

3D deconvolution of adaptive-optics corrected retinal images

G. Chenegros ^a, L. M. Mugnier ^a, F. Lacombe ^b, and M. Glanc ^c

^aOffice National d'Études et de Recherches Aérospatiales, Optics Department, BP 72, F-92322 Châtillon cedex, France

^bMauna Kea Technologies, 9 rue d'Enghien 75010 Paris, France

^cLaboratoire d'Études Spatiales et d'Instrumentation en Astrophysique, Observatoire de Paris-Meudon, 5 place Jules Janssen, 92195 Meudon cedex, France

ABSTRACT

We report on a deconvolution method developed in a Bayesian framework for adaptive-optics corrected images of the human retina. The method takes into account the three-dimensional nature of the imaging process; it incorporates a positivity constraint and a regularization metric in order to avoid uncontrolled noise amplification. This regularization metric is designed to simultaneously smooth noise out and preserve edges, while staying convex in order to keep the solution unique. We demonstrate the effectiveness of the method, and in particular of the edge-preserving regularization, on realistic simulated data.

Keywords: image restoration, deconvolution, three-dimensional imaging, adaptive optics, retinal imaging, microscopy.

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, 6, 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 somewhat 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 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 sensing³ and in deconvolution of AO-corrected images.⁹

Besides, because deconvolution is an ill-posed inverse problem,^{10, 11, 12} most modern deconvolution methods use regularization in order to avoid an uncontrolled amplification of the noise. The regularization that is commonly used in 3D deconvolution is the classical Tikhonov regularization, which is quadratic (see Sect. 3) and thus tends to oversmooth edges.

In this communication, we introduce an edge-preserving regularization into a 3D deconvolution method and apply it to simulated retinal images.

Further author information: (Send correspondence to G.C.)

G.C.: E-mail: Guillaume.Chenegros@onera.fr; L.M.M.: E-mail: Laurent.Mugnier@onera.fr; F.L.: francois@maunakeatech.com .

2. PROBLEM FORMULATION

The image formation is modelled as a 3D convolution:

$$\mathbf{i} = \mathbf{h} \star \mathbf{o} + \mathbf{n} \quad (1)$$

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

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

$$i_k = \left(\sum_{j=0}^{N-1} h_{k-j} \star o_j \right) + n_k \quad (2)$$

where o_j is the object in plane j , i_k is the k -th recorded image and h_{k-j} is the 2D PSF corresponding to a defocus of $(k-j)$ slices. 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.

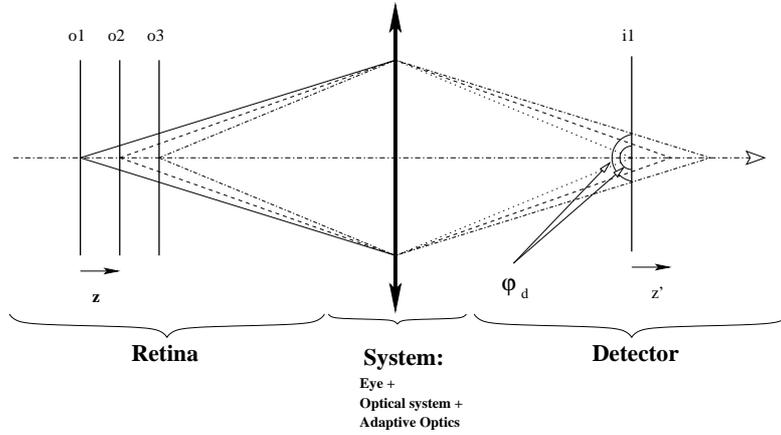


Figure 1. Illustration for 3 object planes. \mathbf{o} is on the left and \mathbf{i} is on the right. φ_d is the defocus phase.

Note that the raw image must be carefully preprocessed to yield an image that closely follows this imaging model. The preprocessing includes the relative recentering of the raw images¹³ 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 $\hat{\mathbf{o}}$ of the observed 3D object \mathbf{o} given the images \mathbf{i} , the 3D PSF \mathbf{h} , and some prior information on the noise statistics and on the object.

3. DECONVOLUTION APPROACH

3.1. 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 Bayesian¹² Maximum *A Posteriori* (MAP) approach, the deconvolution problem can be stated as follows: we look for the most likely object $\hat{\mathbf{o}}$ given the observed image \mathbf{i} and our prior information on \mathbf{o} , which is summarized by a probability density $p(\mathbf{o})$. This reads:

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

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

$$J(o) = J_i(o) + J_o(o), \quad (4)$$

where the negative log-likelihood $J_i = -\ln p(i|o)$ is a measure of fidelity to the data and $J_o = -\ln p(o)$ 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 o truly be the outcome of a stochastic process; rather, $p(o)$ should be designed to embody the available prior information on 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 o is not the outcome of a stochastic process, J_o usually includes a scaling factor or global hyperparameter, denoted by μ in the following, which adjusts the balance between fidelity to the data and fidelity to the prior information.

If no prior knowledge is used, which corresponds to setting $p(o) = \text{constant}$ in Equation (3), one then maximizes $p(i|o)$ (likelihood of the data) so that the solution is a maximum-likelihood (ML) solution. In this case the criterion of Eq. (4) is only constituted of the term J_i . The Richardson-Lucy algorithm,¹⁴ also known as ML-EM (for Maximum Likelihood–Expectation Maximization) is an example of an iterative algorithm which converges towards the minimum of J_i when the noise follows Poisson statistics.

3.2. Noise model

If the noise statistics are additive, stationary, white Gaussian, then the data fidelity term is a simple least-square difference between the actual data i and our model of the data for a given object, $h \star o$:

$$J_i(o) = \sum_{l,m,n} \frac{1}{2\sigma^2} [i(l, m, n) - (h \star o)(l, m, n)]^2, \quad (5)$$

where σ^2 is the total noise variance. In reality, the noise is actually a mixture of non-stationary, Poisson-distributed, photon noise and stationary white Gaussian detector noise. Yet, photon noise statistics can be effectively approximated as non-stationary white Gaussian as soon as the flux level is a few tens of photo-electrons per pixel.¹⁵ Additionally, because the images considered here are illuminated rather uniformly (due to all the out-of-focus object planes contributing to each image), a stationary white Gaussian statistics, with a constant variance equal to the mean number of photo-electrons per pixel, is a reasonable approximation for the noise model.

3.3. Object prior

This section aims at deriving an object prior for objects that are either smooth or piecewise smooth.

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. In this case, a reasonable model of the object’s power spectral density (PSD) can be found¹⁶ and used to derive the regularization criterion J_o , which is then quadratic (or “L2” in short). Additionally, the parameters of the object’s PSD can be estimated automatically (*i.e.*, in an unsupervised way) from the data itself by a Maximum Likelihood method.¹⁷

The disadvantage of a Gaussian prior (or equivalently of a quadratic regularization term), especially for objects with sharp edges such as photoreceptors or vessels, is that it tends to over-smooth edges. A possible remedy is to use an edge-preserving prior that is quadratic for small gradients and linear for large ones.¹⁸ The quadratic part ensures a good smoothing of the small gradients (*i.e.*, of noise), and the linear behavior cancels the penalization of large gradients (*i.e.*, of edges), as explained by Bouman and Sauer.¹⁹ Such priors are called quadratic-linear, or L2–L1 in short.²⁰ Here we use a function that is an isotropic version of the expression suggested by Rey²¹ in the context of robust estimation, used by Brette and Idier²² for image restoration, and recently applied to imaging through turbulence^{2, 15}:

$$J_o(o) = \mu \delta^2 \sum_{l,m,n} \phi(\nabla o(l, m, n)/\delta), \text{ where} \quad (6)$$

$$\phi(x) = |x| - \ln(1 + |x|), \quad (7)$$

and where $\nabla o(l, m, n) = [\nabla_x o(l, m, n)^2 + \nabla_y o(l, m, n)^2]^{1/2}$, and $\nabla_x o$ and $\nabla_y o$ are the object finite-difference gradients along x and y , respectively. This functional J_o is indeed L2–L1 because $\phi(x) \approx x^2/2$ for x close to 0 and $\phi(x)/|x| \rightarrow 1$

for $x \rightarrow \pm\infty$. Thus, parameter δ is a (soft) threshold, in the sense that J_o switches, at each pixel (l, m, n) , between the quadratic and the linear behaviors depending on whether $\nabla o(l, m, n)$ is smaller than or greater than δ .

The global factor μ and the threshold δ have to be adjusted according to the noise level and the structure of the object. These two hyperparameters currently have to be adjusted by hand.

The functional J_o is strictly convex because $\phi''(x) = 1/(1 + |x|)^2 > 0$ and J_i of Eq. (5) is convex because it is quadratic, so that the global criterion $J = J_i + J_o$ is strictly convex. This ensures uniqueness and stability of the solution with respect to noise and also justifies the use of a gradient-based method for the minimization.

4. VALIDATION ON SIMULATED DATA

4.1. Simulation

Figure 2 represents the original simulated object, composed of vessels, ganglion cells and photoreceptors. The vessels are simulated by moving a ring in a random walk, the ganglion cells are simulated by empty globes and photoreceptors are represented by two empty half spheres joined by an empty tube. The cube's height on Figure 2 is approximately $300\mu m$ and the depth and the width of this cube is 301 pixels.

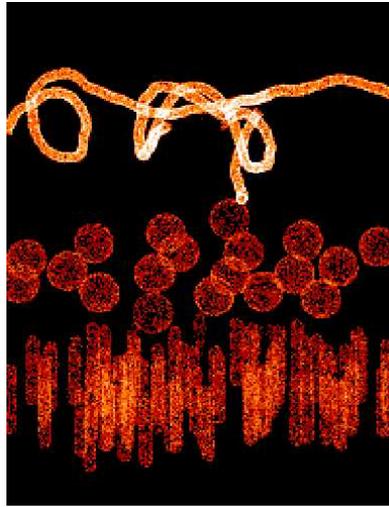


Figure 2. Perspective view of the 3D object used for the simulations

In order to create our simulated images with a five slice object, we average the data from Figure 2 into five $60\mu m$ -thick slices. Additionally, in the simulations presented here, we only use a 128×128 region of the 301×301 volume. The five slices obtained are presented on Figure 3.

The PSF's used to compute the image i are generated with time-invariant aberrations expanded on the Zernike basis; we use 0.2 rd of astigmatism (Z_5), -0.1 rd of astigmatism (Z_6) and -0.5 rd 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 Eq. (2). The noise added is white Gaussian and stationary; its standard deviation is 3% of the maximum intensity in the object o . The five image layers are presented on Figure 4. From these images, it is clear that all object slices contribute to all images (in particular the vessels). The deconvolution aims at disentangling the contribution of each object slice and at improving the resolution within each plane.

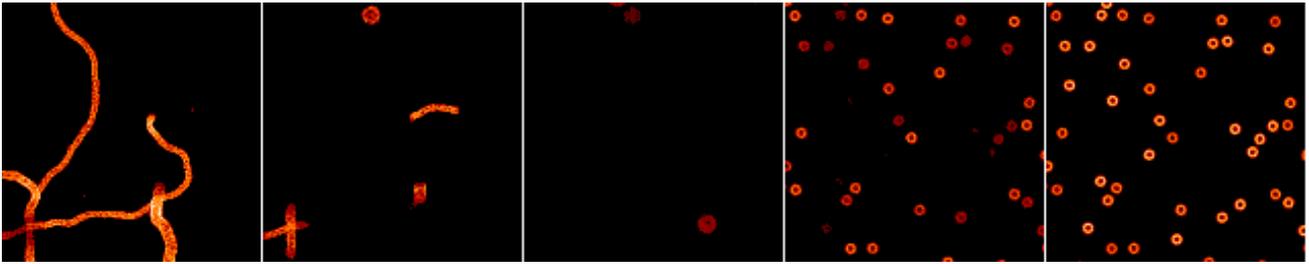


Figure 3. The five object layers (o).

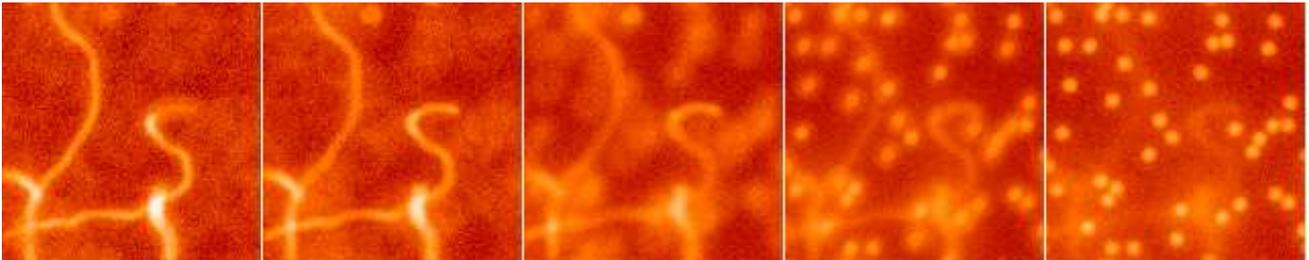


Figure 4. The five image layers (i).

4.2. Restoration

In this section, we present three results obtained with our deconvolution method:

L2 regularization without a positivity constraint: we can see on Figure 5 the five estimated object slices. The restoration error is 10.31 ph/pix (the object is 15.34 ph/pix on average). We can see low frequency oscillations 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. Edges are not preserved (*L2* regularization and no positivity constraint prevents spectral extrapolation).

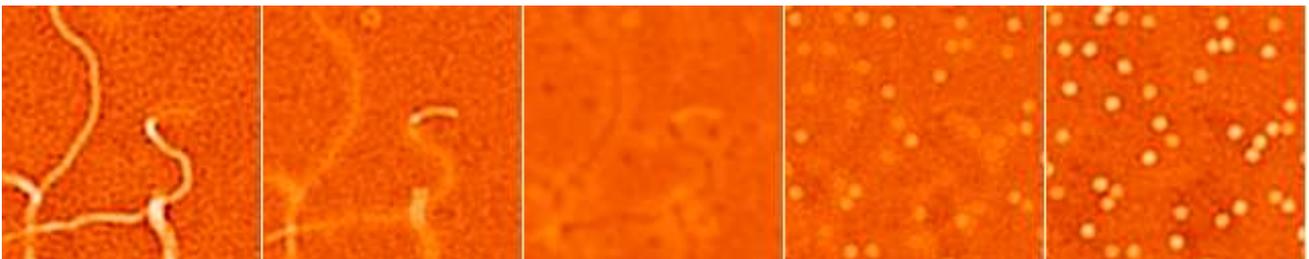


Figure 5. The five estimated object layers with *L2* regularization, without a positivity constraint.

L2 regularization with a positivity constraint: on Figure 6 we can see that the positivity constraint helps the algorithm disentangle the different planes and visibly reduces the low frequency oscillations. More quantitatively, the restoration error is 8.34 ph/pix. Yet, edges are not well preserved.

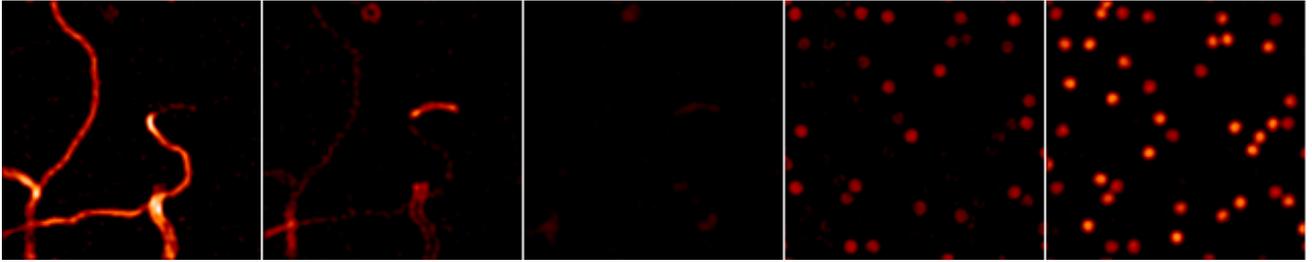


Figure 6. The five estimated object layers with L_2 regularization and a positivity constraint.

L_2 – L_1 regularization with a positivity constraint: on the Figure 7 we can see that the edges are much better preserved and the separation between the different planes is also slightly better on the second restored image plane. The restoration error is 6.33 ph/pix.

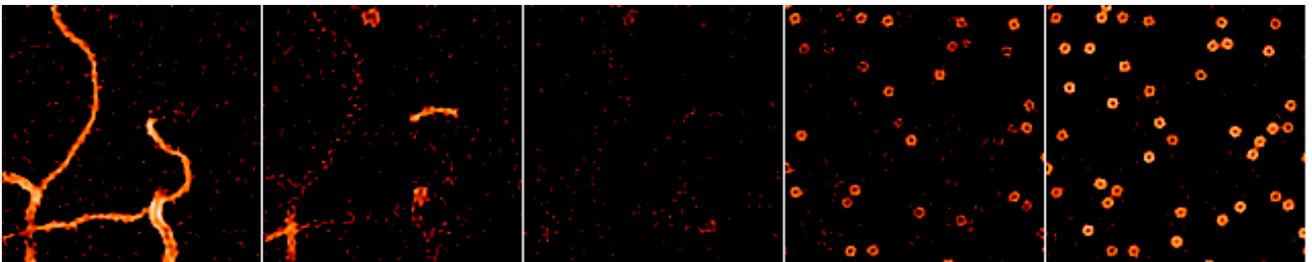


Figure 7. The five estimated object layers with L_2 – L_1 regularization and a positivity constraint.

To study the influence of the positivity constraint we compute three new piles of images (i) with different levels of background and we perform a deconvolution with L_2 – L_1 regularization. In these simulations, we introduce a continuous background in o and we evaluate the positivity constraint effects on the estimated object \hat{o} . These images are 128×128 pixels and the level of the noise is the same as previously but here the PSF used is not oversampled. In Figure 8 we present the results on the last slice where we see the photoreceptors; we can see that the positivity constraint helps reduce noise and ringing in the dark regions of the image, *i.e.*, where it is actually enforced. If the object o contain a continuous background (e.g., the diffusion), the positivity constraint is less efficient as shown in the middle and bottom rows of Figure 8: the higher the background, the less efficient the positivity constraint. In short, the use of a positivity constraint, if the object is on a dark background, is very effective in filling the cone of missing frequencies.

5. 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. The regularization metric simultaneously smoothes noise out and preserves edges, while staying convex in order to keep the solution unique. We have demonstrated the effectiveness of the method, and in particular of the edge-preserving regularization, on realistic simulated data.

Future work includes the processing of experimental data. For this purpose, it is of paramount importance to use a PSF that is a very close approximation of the true one in order not to produce deconvolution artefacts. As this is quite difficult, another appealing perspective is to use blind (or myopic) deconvolution.

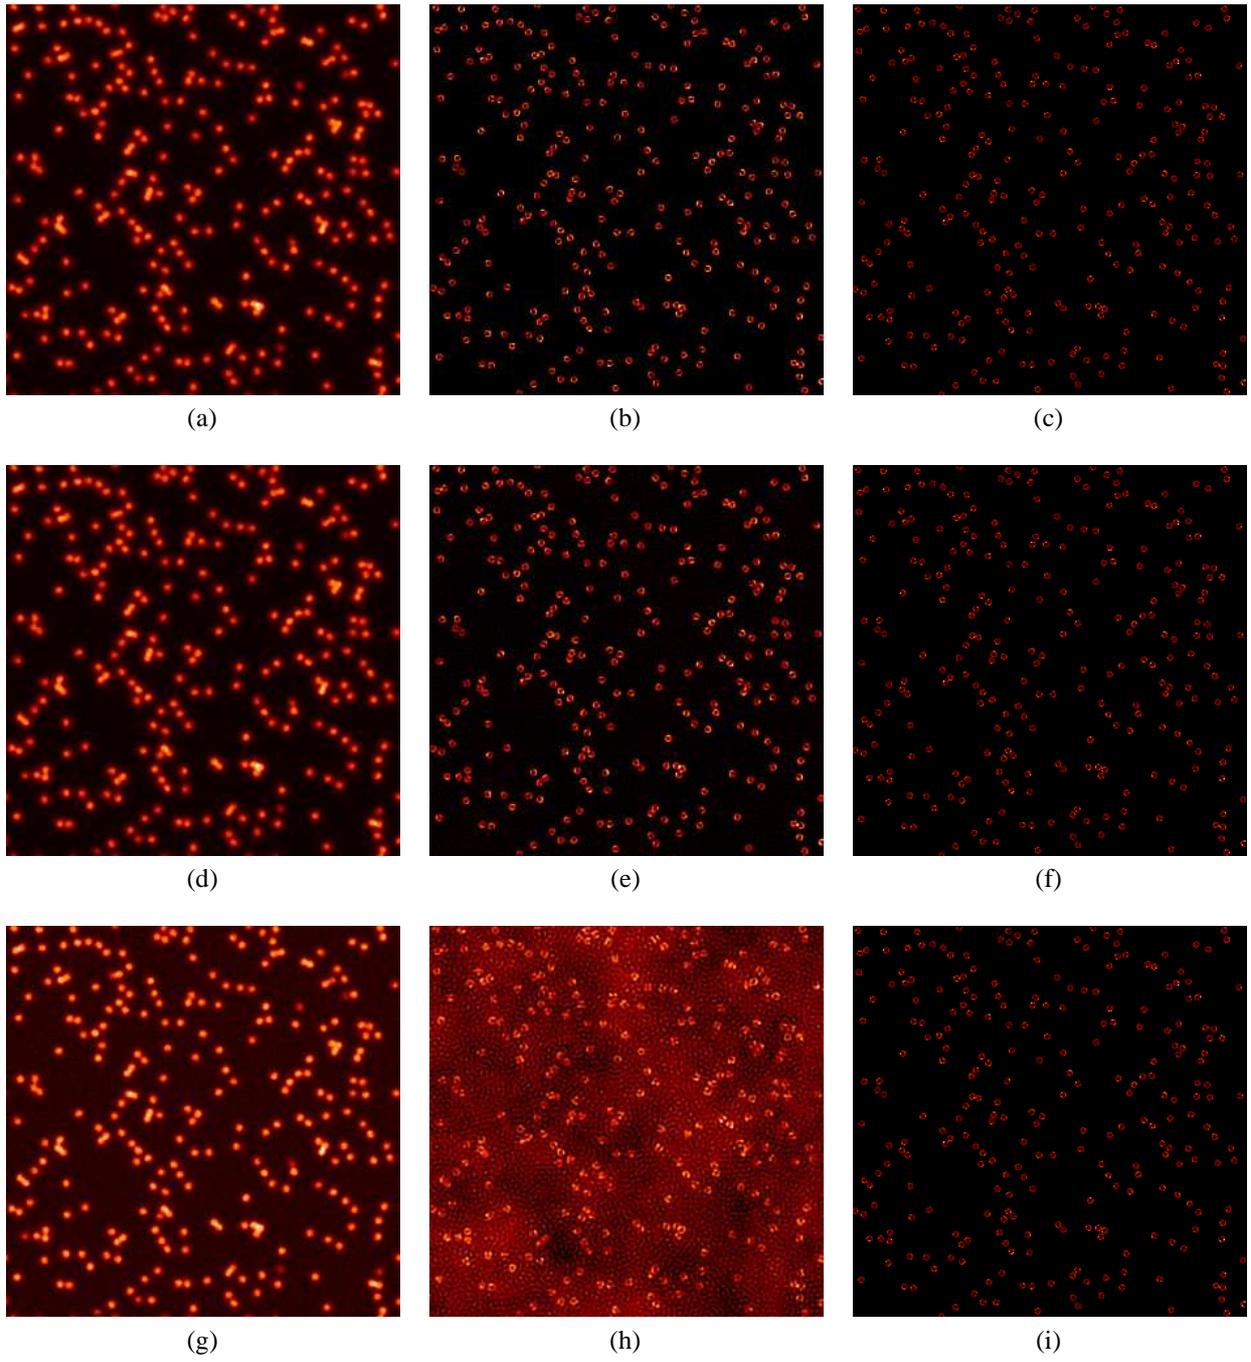


Figure 8. From left to right, a simulated recorded image, the estimated object with a positivity constraint and the real object. From top to bottom, data without continuous background, with a 5% intensity background and with a 50% intensity background.

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., *Approche bayésienne pour les problèmes inverses*, Hermès, Paris, 2001.
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. W. H. Richardson, "Bayesian-based iterative method of image restoration," *J. Opt. Soc. Am.* **62**(1), pp. 55–59, 1972.
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. J.-M. Conan, L. M. Mugnier, T. Fusco, V. Michau, and G. Rousset, "Myopic deconvolution of adaptive optics images using object and point spread function power spectra," *Appl. Opt.* **37**, pp. 4614–4622, July 1998.
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.*, to be published.
18. P. J. Green, "Bayesian reconstructions from emission tomography data using a modified EM algorithm," *IEEE Trans. Med. Imag.* **9**, pp. 84–93, Mar. 1990.
19. C. Bouman and K. Sauer, "A generalized gaussian image model for edge-preserving map estimation," *IEEE Trans. Image Processing* **2**, pp. 296–310, July 1993.
20. J. Idier and L. Blanc-Féraud, "Déconvolution en imagerie," in *Approche bayésienne pour les problèmes inverses*, J. Idier, ed., ch. 6, pp. 139–165, Hermès, Paris, 2001.
21. W. J. J. Rey, *Introduction to robust and quasi-robust statistical methods*, Springer-Verlag, Berlin, 1983.
22. S. Brette and J. Idier, "Optimized single site update algorithms for image deblurring," in *Proceedings of the International Conference on Image Processing*, pp. 65–68, IEEE, IEEE Computer Society, (Los Alamitos, California), 1996.