Unsupervised 3D deconvolution method for retinal imaging: principle and preliminary validation on experimental data.

G. Chenegros [a], L. M. Mugnier [a], C. Alhenc-Gelas [a], F. Lacombe [b], M. Glanc [c] and M. Nicolas [a,c]

[a]Office National d'Études et de Recherches Aérospatiales (ONERA), Optics Department, BP 72, F-92322 Châtillon cedex, France

[b]Mauna Kea Technologies, 9 rue d'Enghien 75010 Paris, France

[c]Observatoire de Paris-Meudon, Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique (LESIA), 5 place Jules Janssen, 92195 Meudon cedex, France

ABSTRACT

High resolution wide-field imaging of the human retina calls for a 3D deconvolution. In this communication, we report on a regularized 3D deconvolution method, developed in a Bayesian framework in view of retinal imaging, which is fully unsupervised, *i.e.*, in which all the usual tuning parameters, a.k.a. "hyper-parameters", are estimated from the data. The hyper-parameters are the noise level and all the parameters of a suitably chosen model for the object's power spectral density (PSD). They are estimated by a maximum likelihood (ML) method prior to the deconvolution itself.

This 3D deconvolution method takes into account the 3D nature of the imaging process, can take into account the non-homogeneous noise variance due to the mixture of photon and detector noises, and can enforce a positivity constraint on the recovered object. The performance of the ML hyper-parameter estimation and of the deconvolution are illustrated both on simulated 3D retinal images and on non-biological 3D experimental data.

Keywords: 3D deconvolution, retinal imaging, three-dimensional microscopy, hyper-parameter estimation, unsupervised estimation, regularization, inverse problems.

1. INTRODUCTION

Early detection of pathologies of the human retina call for an in vivo exploration of the retina at the cell scale. Direct observation from the outside suffers from the poor optical quality of the eye. The time-varying aberrations of the eye can be compensated *a posteriori* if measured simultaneously with the image acquisition; this technique is known as *deconvolution from wavefront sensing*[1,2] and has been successfully applied to the human retina.[3] These aberrations can also be compensated for in real-time, by use of adaptive optics (AO).[4] Yet, the correction is always partial.[5–7] Additionally, the object under examination (the retina) is three-dimensional (3D) and each recorded image contains contributions from the whole object's volume. For these two reasons, a deconvolution of the recorded images is necessary.

In two-dimensional (2D) deconvolution, each image is deconvolved separately, *i.e.*, only one object plane is assumed to contribute to each image. This is an appropriate image model in astronomy for instance, but is a crude approximation in microscopy, as it does not properly account for the halo in each image that comes from the parts of the observed object that are out-of-focus.

Three-dimensional deconvolution is an established technique in microscopy, and in particular in conventional fluorescence microscopy.[8] The combination of a conventional microscope with deconvolution is often referred to as deconvolution microscopy or even "digital confocal", because the use of 3D deconvolution can notably improve the resolution of the recorded conventional images, especially in the longitudinal (a.k.a. axial) dimension, while remaining simpler and cheaper than a confocal microscope. Yet, to the best of our knowledge, deconvolution of retinal images has so far been performed with 2D deconvolution techniques, both in deconvolution from wavefront sensing[3] and in deconvolution of AO-corrected images.[9]

Besides, because deconvolution is an ill-posed inverse problem,[10–12] most modern deconvolution methods use regularization in order to avoid an uncontrolled amplification of the noise.

Further author information: G.C. is now with Observatoire de Paris-Meudon /LESIA. E-mail: Guillaume.Chenegros@obspm.fr. L.M.M.'s E-mail is Laurent.Mugnier@onera.fr.

2. 3D DECONVOLUTION METHOD

The image formation is modeled as a 3D convolution:

$$\boldsymbol{i} = \boldsymbol{h} \bigstar \boldsymbol{o} + \boldsymbol{n} \tag{1}$$

where $\boldsymbol{i}$ is the (3D) pile of (2D) recorded images, $\boldsymbol{o}$ is the 3D unknown observed object, $\boldsymbol{h}$ is the 3D point spread function (PSF), $\boldsymbol{n}$ is the noise and $\bigstar$ denotes the 3D convolution operator.

For a system with N images of N object planes, this 3D convolution can be rewritten as:

$$\boldsymbol{i}_k = \left(\sum_{j=0}^{N-1} \boldsymbol{h}_{k-j} \star \boldsymbol{o}_j \right) + n_k \tag{2}$$

where $\boldsymbol{o}_j$ is the object in plane j, $\boldsymbol{i}_k$ is the k-th recorded image, $\boldsymbol{h}_{k-j}$ is the 2D PSF corresponding to a defocus of $(k-j)$ slices, and $\star$ denotes the 2D convolution operator. The PSF is that of the system composed of the eye, the imaging system (including the AO) and the detector. We assume that the whole recording process is fast enough so that the different 2D PSF's differ only by a defocus. Figure 1 illustrates the imaging process in the case of three object and image planes.

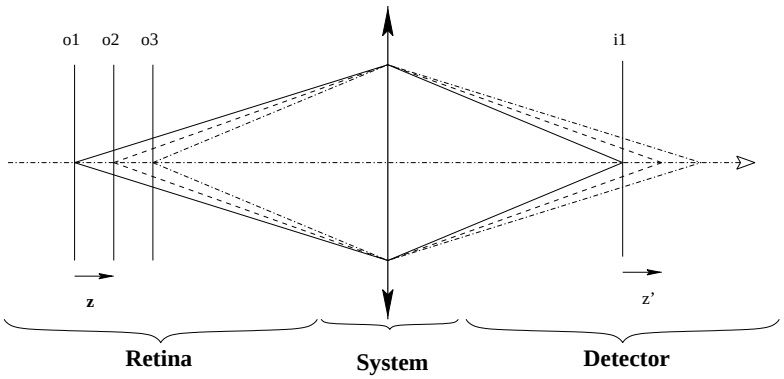

Figure 1. Illustration of 3D imaging of the retina for 3 object planes. The observed 3D object $\boldsymbol{o}$ is on the left and one plane, $\boldsymbol{i}_1$, of the 3D image $\boldsymbol{i}$ is represented on the right. All object planes contribute to each image plane.

Note that the raw image must be carefully pre-processed to yield an image that closely follows this imaging model. The pre-processing includes the relative registration of the raw images[13] and an image tapering in order to reduce the artefacts due to the periodization done by fast Fourier transforms. The goal is to obtain an estimate $\hat{\boldsymbol{o}}$ of the observed 3D object $\boldsymbol{o}$ given the images $\boldsymbol{i}$, the 3D PSF $\boldsymbol{h}$, and some prior information on the noise statistics and on the object.

2.1. Framework

Most deconvolution techniques boil down to the minimization (or maximization) of a criterion. An important task is the definition of a suitable criterion for the given inverse problem.

Following the Bayesian[12] Maximum *A Posteriori* (MAP) approach, the deconvolution problem can be stated as follows: we look for the most likely object $\hat{\boldsymbol{o}}$ given the observed image $\boldsymbol{i}$ and our prior information on $\boldsymbol{o}$, which is summarized by a probability density $p(\boldsymbol{o})$. This reads:

$$\hat{\boldsymbol{o}} = \arg\max_{\boldsymbol{o}} p(\boldsymbol{o}|\boldsymbol{i}) = \arg\max_{\boldsymbol{o}} p(\boldsymbol{i}|\boldsymbol{o}) \times p(\boldsymbol{o}). \tag{3}$$

Equivalently $\hat{\boldsymbol{o}}$ can be defined as the object that minimizes a compound criterion $J(\boldsymbol{o})$ defined as follows:

$$J(\boldsymbol{o}) = J_i(\boldsymbol{o}) + J_o(\boldsymbol{o}), \tag{4}$$

where the negative log-likelihood $J_i = -\ln p(\boldsymbol{i}|\boldsymbol{o})$ is a measure of fidelity to the data and $J_o = -\ln p(\boldsymbol{o})$ is a regularization or penalty term, so the MAP solution can equivalently be called a penalized-likelihood solution. Note that the Bayesian approach does not require that $\boldsymbol{o}$ truly be the outcome of a stochastic process; rather, $p(\boldsymbol{o})$ should be designed to embody the available prior information on $\boldsymbol{o}$, which means that J_o should have higher values for objects that are less compatible with our prior knowledge,[11] e.g. that are very oscillating.

When $\boldsymbol{o}$ is not the outcome of a stochastic process, J_o usually includes a scaling factor or global hyper-parameter, denoted by μ in the following, which adjusts the balance between fidelity to the data and fidelity to the prior information. Because the optimal balance between data and prior depends on the flux of each object plane, in the following we propose a regularization term with one such hyper-parameter *per* object plane. And we shall show in Section 3 how to estimate these hyper-parameters.

2.2. Object prior

The choice of a Gaussian prior probability distribution for the object can be justified from an information theory standpoint as being the least informative, given the first two moments of the distribution.

Additionally, we assume that the different object planes are independent because firstly, the retina is composed of potentially very different slices, and secondly we would like to improve the axial resolution in the restored object, which would be hampered by a regularization in the axial dimension. The regularization criterion is thus quadratic (or "$L2$" in short) and given by:

$$J_{\boldsymbol{o}}(\boldsymbol{o}) = \frac{1}{2}\sum_{l=0}^{N-1}\left(\sum_{f}\frac{\left|\tilde{\boldsymbol{o}}_l(f) - \overline{\tilde{\boldsymbol{o}}_l}(f)\right|^2}{S_{\boldsymbol{o}_l}(f)}\right) \tag{5}$$

where $S_{\boldsymbol{o}_l}$ is the 2D power spectral density (PSD) of object plane l.

A reasonable model of the PSD of each object plane can be found[14] and used with potentially different parameters for each plane in the above equation. It reads:

$$S_{\boldsymbol{o}_l}(f) = \frac{1}{\mu_l\left(1 + \left(\frac{|f|}{f_l^0}\right)^{p_l}\right)}, \tag{6}$$

where

- μ_l is inversely proportional to the energy of object plane l;
- p_l characterizes the spectral richness of object plane l – in practice, p goes from 0 for point like objects to 4 for very smooth extended objects;
- f_l^0 represents the inverse of the 2D object size for finite support objects. This parameter is very difficult to adjust by hand for objects extended beyond the field of view.

With such a fine PSD model per plane ($3 \times NP$ parameters for NP object planes), the MAP method can optimally adjust the balance between the resolution and the robustness to noise not only for each plane for *each spatial frequency in each plane*, provided the PSD model parameters are correctly tuned. Tuning all these hyper-parameters in a supervised fashion (*i.e.*, manually) is unrealistic, so we developed a method to estimate these hyper-parameters in an unsupervised way directly from the data, before performing the deconvolution. Before we derive this method in Section 3, in the next subsection we give some details about the noise model that is incorporated into our method.

2.3. Noise model

The noise is a mixture of photon noise and detector noise. Photon noise is non-stationary, white, Poisson-distributed and can be well modeled as non stationary white Gaussian as soon as the flux level is a few tens of photo-electrons per pixel.[15] Detector noise is, to a good approximation, stationary white Gaussian, so the sum of these two noises can reasonably be modeled as non stationary white Gaussian. The neg-log-likelihood of the data is thus the following weighted least-square difference between the actual data $\boldsymbol{i}$ and our model of the data for a given object, $\boldsymbol{h} \star \boldsymbol{o}$:

$$J_i(\boldsymbol{o}) = \frac{1}{2} \sum_{k=0}^{N-1} \sum_{p,q=0}^{N_{pix}-1} \left(\frac{1}{\sigma_k^2(p,q)} \left| \boldsymbol{i}_k(p,q) - \sum_{l=0}^{N-1} \left[\boldsymbol{h}_{k-l} \star \boldsymbol{o}_l \right](p,q) \right|^2 \right) \tag{7}$$

3. UNSUPERVISED HYPER-PARAMETER ESTIMATION

In this section, we estimate the parameters of the object 3D PSD by Maximum Likelihood (ML). The proposed procedure is an extension to 3D imaging of the method developed by Blanc *et al.* in the context of phase diversity wavefront-sensing[16] and applied by Gratadour *et al.* to the restoration of AO-corrected astronomical images.[17] With the parametric PSD model of Eq. (6), the 3D object PSD is determined by three parameters per object plane. In this subsection (and only here), we make the additional approximation that the noise is stationary, of variance σ_k^2, within each image plane k; this makes the object PSD estimation much more tractable, and is justified by the fact that the images considered here are illuminated rather uniformly due to all the out-of-focus object planes contributing to each image. Therefore, for each plane there are four parameters to be estimated, which will be jointly denoted by Θ_k: $\Theta_k = (\mu_k, f_k^0, p_k, \sigma_k^2)$.

If we denote by $\boldsymbol{H}_{k-j}$ the operator performing the 2D convolution with PSF $\boldsymbol{h}_{k-j}$, the convolution product from Eq. (2) can be rewritten as:

$$\boldsymbol{i}_k \triangleq \sum_{j=0}^{N-1} \boldsymbol{H}_{k-j} \boldsymbol{o}_j + \boldsymbol{n}_k \tag{8}$$

where $\boldsymbol{H}$ is the matrix of the convolution product. The mean of an image plane k is given by :

$$\overline{\boldsymbol{i}_k} = \sum_{j=0}^{N-1} \boldsymbol{h}_{k-j} \star \overline{\boldsymbol{o}_j}, \tag{9}$$

where $\overline{\boldsymbol{o}_j}$ is the mean of the object plane j. In practice we take $\overline{\boldsymbol{o}_j} = 0 \, \forall j$, so that $\overline{\boldsymbol{i}_k} = 0$ too. The covariance matrix of image plane k is given by:

$$C_{i_k} = \sum_{j=0}^{N-1} \boldsymbol{H}_{k-j} C_{o_j} \boldsymbol{H}_{k-j}^t + C_{n_k} \tag{10}$$

where C_{o_j} is the covariance matrix of the object plane j and C_{n_k} is the covariance matrix of image plane k.

Because $\boldsymbol{o}_j - \overline{\boldsymbol{o}_j}$ is assumed stationary, the object covariance matrix C_{o_j} is Toeplitz (more precisely, block-Toeplitz with Toeplitz blocks). Because the noise's covariance matrix is $C_{n_k} = \sigma_k^2$.I, the image covariance matrix C_{i_k} is Toeplitz too. Thus all matrices of Eq. (10) are approximately diagonalized by fast Fourier transforms (FFT), which yields the following expression for the PSD S_{i_k} of image $\boldsymbol{i}_k$:

$$S_{i_k}(f, \Theta_k) = \sum_{j=0}^{N-1} \left| \tilde{\boldsymbol{h}}_{k-j} \right|^2 S_{o_j}(f, \Theta_j) + S_{n_k}(f, \Theta_k) \tag{11}$$

Because only transfer functions $\tilde{\boldsymbol{h}}_{k-j}$ close to focus ($k - j \simeq 0$) carry high frequency information, the latter equation can be approximated by:

$$S_{i_k}(f, \Theta_k) \simeq \left(\sum_{j=0}^{N-1} \left| \tilde{\boldsymbol{h}}_{k-j} \right|^2 \right) S_{o_k}(f, \Theta_k) + S_{n_k}(f, \Theta_k) \tag{12}$$

except at very low frequencies.

With the above Gaussian model for the object prior probability distribution and for the noise distribution, Eq. (8) shows that the probability distribution of each image plane is Gaussian too. The likelihood of image $\boldsymbol{i}_k$ is given by:

$$p(\boldsymbol{i}_k|\Theta_k) = \frac{1}{(2\pi)^{N/2}\sqrt{\det(C_{i_k})}}.\exp\left[\frac{1}{2}(\boldsymbol{i}_k - \overline{\boldsymbol{i}_k})^t C_{i_k}^{-1}(\boldsymbol{i}_k - \overline{\boldsymbol{i}_k})\right] \tag{13}$$

The negative log-likelihood is thus given by :

$$L(\Theta_k) = -\ln(p(\boldsymbol{i}_k|\Theta_k) = \frac{1}{2}\sum_f \ln(S_{i_k}(f,\Theta_k)) + \frac{1}{2}\sum_f \frac{\left|\tilde{\boldsymbol{i}}_k(f) - \overline{\tilde{\boldsymbol{i}}}_k(f)\right|^2}{S_{i_k}(f,\Theta_k)} \tag{14}$$

where $\overline{\boldsymbol{i}}_k$ is given by Eq. (9) and $S_{i_k}(f,\Theta_k)$ by Eq. (12).

Finally, for each image plane we are able to determine the three object PSD parameters by minimizing $L(\Theta_k)$ numerically, as a function of $\Theta_k = (\mu_k, f_k^0, p_k, \sigma_k^2)$. These parameters can then be used in the deconvolution.

4. VALIDATION ON SIMULATED DATA

4.1. Simulated object and images

In order to validate our deconvolution method by simulations, we created a simulated object which complies with the overall structure of a retina. Figure 2 represents the original simulated object, composed of vessels, ganglion cells and photo-receptors. The vessels are simulated by moving a ring in a random walk, the ganglion cells are simulated by empty globes and photo-receptors are represented by two empty half spheres joined by an empty tube. The cube presented on Figure 2 is 300 µm x 300 µm x 300 µm.

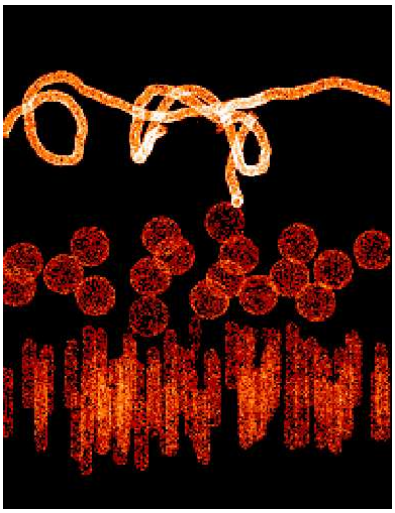

Figure 2. Perspective view of the 3D object used for the simulations.

In the simulations presented here, we use a five slice object obtained by averaging the data from Figure 2 into five 13 µm-thick slices from which we select a 128×128 region of interest. The five slices obtained are presented on Figure 3. The PSF's used to compute the 3D image $\boldsymbol{i}$ are currently purely diffractive. They are generated with a set of aberrations expanded on the Zernike basis; we use 0.2 rd RMS of one astigmatism (Z_5), -0.1 rd RMS of the other astigmatism (Z_6) and -0.5 rd RMS of spherical aberration (Z_{11}). These PSF's are oversampled, with respect to the Nyquist frequency, by a factor of 1.5. With the object and the PSF, we simulate the image by means of Equation (2). The noise added is white

Figure 3. The five object layers (black corresponds to 0 ph/pix).

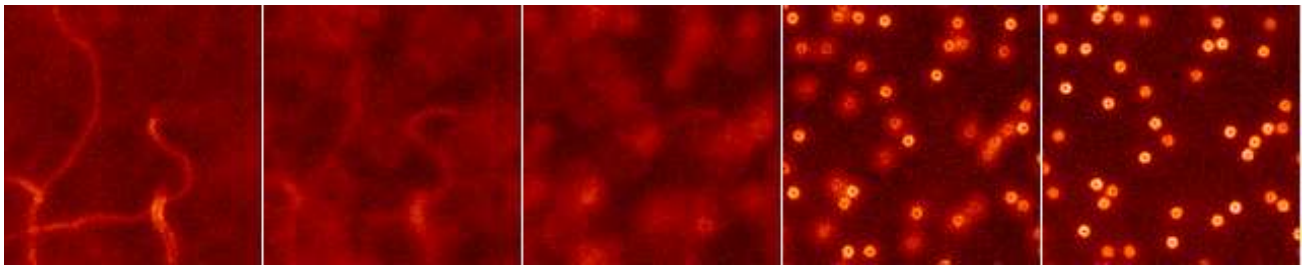

Figure 4. The five image layers.

Gaussian and stationary; its standard deviation is 3% of the maximum intensity in the object $\boldsymbol{o}$ (corresponding roughly to 1000 photo-electrons per pixel for photon-limited data). The five image layers are presented on Figure 4. From these images, it is clear that all object slices contribute to all images. With the relatively small chosen separation between planes ($13\,\mu$m), the two first images are visually very similar whereas the corresponding object slices are very different. The deconvolution aims at disentangling the contribution of the different object slices (*i.e.*, improving the axial resolution) and at improving the lateral resolution within each plane.

4.2. Unsupervised PSD estimation

We now want to validate the unsupervised PSD estimation method by simulations. A comparison between the "true" object PSD and the estimated object PSD is given on Figures 5 and 6. On Figure 5 one can see that the PSD estimated from the first image plane is quite close to the "true" object PSD, which is fitted from the first object plane. One notices an overestimation of the low frequencies by our method but the intermediate and high frequencies are correctly estimated. And these frequencies are the ones that are important for the object reconstruction: the regularization hardly matters for the very low frequencies.

Figure 6 illustrates how well the unsupervised estimation method is able to separate the noise PSD and the object PSD from the empirical image PSD. The quality of the object PSD estimation with our method is perfectly compatible with the 3D deconvolution method as we can see in the next section.

The quality of the μ parameter estimation is studied by plotting the restoration error of the 3D object versus the value of μ relative to the estimated value. Figure 7 and 8 show these plots for a 3D deconvolution without and with positivity constraint respectively. We can observe that without positivity constraint the optimal value of μ is the estimated value. With the positivity constraint, it is necessary to under-regularize the deconvolution because the positivity constraint is itself a regularization. In practice we can take for μ the estimated value divided by 10.

4.3. Restoration of simulated data

In this section, we present two results obtained with our 3D unsupervised deconvolution method. The first simulation, presented on Figure 9 and the second, presented on Figure 10, shows the deconvolution results obtained with $L2$ regularization without and with positivity constraint respectively. The object PSD is automatically estimated by our method. We

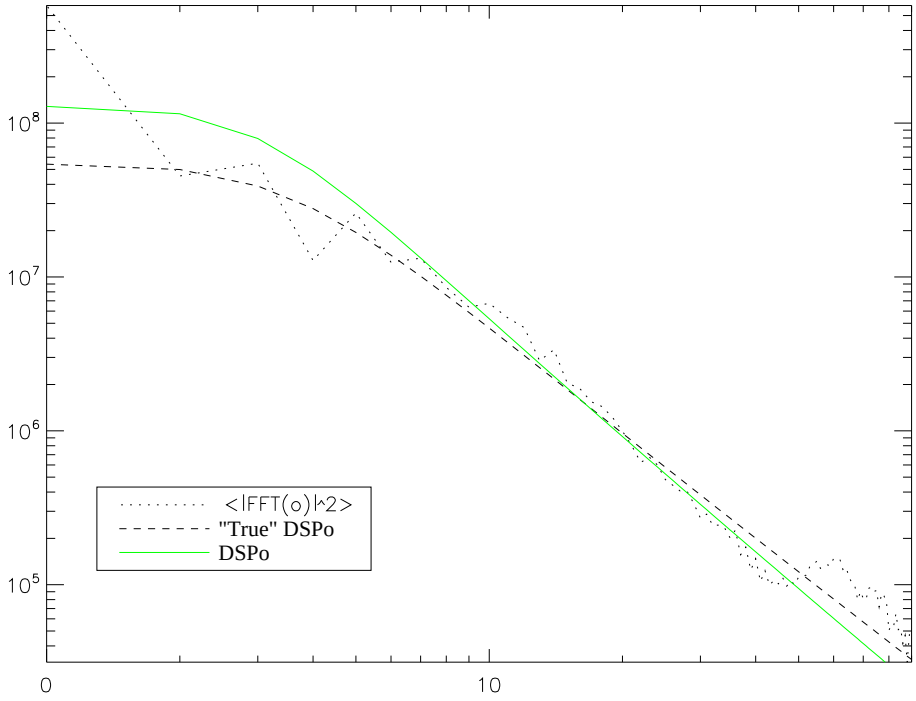

Figure 5. Assessment of the object's PSD estimation quality. The black dashed line plot corresponds to the true object's PSD fitted directly on the object, the black dotted line plot corresponds to the empirical true object's PSD and the green plot corresponds to the object's PSD estimated from the image with our method.

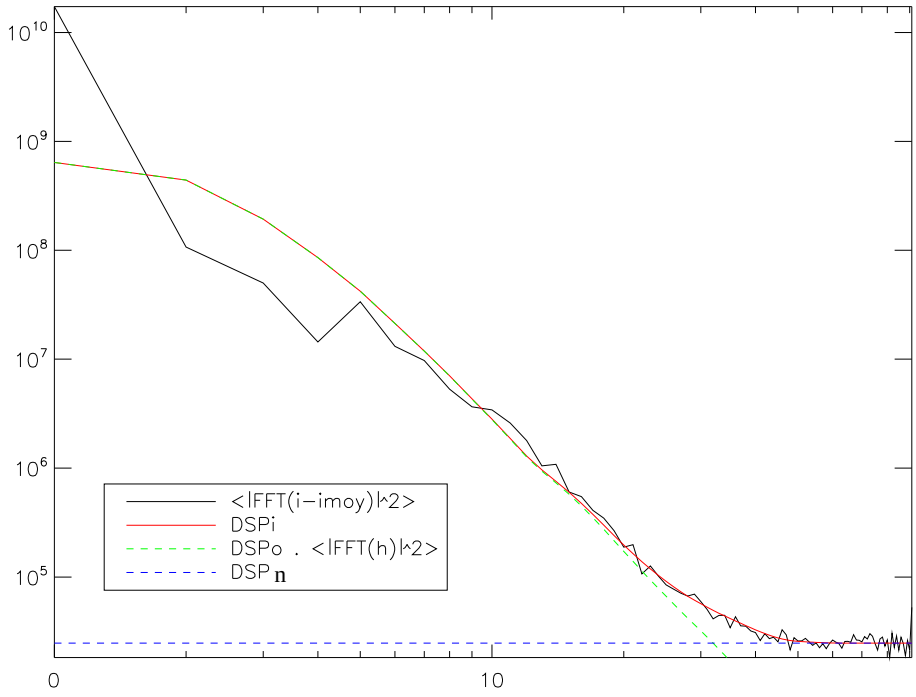

Figure 6. Empirical PSD of the first image plane (solid black line) and its fit by our parametric image PSD model (red line). The latter is the sum of an object contribution (green dashed line) and a noise contribution (blue dashed line), estimated jointly by our method.

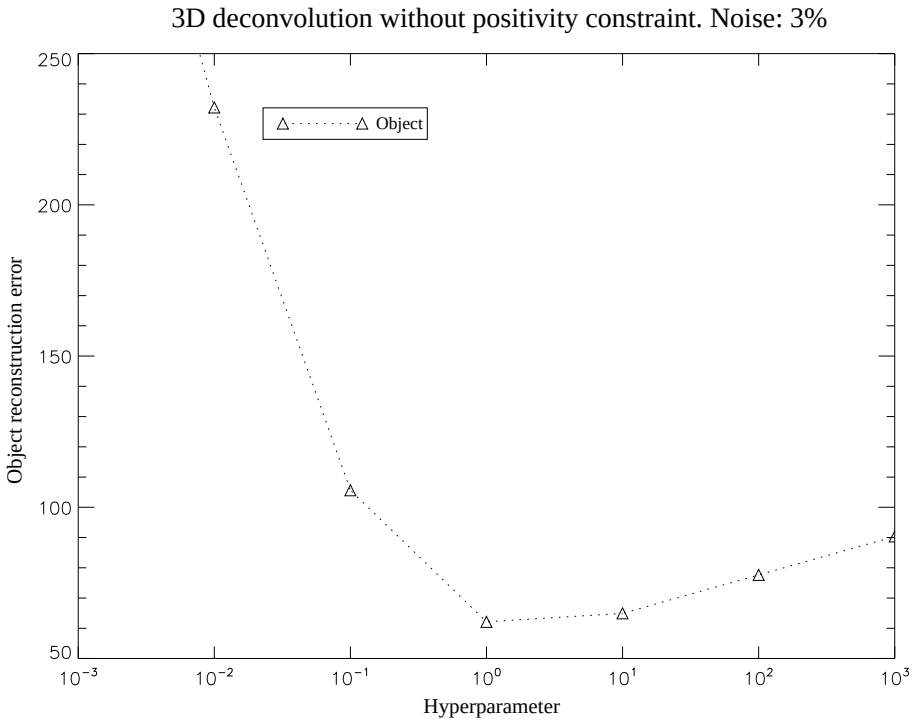

Figure 7. Plot of the 3D object restoration error without positivity constraint versus the relative μ value.

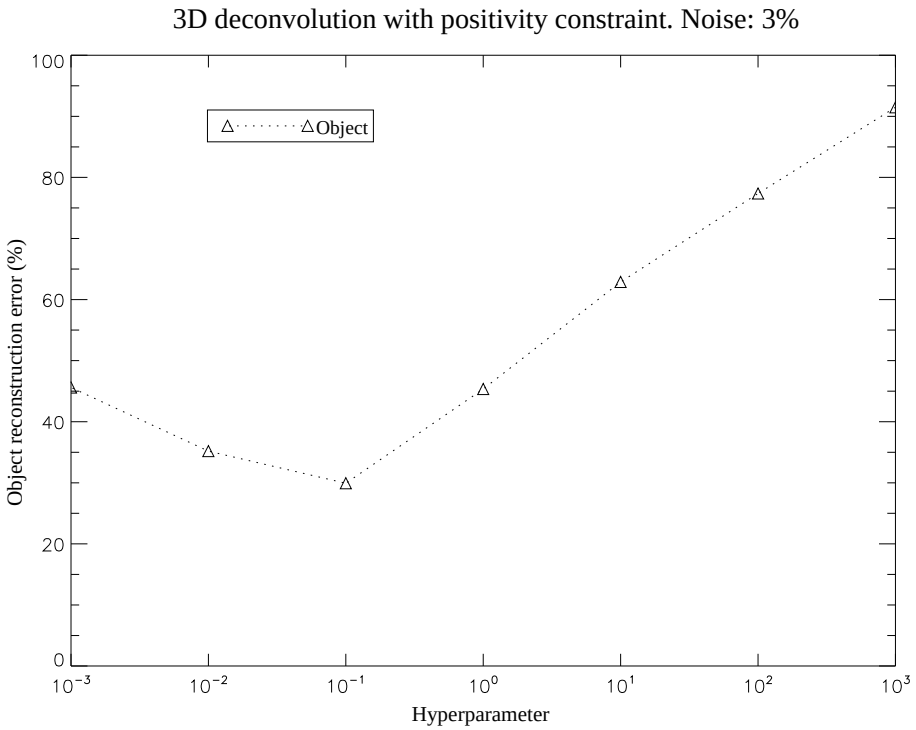

Figure 8. Plot of the 3D object restoration error with positivity constraint versus the relative μ value.

can see ghosts of vessels (in the middle plane for example) on Figure 9 and a residual blur: the missing cone of 3D frequencies makes it difficult for the restoration procedure to correctly disentangle the contribution of all planes. The positivity constraint used in Figure 10 helps the algorithm disentangle the different planes and visibly reduces ghosts of vessels in middle plane. More quantitatively, the RMS restoration error is 8.34 ph/pix with the positivity constraint and 10.31 ph/pix without (the object average level is 15.34 ph/pix).

Figure 9. The five estimated object layers with $L2$ regularization without positivity constraint.

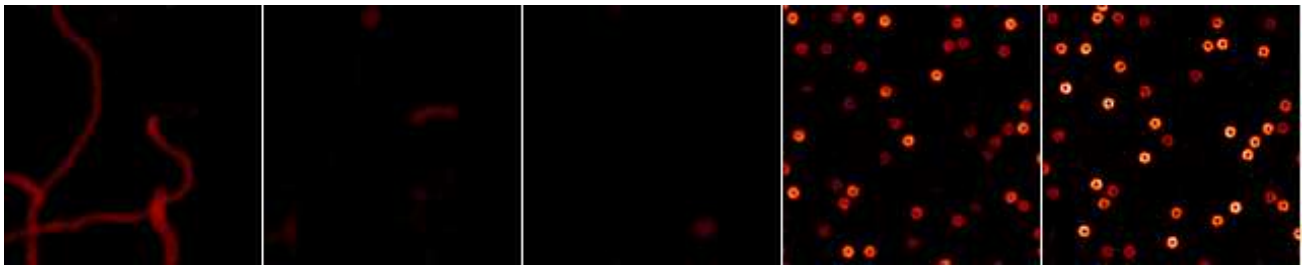

Figure 10. The five estimated object layers with $L2$ regularization under positivity constraint (black corresponds to 0 ph/pix).

5. VALIDATION ON EXPERIMENTAL DATA

5.1. Experimental bench

To validate our 3D deconvolution method, we developed a new 3D optical bench at ONERA. This bench works without adaptive optics to simplify the image acquisition and just focus in the deconvolution problems. The experimental setup is shown on Figure 11.

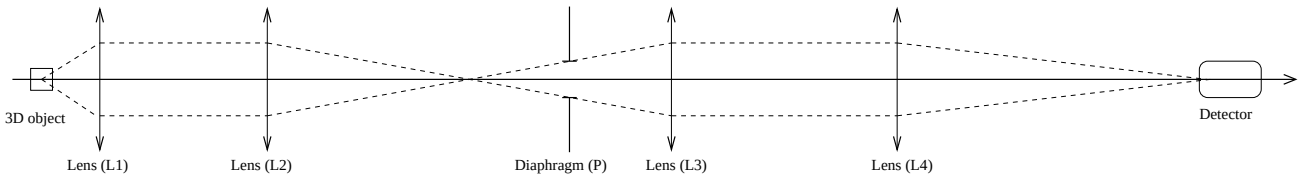

Figure 11. 3D imaging bench.

Some specifications of our bench are:

1. the images must be at least Shannon-sampled;
2. the detector must be translated in an automated way along the optical axis;

3. the aberrations must be the same in the field (no anisoplanatism).

The bench presented on Figure 11 takes into account all these constraints. L_1 is a microscope objective (16X, NA=0.32); L_2 is a biconvex lens ($f = 62.9$ mm and $d = 25.4$ mm); P is the pupil of the system; L_3 and L_4 perform a re-imaging of the object through the diaphragm P. We can use two objects with our bench: a pinhole (1 µm diameter) to measure the 3D PSF, and a micrometer rule tilted with respect to the optical axis (by 35 degrees) as a 3D object.

5.2. Restoration of experimental data

We present here the rule images and the 3D PSF. The 3D rule image has 30 different planes. The Figure 12 presents two planes, taken at two different depths (the distance between the two planes is 3.2 cm).

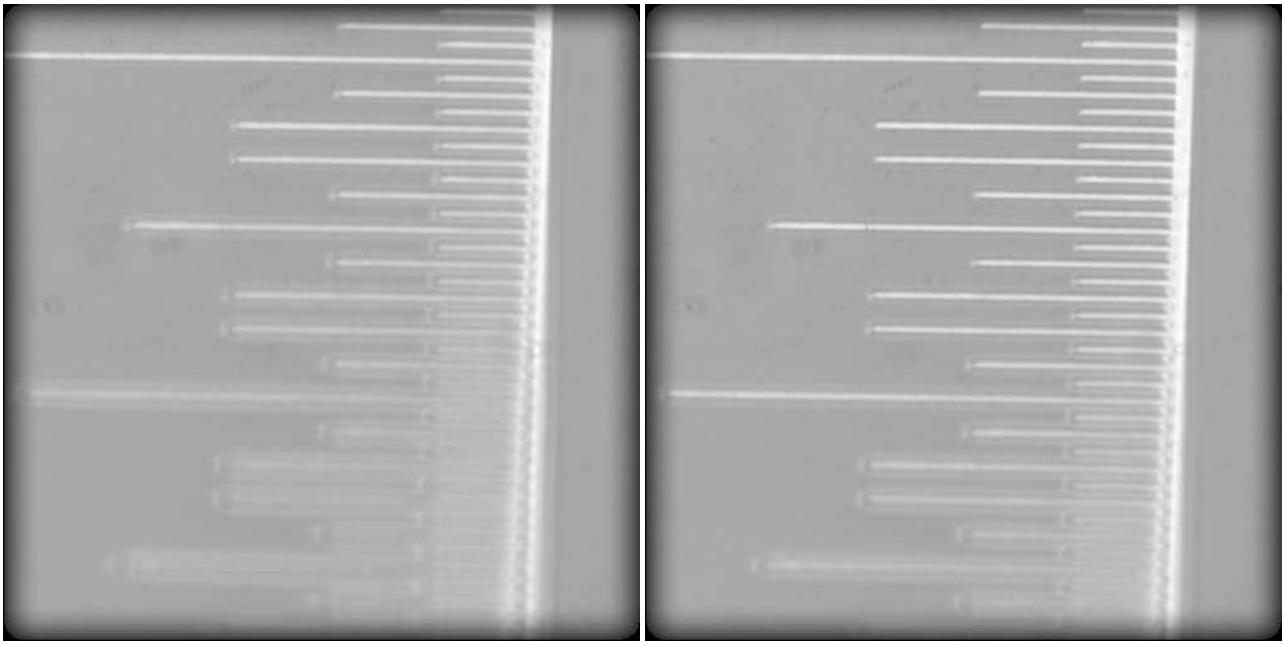

Figure 12. Two planes of the 3D rule image. On the left, the detector is focused at the very top of the rule. On the right, the detector is focused on about one quarter of the image from the top of the rule.

We performed an unsupervised 3D deconvolution of this 3D image of the rule with the PSD model introduced in section 3 and the positivity constraint. The object PSD is estimated from each image plane with the method presented in section 3. With the positivity constraint, the parameter μ is divided by 10 in practice as shown in simulations. One can notice on Figure 13 the lateral resolution increase due to the deconvolution method (the width of each line is 6 pixel in an image and only 1 pixel on the solution). The restored width is compatible with real dimensions of the rule because each line is about 1 µm wide. This restored object is very encouraging and demonstrates the effectiveness of our 3D unsupervised deconvolution method.

6. CONCLUSION AND PERSPECTIVES

A 3D deconvolution method has been derived in a Bayesian framework for the restoration of adaptive-optics corrected images of the human retina; it incorporates a positivity constraint and a regularization metric in order to avoid uncontrolled noise amplification. An unsupervised method has been proposed to estimate the 3D object PSD from the 3D image and the 3D PSF. We have demonstrated the effectiveness of the method, on realistic simulated data and on experimental data obtained the 3D optical bench developed at ONERA.

Future work includes the processing of adaptive-optics corrected retinal images. For this purpose, it is of paramount importance to estimate the PSF precisely in order not to produce deconvolution artefacts. This can be achieved by reconstructing the residual wavefront from the wavefront sensor data of the adaptive-optics loop.

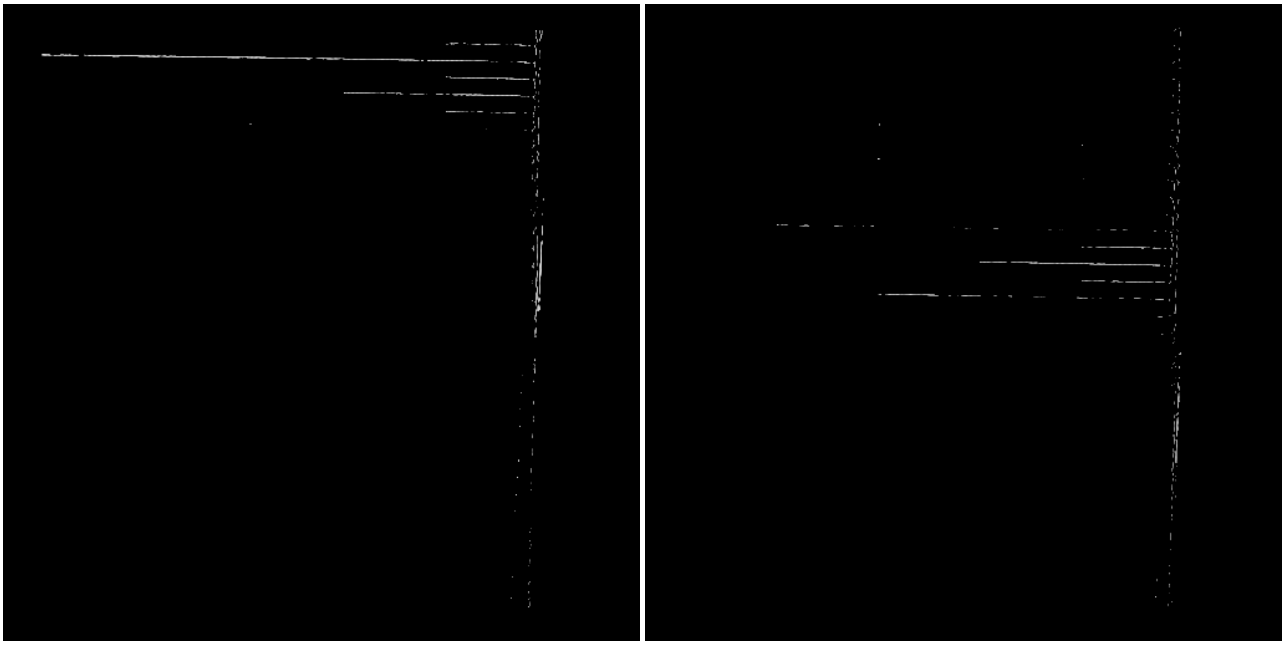

Figure 13. Two object planes restored with positivity constraint. On the left, the detector is focused at the very top of the rule. On the right, the detector is focused on about one quarter of the image from the top of the rule.

REFERENCES

1. J. Primot, G. Rousset, and J.-C. Fontanella, "Deconvolution from wavefront sensing: a new technique for compensating turbulence-degraded images," *J. Opt. Soc. Am. A* **7**(9), pp. 1598–1608, 1990.
2. L. M. Mugnier, C. Robert, J.-M. Conan, V. Michau, and S. Salem, "Myopic deconvolution from wavefront sensing," *J. Opt. Soc. Am. A* **18**, pp. 862–872, Apr. 2001.
3. D. Catlin and C. Dainty, "High-resolution imaging of the human retina with a Fourier deconvolution technique," *J. Opt. Soc. Am. A* **19**, pp. 1515–1523, Aug. 2002.
4. G. Rousset, J.-C. Fontanella, P. Kern, P. Gigan, F. Rigaut, P. Léna, C. Boyer, P. Jagourel, J.-P. Gaffard, and F. Merkle, "First diffraction-limited astronomical images with adaptive optics," *Astron. Astrophys.* **230**, pp. 29–32, 1990.
5. M. C. Roggemann, "Limited degree-of-freedom adaptive optics and image reconstruction," *Appl. Opt.* **30**(29), pp. 4227–4233, 1991.
6. J. M. Conan, P. Y. Madec, and G. Rousset, "Image formation in adaptive optics partial correction," in *Active and Adaptive Optics*, F. Merkle, ed., *ESO Conference and Workshop Proceedings* **48**, pp. 181–186, ESO/ICO, (Garching bei München, Germany), 1994.
7. J.-M. Conan, *Étude de la correction partielle en optique adaptative*. PhD thesis, Université Paris XI Orsay, Oct. 1994.
8. J. G. McNally, T. Karpova, J. Cooper, and J. A. Conchello, "Three-dimensional imaging by deconvolution microscopy," *Methods* **19**, pp. 373–385, 1999.
9. J. C. Christou, A. Roorda, and D. R. Williams, "Deconvolution of adaptive optics retinal images," *J. Opt. Soc. Am. A* **21**, pp. 1393–1401, Aug. 2004.
10. A. Tikhonov and V. Arsenin, *Solutions of Ill-Posed Problems*, Winston, DC, 1977.
11. G. Demoment, "Image reconstruction and restoration: Overview of common estimation structures and problems," *IEEE Trans. Acoust. Speech Signal Process.* **37**, pp. 2024–2036, Dec. 1989.
12. J. Idier, ed., *Bayesian Approach for Inverse Problems*, ISTE, London, 2008.
13. D. Gratadour, L. M. Mugnier, and D. Rouan, "Sub-pixel image registration with a maximum likelihood estimator," *Astron. Astrophys.* **443**, pp. 357–365, Nov. 2005.

14. J.-M. Conan, L. M. Mugnier, T. Fusco, V. Michau, and G. Rousset, "Myopic deconvolution of adaptive optics images by use of object and point spread function power spectra," *Appl. Opt.* **37**, pp. 4614–4622, July 1998.
15. L. M. Mugnier, T. Fusco, and J.-M. Conan, "MISTRAL: a myopic edge-preserving image restoration method, with application to astronomical adaptive-optics-corrected long-exposure images.," *J. Opt. Soc. Am. A* **21**, pp. 1841–1854, Oct. 2004.
16. A. Blanc, L. M. Mugnier, and J. Idier, "Marginal estimation of aberrations and image restoration by use of phase diversity," *J. Opt. Soc. Am. A* **20**(6), pp. 1035–1045, 2003.
17. D. Gratadour, D. Rouan, L. M. Mugnier, T. Fusco, Y. Clénet, E. Gendron, and F. Lacombe, "Near-IR AO dissection of the core of NGC 1068 with NaCo," *Astron. Astrophys.* **446**, pp. 813–825, Feb. 2006.