$. The blur kernel $h_p(s, t)$ is an estimate of $h(s, t, x_0, y_0)$, where (x_0, y_0) is the center of the current window z_p, and $\|.\|$ is the ℓ_2 norm. The first and second terms of Eq. (3) measure the difference between the input blurred patches (z_p and $\hat{g}_p$) and the sharp patch u_p blurred by kernels h_p and $\hat{h}_p$. This difference should be small for the correct solution. Ideally, it should correspond to the noise variance in the image. Although u_p is a restored patch, note that it is not used by our method, but discarded. This is because our method does not work by performing local deconvolutions and sewing restored patches together, which in practice would produce artifacts on the seams. Instead, we perform a global restoration method explained in Sec. 3.5. The two remaining terms of Eq. (3) are regularization terms with positive weighting constants λ_1 and λ_2, which we have set following the fine-tuning procedure described in Ref. 4. The tuning procedure consists of an optimization process where an artificially degraded retinal image is restored by varying λ and measuring a restoration error. This way an optimal λ is obtained. Typical values are $\lambda_1 = \sim 10^3$ and $\lambda_2 = \sim 10^1$. The third term is the total variation of u_p. It improves stability of the minimization and from a statistical perspective, it incorporates prior knowledge about the solution. The last term is a condition linking the convolution kernels which also improves the numerical stability of the minimization. The functional is alternately minimized in the subspaces corresponding to the images and the PSFs. The estimated PSFs for the artificially degraded [Fig. 3(a)] image are shown in Fig. 6(a).

3.3 Identifying and Correcting Nonvalid PSFs

3.3.1 Strategy based on eye-domain knowledge

The local PSF estimation procedure does not always succeeds. Consequently, such nonvalid PSFs must be identified, removed, and replaced. In our case, we replace them by an average of adjacent valid kernels. The main reason why the PSF estimation may fail is due to the existence of textureless or nearly homogenous regions bereft of structures with edges (e.g., blood vessels) to provide sufficient information.[14] To identify these nonvalid PSFs, we devised an eye-domain knowledge strategy. The incorporation of proper a priori assumptions and domain knowledge about the blur into the method provides an effective mechanism for a successful identification of poorly estimated PSFs.

The optics of the eye is part of the imaging system, therefore it is reasonable to assume that the PSF of the imaging system is determined by the PSF of the eye. The retinal camera can indeed be close to diffraction limited with a very narrow PSF, but the optics of the eye is governed by optical aberrations that change across the visual field[18] that lead to an SV PSF. The typical PSFs of the human eye, as reported in the literature,[18,20] display distinct shapes in many cases displaying long tails' evidence of the inherent optical aberrations. What is common to all PSFs of the human eye is that the energy is spread from a central lobe and

decreases outwardly. In this paper, we assume that PSFs that do not follow this general pattern are nonvalid PSFs which correspond to patches where the estimation failed.

In order to prove this, we designed an artificial experiment where we compare the estimated PSFs with the ground-truth PSFs by using a kernel similarity measure S proposed by Hu and Yang.[21] The measure is based on the peak of the normalized cross correlation between the two PSFs. The authors showed in the paper that this measure is more accurate than the root mean square error, especially because it is shift invariant. The measure is defined as the blur kernel similarity $S(h, \tilde{h})$ of two kernels, h and $\tilde{h}$,

$$S(h, \tilde{h}) = \max_{\gamma} \rho(h, \tilde{h}, \gamma), \qquad (4)$$

where $\rho(.)$ is the normalized cross-correlation function and γ is the possible shift between the two kernels. Let τ represent element coordinates, $\rho(.)$ is given by

$$\rho(h_p, \tilde{h}_p, \gamma) = \frac{\sum_{\tau} h_p(\tau) \cdot \tilde{h}_p(\tau + \gamma)}{\|h_p\| \cdot \|\tilde{h}_p\|}, \qquad (5)$$
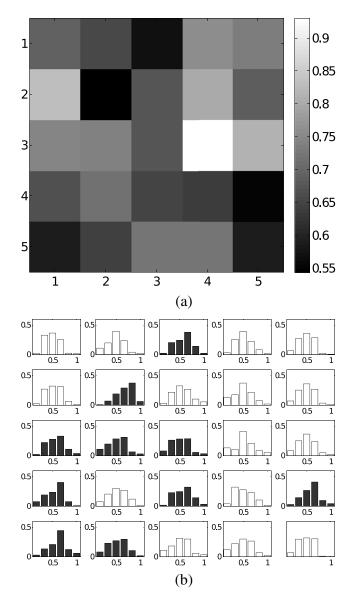
where $\| \cdot \|$ is the ℓ_2-norm. Larger similarity values reflect more accurate PSF estimation, thus better image restoration. The graphical representation of the similarity measure S is shown in Fig. 4(a). The darkest squares correspond to PSFs with the lowest similarity score.

The way we determine a nonvalid PSF is by characterizing the energy distribution along the local PSF space. We add the PSF values along concentric squares of radius r to build an energy distribution histogram $f(r)$, which is normalized to sum to 1. In Fig. 4(b), we show the histograms for the energy distribution characterization of the estimated PSFs from the artificially degraded retinal image [Fig. 3(a)]. To identify nonvalid PSFs, i.e., PSFs that do not follow the pattern we have previously described, we compute a shape descriptor for $f(r)$ defined as the probability $P(r \leq r_m)$,
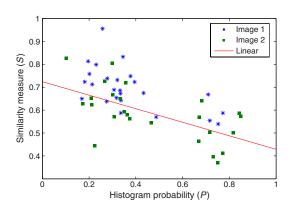
$$P(r \leq r_m) = \sum_{0}^{r_m} f(r) \cdot r, \qquad (6)$$

where r_m is the mode of $f(r)$. This descriptor gives an indication of how much a histogram is spread in relation to the peak (r_m) of the histogram. Particularly, it yields low values when the mode is located in the first few bins of the histogram and large values otherwise. This descriptor can be correlated with the PSF similarity measure to determine how it discriminates valid and nonvalid PSFs.

In Fig. 5, we plot the similarity measure S against $P(r \leq r_m)$ for the estimated PSFs from two different retinal images which have been artificially degraded. The correlation coefficient for the two variables is -0.64. In addition to the correlation, from the plot we note that most PSFs with high $P(r \leq r_m)$ have low kernel similarity (S) values, which is an indication that these are nonvalid or poorly estimated PSFs. A machine learning algorithm or clustering technique that automatically classifies nonvalid PSFs is out of the scope of this paper. Instead, our aim is to show that the energy characterization approach is sufficient to identify nonvalid PSFs, even at the expense of a few valid ones. As we show in Sec. 4.1, this is

(a)

(b)

Fig. 4 (a) PSF similarity measure. (b) Characterization of estimated local PSFs by energy distribution. Histograms plotted in white bars in (b) have been labeled as valid PSFs based on the histogram probability descriptor (P) described in the text.

Fig. 5 PSF similarity measure (S) versus histogram probability (P). An S value of 1 represents an accurately estimated kernel. Low P values are associated with correctly estimated kernels.

not critical because the PSF changes smoothly throughout the FOV.

The histograms plotted in white bars in Fig. 4(b) correspond to PSFs with $P(r \leq r_{\mathrm{m}}) < 0.3$, which we have labeled as valid PSFs. This means that we are favoring histograms skewed toward the left side. This is based on our assumption of the pattern that valid PSFs should follow. This correlates well with the similarity measure. Note that most PSFs with low similarity measure [darkest squares in Fig. 4(a)] have been correctly identified as nonvalid PSFs [black histograms in Fig. 4(b) and PSFs labeled with boxes in Fig. 6(a)].

The procedure for correcting the nonvalid local PSFs consists of replacing them with the average of adjacent valid kernels. Without this correction, the reconstruction develops ringing artifacts [see for example Fig. 11(b)]. The new set of valid local PSFs after replacing the nonvalid ones for the artificially degraded image is shown in Fig. 6(b).

3.4 PSF Interpolation

The computation of the SV PSF h is carried out by interpolating the local PSFs estimated on the regular grid of positions. The PSF values at intermediate positions are computed by bilinear interpolation of four adjacent known PSFs,[22] as shown in Fig. 7. Indexing any four adjacent grid points as $p = 1, \ldots, 4$ (starting from the top-left corner and continuing clockwise), the SV PSF in the position between them is defined as

$$h(s, t; x, y) = \sum_{p=1}^{4} \alpha_p(x, y) h_p(s, t), \tag{7}$$

where α_p is the coefficients of bilinear interpolation. Let us denote x_1 and x_2 as minimum and maximum x-coordinates of the subwindow, respectively. Analogously, y_1 and y_2 are the y-coordinates. Using auxiliary quantities

$$t_x = \frac{x - x_1}{x_2 - x_1}, \qquad t_y = \frac{y - y_1}{y_2 - y_1}, \tag{8}$$

the bilinear coefficients are

$$\alpha_1 = (1 - t_y)(1 - t_x), \qquad \alpha_2 = (1 - t_y)t_x,$$
$$\alpha_3 = t_y(1 - t_x), \qquad \alpha_4 = t_y t_x. \tag{9}$$

In light of the definition for an SV PSF in Eq. (7), we can compute the convolution of Eq. (1) as a sum of four convolutions of the image weighted by coefficients $\alpha_p(x, y)$

$$[Hu](x, y) = \int u(s, t)h(s, t, x - s, y - t)\mathrm{d}s\mathrm{d}t, \tag{10}$$

$$= \int u(s, t) \sum_{p=1}^{4} \alpha_p(s, t)h_p(x - s, y - t)\mathrm{d}s\mathrm{d}t, \tag{11}$$

$$= \sum_{p=1}^{4} \int (\alpha_p(s, t)u(s, t))h_p(x - s, y - t)\mathrm{d}s\mathrm{d}t, \tag{12}$$

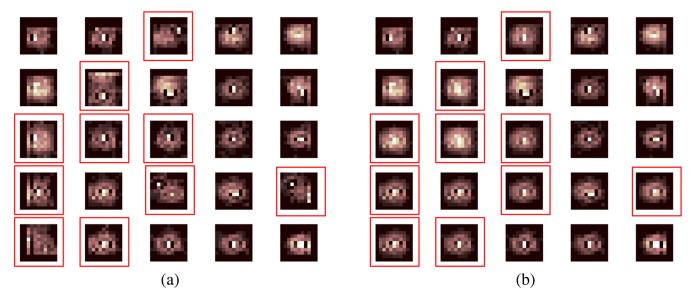$$= \left[\sum_{p=1}^{4} [\alpha_p u] * h_p\right](x, y). \tag{13}$$

In the same fashion, the operator adjoint to H (SV counterpart of correlation) denoted by H^* can also be defined in terms of the sums of four convolutions weighted by the α_p coefficients. These two operators are needed in all first-order minimization algorithms as the one used in the restoration stage (see Ref. 23 for further details).

3.5 Restoration

Having estimated a reliable SV PSF, we proceed to deblur the image. Image restoration is typically formulated within the Bayesian paradigm, in which the restored image is sought as the most probable solution to an optimization problem.

(a)　　　　　　　　　　(b)

Fig. 6 (a) Estimated set of 5×5 local PSFs. (b) New set of local PSFs with nonvalid PSFs replaced [compared with ground-truth set Fig. 3(b)]. Nonvalid PSFs have been labeled with a red square.

The restoration can be described as the minimization of the functional

$$\min_u \left[\frac{1}{2} \|z - Hu\|^2 + \lambda \int |\nabla u| \mathrm{dxdy} \right],$$
(14)

where z is the blurred observed image, H is the blurring operator [Eq. (1)], u is the unknown sharp image, and λ is a positive regularization constant, which we have set according to a fine-tuning procedure.[4] The tuning procedure consists in artificially degrading a retinal image and restoring it with Eq. (14) by varying λ. Because the sharp original image is known we can compare it against the restored image using a metric like the peak-signal-to-noise ratio to determine an optimal value of λ. The first term penalizes the discrepancy between the model and the observed image. The second term is the regularization term which serves as a statistical prior. As regularization we use total variation, a technique that exploits the sparsity of image gradients in natural images. At present, solving the convex functional of Eq. (14) is considered a standard way to achieve close to state-of-the-art restoration quality without excessive time requirements.[24] We used an efficient method[25] to solve Eq. (14) iteratively as a sequence of quadratic functionals
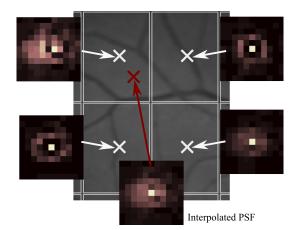
$$u_{i+1} = \arg \min_u \left[\frac{1}{2} \|z - Hu\|^2 + \lambda \int \frac{|\nabla u|^2}{2|\nabla u_i|} + \frac{|\nabla u_i|}{2} \mathrm{dxdy} \right].$$
(15)

The functional of Eq. (15) bounds the original function in Eq. (14) and has the same value and gradient in the current u_i, which guarantees convergence to the global minimum. To solve Eq. (15), we used the conjugate gradient method.[17] The initial value of u_i for $i = 0$ is set to be equal to z. In order to avoid numerical instability for areas with small gradient ($|\nabla u_i|$ approaching zero), we use a relaxed ϵ-form of the minimized functional in Eq. (14), which implies that $|\nabla u_i| < \epsilon$ is equal to ϵ.

As regards the restoration of color RGB retinal images, we consider the following. The most suitable channel for PSF estimation is the green because it provides the best contrast.[26] This is mainly due to the spectral absorption of the blood in this band, which yields the dark and well contrasted blood vessels.[27] Conversely, the blue channel encompasses the wavelengths most scattered and absorbed by the optical media of the eye,[28] therefore the image in this band has very low energy and a relatively high level of noise. In the spectral zone of wavelengths larger than 590 nm, the light scattering on the red blood cells and the light reflection from the eye structures behind the vessel are dominant.[29] This produces the red band to be saturated and of poor contrast. As a result, we estimate the SV PSF from the green channel of the RGB color image, and later deconvolve every R, G, and B channels with the estimated SV PSF to obtain a restored RGB color image.

4 Experiments and Results

We performed several experiments on artificially and naturally degraded images to illustrate the appropriateness of the SV approach for restoring blurred retinal images. Moreover, to achieve an artifact-free restoration, we used our strategy for detecting and replacing the nonvalid local PSFs.
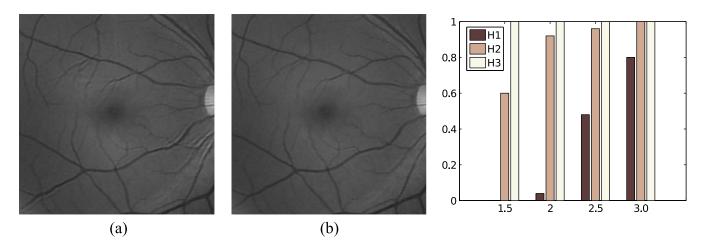

Fig. 7 Because the blur changes gradually, we can estimate convolution kernels on a grid of positions and approximate the PSF in the rest of the image (bottom kernel) by interpolation from four adjacent kernels.

4.1 Artificially Degraded Images

For the artificial experiment, we take a pair of images and degrade them with a 5×5 grid of realistic PSFs plus Gaussian noise ($\sigma = 10^{-6}$). The grid of PSFs was built upon realistic PSFs estimated from real degraded retinal images following the approach of Ref. 4. From the two input images, we restore one, and the other is used exclusively for the purpose of PSF estimation. We estimate the local PSFs by dividing the image into overlapping patches on a 5×5 grid [as shown in Fig. 3(a)]. The estimated PSFs are shown in Fig. 6(a). Because the PSF estimation may fail, we identify the nonvalid PSFs as described in Sec. 3.3. We replace them with with the average of adjacent valid kernels.

In Fig. 8(a), we show the restored artificial image with the directly estimated PSFs. The effect of nonvalid PSFs is evident in the poor quality of the restoration and the ringing artifacts. In Fig. 8(b), we show the restoration with the proposed method, where the nonvalid PSFs have been identified, removed, and replaced by the average of adjacent PSFs. To evaluate the restoration, we use the cumulative error histogram on a patch basis. The error[5] is the difference between a recovered image I_r with the estimated kernels and the known ground-truth sharp image I_g over the difference between the deblurred image I_{kg} with the ground-truth kernels. The error is given by $\|I_r - I_g\|/\|I_{kg} - I_g\|$. In Fig. 8(c), we show the cumulative error histogram for three restorations. H1 is the restoration with the directly estimated PSFs. It is important to note that shifted local PSFs warp the image which introduce additional artifacts and is the reason for such a low performance with approximately 40% of patches with an error lower than 2.5. After shifting the centroid of the PSFs to the geometrical center [restoration H2 in Fig. 8(c)], the reconstruction error is reduced significantly, about 60% of patches have an error lower than 1.5. Finally, the restoration (H3) with the removal of nonvalid PSFs increases significantly with all patches now displaying an error lower than 1.5. This means that after the nonvalid PSFs have been replaced the restoration quality is significantly increased.

To determine the limitations and robustness of the proposed method, we carried out several tests. First, we need to determine an optimal patch size for accurately estimating the local PSF. A patch that is too small compared to the kernel size may not have

Fig. 8 (a) Restoration with the set of directly estimated PSFs shown in Fig. 6(a) (notice the artifacts due to nonvalid PSFs) and (b) restoration with the new set of PSFs with nonvalid PSFs replaced shown in Fig. 6. (c) Error histogram for evaluating the reconstruction using: H1—directly estimated PSFs, H2—PSFs shifted toward the geometrical center, and H3—the new set of valid PSFs.

enough information and is likely to favor the trivial solution (convolution with a delta). Conversely, a patch that is bigger than necessary may hinder the SV approach in addition to increasing the computational burden. In Fig. 9(a), we show the mean similarity measure versus the image patch size. It is important to note that a patch size of roughly four times the kernel size may be sufficient for an accurate PSF estimation. Increasing too much the patch size hinders the PSF estimation.

The second aspect to consider is under- and over-estimation of the local PSF size. In Fig. 9(b), we show the mean patch reconstruction error versus PSF size. As expected the error is minimum when the size is similar to the true size, yet under-estimating the PSF size is worse than over-estimating.

In relation to the SV characterization of the blur, we performed the PSF estimation with a different number of kernels. Initially with a coarse 2×2 grid increasing up to a 10×10 fine grid (the 5×5 grid is the ground-truth). A similar behavior is observed. Under-estimating the proper PSF grid size has a negative effect in that the SV nature of the blur is hardly identified which yields an error above 6 for the whole image. In this case, the error is not computed on a patch-basis because of the variable grid size.

4.2 Naturally Degraded Images

All of the naturally degraded images used in the experiments were acquired in pairs, typically with a time span between acquisitions of several minutes. Initially, to show the limits of the space-invariant approach we restored the blurred retinal image from Fig. 10(a) with a single global PSF with the space-invariant method we proposed in Ref. 4. This image corresponds to the eye fundus of a patient with strong astigmatism, which induces an SV blur as depicted by the image details shown in Figs. 10(b)–10(e). The restoration is shown in Fig. 11(a) and we can clearly observe various artifacts despite an increase in sharpness in a small number of areas. In view of this, it is evident that the space-invariant assumption does not hold in such cases. In the following, we move to the SV approach.

To carry out the SV restoration, we estimated the local PSFs on a 5×5 grid of image patches. From the estimated PSFs shown in Fig. 12(a), we notice a clear variation in shape mainly from the top-right corner where they are quite narrow, to the bottom left corner where they are more spread and wide. This variation is consistent with the spatial variation of the blur observed in the retinal image of Fig. 10(a). We restored the
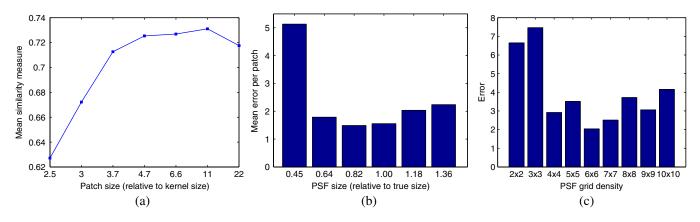

Fig. 9 (a) Mean similarity measure versus patch size. (b) Mean reconstruction error versus local PSF size. (c) Reconstruction error versus PSF grid size.
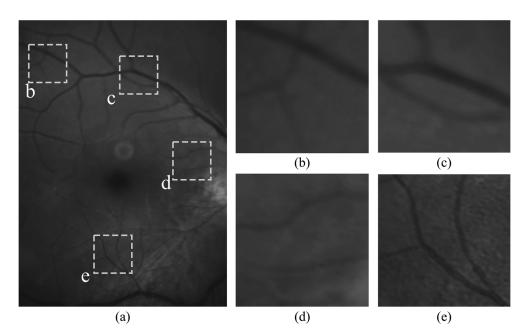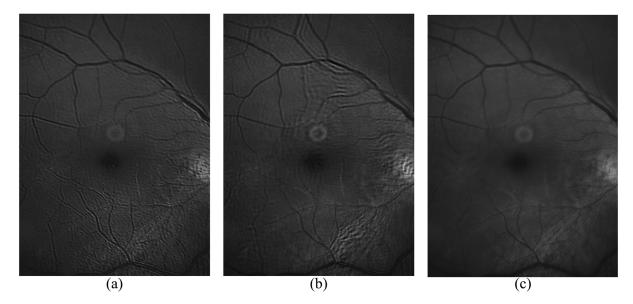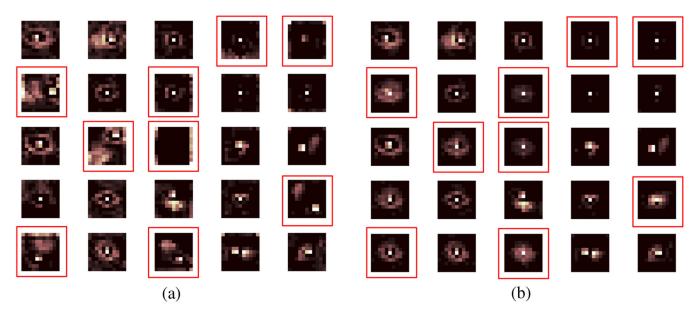
Fig. 10 (a) Retinal image degraded with real SV blur given by strong astigmatism. (b), (c), (d), and (e) zoomed regions to show the SV nature of the blur.

image with these local PSFs that were estimated directly without any adjustment. The restored image is shown in Fig. 11(b). One immediately obvious feature is that in several areas the restoration is rather poor, displaying ringing artifacts, whereas in others it is to some extent satisfactory. The local poor-restoration is linked to areas where the PSF estimation failed. By removing and correcting these nonvalid local PSFs, we obtained a noteworthy restoration shown in Fig. 11(c). Notice the overall improvement in sharpness and resolution with small blood vessels properly defined as shown by the image-details in the third column of Fig. 13. It could be said that without the replacement of the nonvalid PSFs the image quality after restoration is certainly worse than the original degraded image (see second column of Fig. 13).

To further demonstrate the capabilities of our method, additional restoration results on real cases from the clinical practice are shown in the following figures. As we mentioned in Sec. 1, a typical source of retinal image degradation comes from patients with corneal defects in which the cornea has an irregular structure [Fig. 1(a)]. This induces optical aberrations, which are mainly responsible for the SV blur observed in the retinal image. The image details shown in Fig. 1(b) reveal a significant improvement in which the retinal structures are much sharper and enhanced. In Fig. 14(a), a full color retinal image is shown, in which three small hemorrhages are more easily discernible in the restored image, along with small blood vessels. Another retinal image, shown in Fig. 14(b), reveals a clear improvement in resolution with a much finer definition of blood vessels.

Fig. 11 (a) Space-invariant restoration, (b) SV restoration with directly estimated PSFs, and (c) SV restoration with the new set of PSFs. The reader is strongly encouraged to view these images in full resolution at http://www.goapi.upc.edu/usr/andre/sv-restoration/index.html.

Fig. 12 (a) Estimated set of local PSFs from naturally degraded retinal image shown in Fig. 10(a). (b) Set of local PSFs after replacing nonvalid PSFs. Nonvalid PSFs have been labeled with a red square.

In addition, we processed retinal angiography images to test our method against a different imaging modality. Ocular angiography is a diagnostic test that documents, by means of photographs, the dynamic flow of dye in the blood vessels of the eye.[30] The ophthalmologists use these photographs both for diagnosis and as a guide to patient treatment. Ocular angiography differs from fundus photography in that it requires an exciter–barrier filter set (for further details see Ref. 30). The retinal angiography shown in Fig. 15 is degraded with a mild SV blur that hinders the resolution of small—yet important—details. The restoration serves to overcome this impediment; this can be observed from the zoomed-detail of the restored image.

The image enhancement may be useful for the improvement of recent analysis techniques for automated flow dynamics and identification of clinical relevant anatomy in angiographies.[31]

Finally, another way to demonstrate the added value of deblurring the retinal images is to extract important features, in this case detection of blood vessels. Such a procedure is commonly used in many automated disease detection algorithms. The improvement in resolution paves the way for a better segmentation of structures with edges. This is in great part due to the effect of the total variation regularization because it preserves the edge information in the image. To carry out the detection of the retinal vasculature, we used Kirsch's method.[32] It is a matched filter algorithm that computes the gradient by convolution with the image and eight templates to account for all possible directions. This algorithm has been widely used for
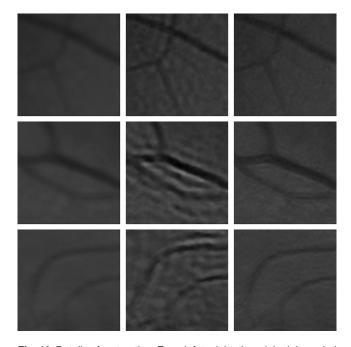

Fig. 13 Details of restoration. From left to right: the original degraded image, the SV restoration without correction of PSFs and the SV restoration with the correction.
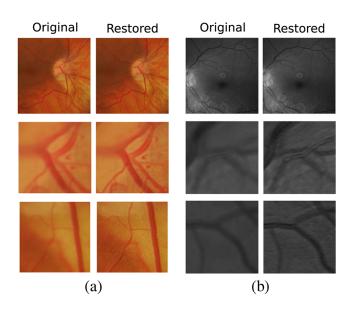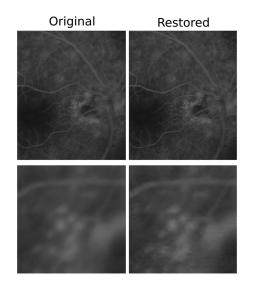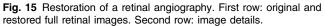

Fig. 14 Other retinal images restored with the proposed method. (a) First row: original and restored full-size retinal images. (b) Second and third rows: image details.
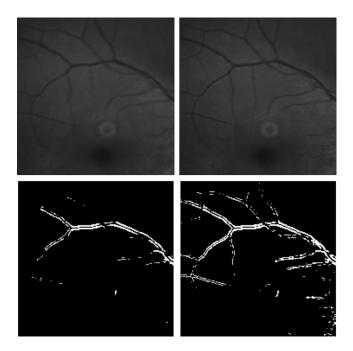
Original   Restored

Fig. 15 Restoration of a retinal angiography. First row: original and restored full retinal images. Second row: image details.

Fig. 16 First row (from left to right): original and restored retinal images. Second row: detection of blood vessels. Notice how the small blood vessels are better detected in the restored image.

detecting the blood vessels in retinal images.[33] In Fig. 16, we show the detection of the blood vessels from a real image of poor quality image and its restored version using our proposed method. A significant improvement in blood vessel detection is achieved. Smaller blood vessels are detected in the restored image, whereas the detection from the original image barely covers the main branch of the vasculature.

5 Conclusion

In this paper, we have introduced a method for restoring retinal images affected by SV blur by means of blind deconvolution. To do so, we described a spatially variant model of blur in terms of a convolution with a PSF that changes depending on its position. Since the SV degradation changes smoothly across the image,

we showed that the PSF need not be computed for all pixels, which is quite a demanding task, but for a small set of discrete positions. For any intermediate position bilinear interpolation suffices. In this way, we achieve an SV representation of the PSF.

The estimation of accurate local PSFs proved difficult due to the very nature of the images; they usually contain textureless or nearly homogenous regions that lack retinal structures, such as blood vessels, to provide sufficient information. In this regard, we proposed a strategy based on eye-domain knowledge to adequately identify and correct such nonvalid PSFs. Without this, the restoration results are artifact-prone with an overall image quality that is worse than the original image. The proposal has been tested on artificially and naturally degraded retinal images coming from the clinical practice. The details from the restored retinal images show an important enhancement, which is also demonstrated with the improvement in the detection of the retinal vasculature.

In summary, it seems clear that the SV restoration of blurred retinal images is significant enough to leverage the images' clinical use. Improving the visibility of subtle details like small hemorrhages or small blood vessels may prove useful for disease screening purposes, follow-up monitoring, or early disease detection. With the new challenges faced by clinical services in the 21st century, automated medical image analysis tools are mandatory—this work is a step toward that direction.

References

1. H. Bartling, P. Wanger, and L. Martin, "Automated quality evaluation of digital fundus photographs," *Acta Ophthalmol.* **87**(6), 643–647 (2009).
2. P. Godara et al., "Adaptive optics retinal imaging: emerging clinical applications," *Optom. Vision Sci.* **87**(12), 930–941 (2010).
3. J. Arines and E. Acosta, "Low-cost adaptive astigmatism compensator for improvement of eye fundus camera," *Opt. Lett.* **36**(21), 4164–4166 (2011).
4. A. G. Marrugo et al., "Retinal image restoration by means of blind deconvolution," *J. Biomed. Opt.* **16**(11), 116016 (2011).
5. A. Levin et al., "Understanding blind deconvolution algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(12), 2354–2367 (2011).
6. P. Bedggood et al., "Characteristics of the human isoplanatic patch and implications for adaptive optics retinal imaging," *J. Biomed. Opt.* **13**(2), 024008 (2008).
7. R. Tutt et al., "Optical and visual impact of tear break-up in human eyes," *Invest. Ophthalmol. Visual Sci.* **41**(13), 4117–4123 (2000).
8. J. Xu et al., "Dynamic changes in ocular zernike aberrations and tear menisci measured with a wavefront sensor and an anterior segment OCT," *Invest. Ophthalmol. Vis. Sci.* **52**(8), 6050–6056 (2011).
9. T. Costello and W. Mikhael, "Efficient restoration of space-variant blurs from physical optics by sectioning with modified Wiener filtering," *Digital Signal Process.* **13**(1), 1–22 (2003).
10. J. Bardsley et al., "A computational method for the restoration of images with an unknown, spatially-varying blur," *Opt. Express* **14**(5), 1767–1782 (2006).

11. S. Harmeling, M. Hirsch, and B. Scholkopf, "Space-variant single-image blind deconvolution for removing camera shake," *Adv. Neural Inf. Process. Syst.* **23**, 829–837 (2010).

12. O. Whyte et al., "Non-uniform deblurring for shaken images," *Int. J. Comput. Vis.* **98**(2), 168–186 (2012).

13. A. Gupta et al., "Single image deblurring using motion density functions," in *European Conf. on Computer Vision (ECCV 2010)*, K. Daniilidis, P. Maragos, and N. Paragios, Eds., pp. 171–184, Springer-Verlag, Crete, Greece (2010).

14. M. Tallón et al., "Space-variant blur deconvolution and denoising in the dual exposure problem," *Inf. Fusion* **14**(4), 396–409 (2012).

15. N. Salem and A. Nandi, "Novel and adaptive contribution of the red channel in pre-processing of colour fundus images," *J. Franklin Inst.* **344**(3–4), 243–256 (2007).

16. A. G. Marrugo and M. S. Millán, "Retinal image analysis: preprocessing and feature extraction," *J. Phys.: Conf. Ser.* **274**(1), 012039 (2011).

17. G. Golub and C. Van Loan, *Matrix Computations*, Vol. 3, Johns Hopkins University Press, Baltimore, Maryland (1996).

18. R. Navarro, E. Moreno, and C. Dorronsoro, "Monochromatic aberrations and point-spread functions of the human eye across the visual field," *J. Opt. Soc. Am. A* **15**(9), 2522–2529 (1998).

19. F. Sroubek and J. Flusser, "Multichannel blind deconvolution of spatially misaligned images," *IEEE Trans. Image Process.* **14**(7), 874–883 (2005).

20. R. Navarro, "The optical design of the human eye: a critical review," *J. Optom.* **2**(1), 3–18 (2009).

21. Z. Hu and M.-H. Yang, "Good regions to deblur," in *European Conf. on Computer Vision (ECCV 2012)*, pp. 59–72, Springer-Verlag (2012).

22. J. G. Nagy and D. P. O'Leary, "Restoring images degraded by spatially variant blur," *SIAM J. Sci. Comput.* **19**(4), 1063–1082 (1998).

23. M. Sorel, F. Sroubek, and J. Flusser, "Towards super-resolution in the presence of spatially varying blur," in *Super-Resolution Imaging*, P. Milanfar, Ed., pp. 187–218, CRC Press, Boca Raton, Florida (2010).

24. P. Campisi and K. Egiazarian, *Blind Image Deconvolution: Theory and Applications*, CRC Press, Boca Raton, Florida (2007).

25. A. Chambolle and P. L. Lions, "Image recovery via total variation minimization and related problems," *Numer. Math.* **76**(2), 167–188 (1997).

26. A. Hoover and M. Goldbaum, "Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels," *IEEE Trans. Med. Imaging* **22**(8), 951–958 (2003).

27. L. Gao, R. T. Smith, and T. S. Tkaczyk, "Snapshot hyperspectral retinal camera with the Image Mapping Spectrometer (IMS)," *Biomed. Opt. Express* **3**(1), 48–54 (2012).

28. M. Hammer and D. Schweitzer, "Quantitative reflection spectroscopy at the human ocular fundus," *Phys. Med. Biol.* **47**(2), 179–191 (2002).

29. V. Vucea et al., "Blood oxygenation measurements by multichannel reflectometry on the venous and arterial structures of the retina," *Appl. Opt.* **50**(26), 5185–5191 (2011).

30. P. Saine and M. Tyler, *Ophthalmic Photography: Retinal Photography, Angiography, and Electronic Imaging*, Butterworth-Heinemann, Woburn, Massachusetts (2002).

31. T. Holmes et al., "Dynamic indocyanine green angiography measurements," *J. Biomed. Opt.* **17**(11), 116028 (2012).

32. R. A. Kirsch, "Computer determination of the constituent structure of biological images," *Comput. Biomed. Res.* **4**(3), 315–328 (1971).

33. M. Al-Rawi, M. Qutaishat, and M. Arrar, "An improved matched filter for blood vessel detection of digital retinal images," *Comput. Biol. Med.* **37**(2), 262–267 (2007).

Andrés G. Marrugo received the BE degree (summa cum laude) in mechatronical engineering from the Technological University of Bolívar, Cartagena, Colombia, in 2008, the MSc in photonics, and the PhD in optical engineering from the Technical University of Catalonia, Barcelona, Spain, in 2009 and 2013, respectively. He is currently with the Technological University of Bolívar. He is a member of SPIE, the Spanish Optical Society, and the European Optical Society.

María S. Millán received the PhD in physics in 1990 and is a full professor of the College of Optics and Optometry in the Technical University of Catalonia, Barcelona, Spain. Her research work on image processing involves optical and digital technologies, algorithms, and development of new applications to industry and medicine. She has been a representative of the Spanish Territorial Committee in the International Commission for Optics. She is a member of the Optical Society of America, a fellow member of SPIE and a fellow member of the European Optical Society.

Michal Šorel received the MSc and PhD degrees in computer science from Charles University in Prague, Czech Republic, in 1999 and 2007, respectively. From 2012 to 2013 he was a postdoctoral researcher at Heriot-Watt University in Edinburgh, Scotland, and the University of Bern, Switzerland. Currently he is a research fellow in the Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic.

Filip Šroubek received the MS degree in computer science from the Czech Technical University, Prague, Czech Republic, in 1998 and the PhD degree in computer science from Charles University, Prague, Czech Republic, in 2003. From 2004 to 2006, he was on a postdoctoral position in the Instituto de Optica, CSIC, Madrid, Spain. In 2010 and 2011, he was the Fulbright visiting scholar at the University of California, Santa Cruz. He is currently with the Institute of Information Theory and Automation of the Academy of Sciences of the Czech Republic.