EURASIP Journal on Advances in Signal Processing
a SpringerOpen Journal

RESEARCH Open Access

Efficient shift-variant image restoration using deformable filtering (Part II): PSF field estimation

David Miraut[1], Johannes Ballé[2] and Javier Portilla[3*]

Abstract

We present a two-step technique for estimating the point spread function (PSF) field from a single star field image affected by shift-variant (SV) blur. The first step estimates the best-fitting PSF for each block of an overlapping block grid. We propose a local image model consisting of a pattern (the PSF) being replicated at arbitrary locations and with arbitrary weights. We follow an efficient alternate marginal optimization approach for estimating (1) the most likely pattern, and (2) the locations where it appears in the block, with sub-pixel accuracy. The second step uses linear dimensionality reduction and nonlinear spatial filtering for estimating the entire PSF field from the grid of local PSF estimates. We simulate SV blur on realistic synthetic star fields to assess the accuracy of the method for this kind of images, for different blurs, star densities, and Poisson counts. The results indicate a moderately low error and very robust behavior against noise and artifacts. We also apply our method to real astronomical images, and demonstrate that the method provides relevant information about the underlying structure of the actual telescope and atmosphere PSF fields. We use a variant of the method proposed in Part I to compensate for the observed blur.

Keywords: PSF estimation, PSF field, PSF field estimation, Shift variant blur, Deformable kernel, Dimensionality reduction, Maximum likelihood, Sparsity, Star fields

1 Introduction

Shift-variant (SV) image restoration requires knowledge of the point spread function (PSF) at each image location. If we have access to the imaging device, and the capture conditions are known, it may be possible to obtain the PSF field by pre-calibration. In many practical situations, though, we do not have access to this information. Moreover, the PSF field may change due to factors that are beyond our control (e.g., atmospheric turbulences, device temperature, vibrations, fog, relative movement of camera and objects). A given PSF field must then be estimated solely from the observed image(s) we want to restore.

The most common approach for estimating the blurring kernel from a single image consists of formulating a joint optimization problem, often based on statistical models of the image and the degradation, which is used to estimate both the uncorrupted image and the kernel. Typically, a starting guess is refined by alternating between estimating the image (assuming the kernel is known) and the PSF kernel (assuming the image is known). Apart from other potential problems like convergence and stability, adding such an outer loop to the estimation makes it especially heavy in computational terms, as each of these marginal estimations is, by itself, computationally expensive. To the best of the authors' knowledge, the intrinsic complexity of these approaches has prevented their application to the general case of SV blur kernels (although special cases have been treated, like considering foreground and background layers, e.g., [1]).

Furthermore, as blind restoration problems are intrinsically highly ill posed, stable solutions have been obtained mostly using some prior information about the kernel, adapting the problem solution to particular situations. Many methods benefit computationally from using a restricted (e.g., parametric) PSF model: Prior knowledge is applied in camera motion [2,3], parametric models are used in defocus and lens aberration correction [4,5],

*Correspondence: javier.portilla@csic.es
[3]Instituto de Óptica, Consejo Superior de Investigaciones Científicas, Madrid, Spain
Full list of author information is available at the end of the article

sparse characteristics are exploited for frosted glass [6], etc. In the specific case of PSF field estimation in astronomical images, several authors have studied orthogonal representations to characterize observed PSFs from stars [7,8]. The underlying motivation is to provide a robust tool against noise, flexible and possibly adaptive, without imposing a narrow structure to the PSF field. They may also facilitate PSF field interpolation (see, e.g., [9,10]). In the data-adaptive case, the linear dimensionality reduction (by singular value decomposition (SVD) or principal component analysis) deserves to be mentioned, which connects directly to the idea of deformable filters [11] approached in the companion paper, Part I. Function bases are optimal in a least-square sense, and they are ranked in terms of energy contribution to the PSF field description, so they can be selected to filter measurement noise in the estimation in a way that can be easily automated with techniques like Generalized Cross Validation [12]. Not surprisingly, such dimensionality reduction techniques have previously been used to estimate the PSF field [13]. However, they have typically been used for special cases, such as for modeling particular devices (for instance, the Large Synoptic Survey Telescope [14] or the Advanced Camera for Surveys on Hubble space Telescope [15]), adapted to particular stars as the ideal PSF reference, or for modeling gravitational lenses [16,17]. Hence, these approaches lack generality in the sense that they cannot be used as a general, "knowledge blind" tool for astronomical PSF field estimation. Only a few of the referred methods (like [18]) intend to apply their PSF estimation results to image deblurring, and none of them draw the clear connection between image restoration and PSF field estimation.

To alleviate the previously mentioned computational bottleneck, many techniques analyze the image in search of local features which provide information of the PSFs across the image, such as edges and corners [4]. These groups of pixels give clues which are used to compute point and line spread functions without user intervention. The idea of extracting *local* information about the PSF to build *local* PSF estimates is especially relevant for the SV blur case. Considering the lack of generality of the most common approaches outlined above, and taking the benefits of utilizing local estimates into account, we design here a novel approach to astronomical PSF field estimation. It is based on two simple observations relating to the fundamental characteristics of astronomical imagery, as opposed to typical photographic images.

First, astronomical images typically contain an abundance of stars, which can be modeled well as ideal point light sources. While only a narrow set of typical photographic images present enough repeated patterns, such as bright distant lights, to be used directly for characterizing the PSF field, the presence of stars in astronomical images allows for a particularly simple way to locally estimate the PSF, by weighted averaging, after sub-pixel localization. Note that, even under the smoothness assumption, obtaining a more or less dense set of local PSF estimates is not enough to fully characterize the blur, and a certain type of regularization and interpolation of that information to the rest of the image locations is necessary.

This is where our second observation comes in: we take advantage of the fact that the PSF field in astronomical images is usually simpler compared to typical photographic images, as there is no foreground–background structure. To a first approximation, PSFs result from the composition of the telescope and atmosphere PSF fields (for Earth-based observations), and they usually vary smoothly across the image field. Typical photographic images, on the other hand, may present a very rich variability in PSF field structures, due to differences in focus and/or relative speed of objects in the scene. Only in particular cases (e.g., camera movement and/or defocus, still long distance objects) one may expect to find a smooth PSF field. Furthermore, only in an even narrower set of cases, typical photographic images present enough repeated patterns (e.g., bright distant lights) to be used directly for characterizing the PSF field, similarly to the star field case.

Hence, our approach is twofold, and sequential. The first step consists of estimating the *most likely* local PSF for each block of a set of overlapping image blocks covering the whole image, according to a simple local image model. The second step extends and refines these local PSF estimates through linear dimensionality reduction, nonlinear filtering of outliers, and spatial interpolation. Whereas the first part is completely model-based, the second part is presented here, to a large extent, in an *ad hoc* way. This does not diminish its strong connection with the ideas presented in the companion paper, Part I, about using the concept of deformable kernels to deal with smooth PSF fields.[a]

There is a third fundamental characteristic of astronomical images that needs to be mentioned. Astronomical images have a huge dynamic range compared to typical photographic images, and very often, the search for relevant information requires processing low-contrast details in a contrast-adaptive fashion. Therefore, traditional error measures, such as quadratic error, must fail to describe the quality of a restoration method for star fields. To evaluate the quality of a restoration method on this kind of images is, thus, challenging. Quantitatively assessing the quality of the deblurring results has therefore not been tried here.

The rest of the article is organized as follows. Section 2 presents a simple image model for SV blurred star fields to be used as a local approximation for image blocks. Then, Section 3 describes an alternate optimization algorithm

for obtaining, through maximum likelihood, an estimate of the local PSFs. Section 4 describes a generic approach for estimating a smooth PSF field from a regular grid of local PSF estimates. A broad set of experiments, and their corresponding results, are described and discussed in Section 5. First (Subsection 5.1), objective measurements of the PSF field estimation quality are made by means of simulated star fields using two synthetic PSF fields. Then (Subsection 5.2), real astronomical images are analyzed to estimate their PSF field. These estimates are used for restoration of real images in Section 5.2.1. Section 6 concludes the article.

2 A smooth pattern field model for images

In this section, we describe a simple and yet effective image model for star fields subject to smoothly varying SV blur, e.g., due to a combination of atmospheric blur and non-ideal telescope optics. The model is conceptually much wider and potentially applicable to very different situations, like when observing repeated patterns in typical photographic images, possibly subject to certain variations depending on their image location. Although future instantiations of this model may attack the problem of directly characterizing the smooth pattern field for the whole image, we follow a simpler approach here, consisting of characterizing local regions of the image first by using a single reference pattern. Despite this simplification, we do consider in the local model that there is an error due to modeling the actual pattern within a block as constant. Figure 1 illustrates the idea of obtaining local PSF estimates on a regular spatial grid, using overlapping blocks.

2.1 Image model

We define a *pattern field* $p_{\boldsymbol{d}}(\boldsymbol{x}; \boldsymbol{d}(\boldsymbol{x}^0))$ as a special type of four-dimensional function which associates a certain spatial pattern $p_{\boldsymbol{d}}(\boldsymbol{x})$ to each image location $\boldsymbol{x}^0$. The spatial pattern $p_{\boldsymbol{d}}(\boldsymbol{x})$ depends on a vector of hidden parameters $\boldsymbol{d}(\boldsymbol{x}^0)$ (usually unknown *a priori*) which, in turn, depends on the image location $\boldsymbol{x}^0$. If the pattern depends smoothly (i.e., continuously and with continuous first derivative) on the parameter vector $\boldsymbol{d}$, and each component of $\boldsymbol{d}$ is a smooth function of $\boldsymbol{x}^0$, then $p_{\boldsymbol{d}}(\boldsymbol{x}; \boldsymbol{d}(\boldsymbol{x}^0))$ can be viewed as a *smooth pattern field*. The reader familiar with SV blur will already have noted the applicability of such smoothly varying pattern fields to characterize the PSF at any given image location, in some favorable cases (as discussed in the introduction). The star field, then, can be modeled as the result of blurring a set of ideal point sources (Dirac delta functions) with a shift variant PSF field:

$$y(\boldsymbol{x}) = y_{\mathrm{B}}(\boldsymbol{x}) + \sum_{i=1}^{M} a_i \, p_{\boldsymbol{d}}(\boldsymbol{x} - \boldsymbol{x}_i; \boldsymbol{d}(\boldsymbol{x}_i)) \tag{1}$$

This image model $y(\boldsymbol{x})$ is composed of a number of instances of the pattern field, each representing a local PSF, in this case, located at a position $\boldsymbol{x}_i$ and scaled by a

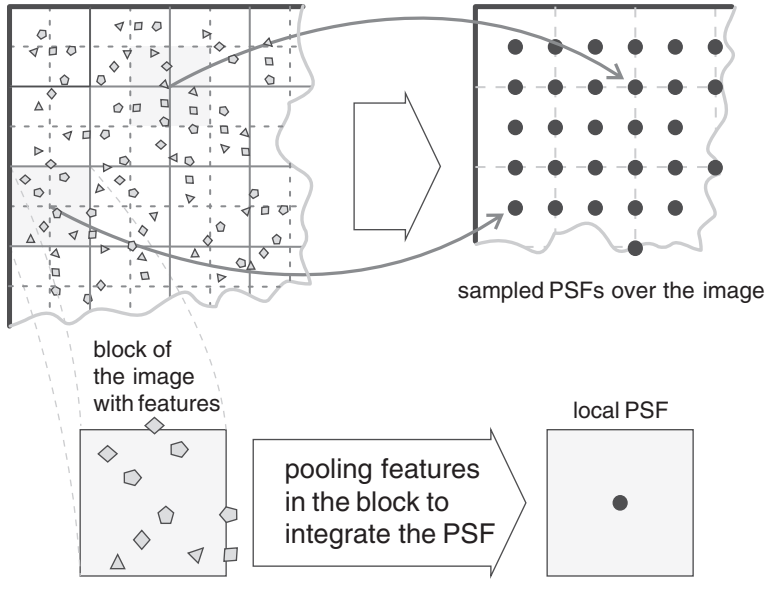

Figure 1 General strategy applied for estimating the PSF field. The first step is to estimate the local PSF on a grid of locations, using overlapping blocks. From these local estimates, the PSF field can be estimated, assuming it varies smoothly across the image. The method is presented here in its most general way (grouping local features).

factor a_i corresponding to the strength of the ideal point source. In order to make the model robust to changes of the background level, it is highly convenient to include a very smooth additional background term $y_{\mathrm{B}}(\boldsymbol{x})$.

In this application of the pattern field model, we have taken the approach of representing the patches by finite-dimensional vectors, corresponding to finite and discrete sampled patches in the image. A proper pixel size of the patch can easily be assigned by inspection of the most blurred areas of the image. Also, for algorithmic convenience, we have imposed that the patches do not overlap (in other applications of the model this may respond to an actual physical constraint). Although this is not a realistic assumption for stars, we have experienced that it does not significantly harm the estimation.

Note that, seen as a generative model, Equation 1 is fairly simple: a highly sparse input distorted by an SV blur, even linear for a given pattern field. However, from an analysis point of view, fitting an observation to this model involves a highly nonlinear procedure, mainly because of the sparse character of its input. Although powerful techniques for sparse estimation exist (e.g., [19]), they usually assume that the input is sparse, but not necessarily in such a way that the input is a set of (isolated) deltas, and much less that the "basis functions" are non-overlapping. On the other hand, our approach takes advantage of this particular structure of the problem.

2.2 Pattern observation model

The pattern field p generally is a nonlinear function for which there is no prior information except that it is smooth with respect to the image coordinates. If it is smooth enough, the Nyquist theorem limits the amount of information that is lost when we spatially *sample* the parameter vector $\boldsymbol{d}$. This takes us to a block-based approach, in which we model the PSF at a certain block location as a component along the PSF at the block center, plus a certain perturbation increasing with the distance to the center: $p_{\boldsymbol{d}}(\boldsymbol{x};\boldsymbol{d}(\boldsymbol{x}^0)) = bp_{\boldsymbol{d}}(\boldsymbol{x};\boldsymbol{d}(\boldsymbol{x}_C)) + \boldsymbol{\varepsilon}$, where $\boldsymbol{x}_C$ is the central location of the block $B(\boldsymbol{x}_C)$, and $\boldsymbol{x}^0 \in B(\boldsymbol{x}_C)$. Consequently, a given pattern observation, numbered i, at some position $\boldsymbol{x}_i$ in a given block, is modeled as:

$$\boldsymbol{y}(\boldsymbol{x}_i) = \boldsymbol{y}_{\mathrm{B}}(\boldsymbol{x}_i) + a_i\,(b_i\boldsymbol{p} + \boldsymbol{\varepsilon}_i) + \boldsymbol{n}_i \tag{2}$$

Here, $\boldsymbol{p}$ is the vectorized pattern at $\boldsymbol{x}_C$, b_i represents the fading of the component along the central pattern towards the boundaries of the block, $\boldsymbol{\varepsilon}_i \sim \mathrm{N}(\boldsymbol{0}, \sigma_{\varepsilon i}^2\mathbf{I})$ represents the deviation of the pattern with respect to the $\boldsymbol{p}$ direction, and $\boldsymbol{n}_i \sim \mathrm{N}(\boldsymbol{0}, \sigma_n^2\mathbf{I})$ is a noise term. Note that b_i will be 1 (or less, in practice, if our knowledge of $\boldsymbol{p}$ is approximate) at the block center, and will decrease as we move away from it, while $\sigma_{\varepsilon i}$ is modeled as an ever increasing but upper bounded function of $||\boldsymbol{x} - \boldsymbol{x}_C||$, thus behaving in a complementary way to b_i.

Despite of the apparent intricacy of this model, it can be reformulated into a much simpler one. First, observe that we can decompose $\boldsymbol{p}$ into

$$\boldsymbol{p} = \boldsymbol{p}_0 + \mu_p\mathbf{1} \tag{3}$$

where $\mu_p\mathbf{1} = [\mu_p, \ldots, \mu_p]^{\mathrm{T}}$ represents the sample pixel mean of the reference pattern, and $\boldsymbol{p}_0$, consequently, its zero mean version.[b] Second, as the background $y_{\mathrm{B}}(\boldsymbol{x})$ is smooth, it can locally be approximated by a constant:

$$\boldsymbol{y}_{\mathrm{B}}(\boldsymbol{x}_i) \approx \mu_{\mathrm{B},i}\mathbf{1} \tag{4}$$

Since we are unconcerned about estimating the image background or the ground level of the zero-mean pattern, we may encapsulate these components of the model, thus simplifying the model as well as the computational procedure:

$$\boldsymbol{y}(\boldsymbol{x}_i) - \underbrace{(\mu_{\mathrm{B},i} + a_i b_i \mu_p)}_{=\mu_i}\mathbf{1} = \underbrace{a_i b_i}_{=c_i}\boldsymbol{p}_0 + \underbrace{a_i\boldsymbol{e}_i + \boldsymbol{n}_i}_{\text{random}} \tag{5}$$

The simplified patch model can be restated as follows:

$$\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1} = c_i\boldsymbol{p}_0 + \boldsymbol{w}_i \tag{6}$$

where $\boldsymbol{w}_i \sim \mathrm{N}(\boldsymbol{0}, (a_i^2\sigma_{\varepsilon_i}^2 + \sigma_n^2)\mathbf{I})$ now is a joint "error" term representing the uncertainty of the observation.

3 Local PSF estimation by alternating marginal likelihood optimization

Since the observed star patterns are independent, the likelihood of observing a number M of patterns in the block is

$$-\log L \propto \sum_{i=1}^{M}\left(K\log(a_i^2\sigma_{\varepsilon i}^2+\sigma_n^2) + \frac{\|\boldsymbol{y}(\boldsymbol{x}_i)-\mu_i\mathbf{1}-c_i\boldsymbol{p}_0\|^2}{a_i^2\sigma_{\varepsilon i}^2+\sigma_n^2}\right) + (N-MK)\log\sigma_n^2 + \sum_{\boldsymbol{x}\in R}\frac{\left|y(\boldsymbol{x})-y_{\mathrm{B}}(\boldsymbol{x})\right|^2}{\sigma_n^2} + const, \tag{7}$$

where K is the dimensionality of the vectorized neighborhoods, N is the number of pixels in the block, and R is the set of pixels that are not included in any pattern. Optimizing the likelihood directly for all the unknowns is a difficult problem. However, it can effectively be attacked by alternating between optimizing the marginal likelihood as a function of $\boldsymbol{p}_0$ and the set of pattern locations, $\{\boldsymbol{x}_i, i = 1\ldots M\}$, respectively.

The algorithm we propose essentially comprises three stages:

1. taking an initial guess at the pattern;
2. finding approximate (full-pixel precision) locations of the pattern, as well as the number of locations, with respect to that guess; and
3. alternating between re-estimating the pattern and refining the locations with sub-pixel precision.

Likelihood maximization with respect to each of μ_i, c_i is a simple least-squares fitting, and can be done independently for each image location $\boldsymbol{x}_i$:

$$\hat{\mu}_i = \frac{\boldsymbol{y}^{\mathrm{T}}(\boldsymbol{x}_i)\mathbf{1}}{K} \tag{8}$$

$$\begin{aligned}\hat{c}_i &= \frac{(\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1})^{\mathrm{T}}\boldsymbol{p}_0}{\|\boldsymbol{p}_0\|^2} \\ &= \frac{\|\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\|}{\|\boldsymbol{p}_0\|}\cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\right).\end{aligned} \tag{9}$$

For estimating a_i we observe in (5) that:

$$\mathrm{E}\{\|\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\|^2\} = a_i^2\mathrm{E}\{\|b_i\boldsymbol{p}_0 + \boldsymbol{\varepsilon}_i\|^2\} + K\sigma_n^2 \tag{10}$$

Assuming $\mathrm{E}\{\|b_i\boldsymbol{p}_0 + \boldsymbol{\varepsilon}\|^2\} = 1$, we simply choose the estimate

$$\hat{a}_i^2 = \left\lfloor \|\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\|^2 - K\sigma_n^2 \right\rfloor_+ \tag{11}$$

where the brackets indicate non-negative clipping. For those selected patches with norms well above the background noise level ($K\sigma_n^2$), the effect of a_i on the variance is to normalize the local energy of the patch, so effectively giving a stronger weight to candidates with low relative error (high SNR). The low energy patches, on the other hand, are effectively damped by the larger influence of σ_n^2 in their associated variance.

Also, for a given set of scaled and noisy pattern candidates, $\{\boldsymbol{y}(\boldsymbol{x}_i) - \hat{\mu}_i\mathbf{1}\}$, the solution for the most likely $\boldsymbol{p}_0$ is the standard least-squares one:

$$\hat{\boldsymbol{p}}_0 \propto \sum_{i=1}^{M} \frac{\|\boldsymbol{y}(\boldsymbol{x}_i) - \hat{\mu}_i\mathbf{1}\|\cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}_i) - \hat{\mu}_i\mathbf{1}\right)}{a_i^2\sigma_{\varepsilon i}^2 + \sigma_n^2}(\boldsymbol{y}(\boldsymbol{x}_i) - \hat{\mu}_i\mathbf{1}) \tag{12}$$

Note that $\hat{c}_i\boldsymbol{p}_0$ is invariant with respect to $\|\boldsymbol{p}_0\|$. Hence, we only need to find the direction of $\hat{\boldsymbol{p}}_0$ in (12). Consequently, our method is invariant with respect to the normalization of the PSF. In order to avoid numerical error, it is still useful to constrain $\|\boldsymbol{p}_0\|$. As it appears most natural in this context, we choose $\|\boldsymbol{p}_0\| = 1$. Recall that, in spite of this non-standard choice when dealing with PSF patterns, it is always possible to go back to standard PSF form (non-negative, normalized volume) by estimating μ_p and re-normalizing.

Methodically, the optimization for the locations $\{\boldsymbol{x}_i\}$ takes a different form for the initial (full pixel precision) and the subsequent (sub-pixel) estimates.

3.1 Estimation of the approximate locations of star patterns

When the $\{\boldsymbol{x}_i\}$ are estimated for the first time, we can make use of the observation that the last two terms in (7) (the "non-pattern" terms) are constant if extended to all pixels of the block. Then, we can remove that constant from the optimization functional to obtain

$$\begin{aligned}-\log L \propto \sum_{i=1}^{M}\Bigg(K\log(a_i^2\sigma_{\varepsilon i}^2 + \sigma_n^2) + \frac{\|\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1} - c_i\boldsymbol{p}_0\|^2}{a_i^2\sigma_{\varepsilon i}^2 + \sigma_n^2} \\ - K\log\sigma_n^2 - \frac{\|\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\|^2}{\sigma_n^2}\Bigg) + const,\end{aligned} \tag{13}$$

where we have applied the approximation (4) again.[c] Consider the expression inside the parentheses, evaluated as a function of a generic image location $\boldsymbol{x}$ and holding $\boldsymbol{p}_0$ fixed:

$$\begin{aligned}A(\boldsymbol{x}) = -K\log\left(\frac{a^2(\boldsymbol{x})\sigma_\varepsilon^2(\boldsymbol{x})}{\sigma_n^2} + 1\right) + \|\boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1}\|^2 \\ \times\left(\frac{1}{\sigma_n^2} - \frac{1 - \cos^2\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}_i) - \mu(\boldsymbol{x})\mathbf{1}\right)}{a^2(\boldsymbol{x})\sigma_\varepsilon^2(\boldsymbol{x}) + \sigma_n^2}\right),\end{aligned} \tag{14}$$

where we have applied

$$\begin{aligned}&\|\boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1} - c(\boldsymbol{x})\boldsymbol{p}_0\|^2 \\ &\quad = \|\boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1}\|^2(1 - \cos^2\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1}\right)),\end{aligned}$$

and where $\mu(\boldsymbol{x}) = \boldsymbol{y}^{\mathrm{T}}(\boldsymbol{x})\mathbf{1}/K$ is the sample pixel average of the patch. This function represents the log-likelihood ratio between two hypotheses: the patch centered on $\boldsymbol{x}$ corresponding to an instance of the pattern versus corresponding to the background. When the patch centered at $\boldsymbol{x}$, $\boldsymbol{y}(\boldsymbol{x})$, contains significantly more energy than that of the background noise, we have a pattern candidate. In that case, all but one of the terms in Equation (14) in the vicinity of $\boldsymbol{x}$ are smooth functions of $\boldsymbol{x}$, because they correspond to energy measurements which are insensitive to the exact location of the pattern, whenever it remains entirely within the patch. The exception is the covariance factor term, which is very selective to the exact pattern location, even under noise. Furthermore, in this vicinity $A(\boldsymbol{x})$ must yield high positive values (because $1/\sigma_n^2 \gg 1/(a^2(\boldsymbol{x})\sigma_\varepsilon^2(\boldsymbol{x}) + \sigma_n^2)$), peaking in an isolated maximum where the covariance factor gets closest to 1, i.e., when the pattern candidate is aligned with respect to the pattern guess $\boldsymbol{p}_0$. Therefore, a greedy algorithm, under the assumption of no pattern overlapping in the image, is optimal for finding the most likely locations for the current pattern guess. To estimate the set of M locations $\{\boldsymbol{x}_i\}$, we first pick the global maximum of $A(\boldsymbol{x})$ in the block as $\boldsymbol{x}_1$ and set M to 1. Next, we rule out the neighborhood around that maximum as a potential location for subsequent picks (to enforce the non-overlapping constraint in the analysis) and find the next-to-optimal maximum $\boldsymbol{x}_2$, increasing M by 1; and so on, until no further locations can be selected, or the next minimum attains a close-to-zero

value, indicating that nothing is gained by considering a star pattern at this location.

Practically, this algorithm fails to choose the correct locations only in the case where the initial assumption of non-overlapping patterns does not hold, i.e., in our case, when two or more stars of similar brightness are too close to each other (if one star clearly dominates, then the ML choice corresponds to centering that dominant star). However, these wrong choices can be easily detected by thresholding the covariance $\cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1}\right)$ between the current guess of the pattern and the image, and thus preventing these locations to contribute to the pattern estimate (in our experiments we used a threshold $c_{\min} = 0.2$).

3.2 Sub-pixel refinement

When estimating $\boldsymbol{x}_i$ in the subsequent iterations, we do not need to adjust M, but we would like to improve the locations $\boldsymbol{x}_i$ to sub-pixel precision. Observing that $\sigma_{\varepsilon i}$ and $\|\boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\|$ are approximately constant for such small changes of $\boldsymbol{x}_i$ (as the pattern is assumed to be approximately centered in the patch, with some background space around it), the optimization functional can be simplified to obtain the following estimate:

$$\hat{\boldsymbol{x}}_i \approx \arg\max_{\boldsymbol{x}_i} \cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\right) \tag{15}$$

As this functional is very simple to compute, we can afford to perform a local search around the previous location estimates simply by computing a standard, bicubic sub-pixel interpolation of the patch and evaluating the functional.[d] We choose to do this in a hierarchical fashion, that is, performing local search in $[-1, 1]$ first with a step size of $\Delta x = 1/2$ in each coordinate, updating, then performing a local search for the relative shift in $[-1/2, 1/2]$ with $\Delta x = 1/4$, and so on, down to a precision of one eighth of a pixel.

3.3 Variance estimates

The overall variance of the pixels of a patch including a pattern depends on the additive background noise term (σ_n^2), and the multiplicative error term ($a_i^2\sigma_{\varepsilon_i}^2$) due to the error on the pattern itself. The latter depends solely on the distance to the block center in relative terms; in absolute terms, it also depends on the scale factor a_i of the instance of the pattern. For estimating σ_n^2, we use the Median Absolute Deviation method with a simple difference filter along each dimension [20].

To model the variance of the error $\boldsymbol{\varepsilon}_i$, we assume that the pattern random vector field has an isotropic covariance structure similar to an AR(1) process. In particular,

$$\sigma_{\varepsilon i}^2 = \frac{1 - \rho^{\frac{2}{L/f_r}\|\boldsymbol{x}_i - \boldsymbol{x}_C\|}}{K} \tag{16}$$

and $\mathrm{E}\{\|b_i\boldsymbol{p}_0 + \boldsymbol{\varepsilon}_i\|^2\} = 1$, such that the energy is stationary across the pattern field. ρ is the parameter that determines the smoothness of the field, and L/f_r is the distance between neighboring block centers (L being the side length of the block, and f_r the overlapping block redundancy factor). Thus, $\mathrm{E}\{\|\boldsymbol{\varepsilon}_i\|^2\} = 0$ at the block center itself and $\mathrm{E}\{\|\boldsymbol{\varepsilon}_i\|^2\} = 1 - \rho^2$ at the neighboring block center. In our experiments, we chose $\rho = 0.98$ as a reasonable smoothness measurement for adjacent PSF in the grid made of the overlapping blocks' centers. In the estimation procedure (12), this structure of $\sigma_{\varepsilon i}^2$ leads to the desirable effect of giving a stronger weighting of the detected patterns that are near the block center.

In addition to uncertainty due to noise and spatial variation of the pattern, there is a methodical uncertainty with respect to the current estimate of $\boldsymbol{p}_0$. Particularly the first guess may be somewhat inaccurate; and since the weights in (12) are based on covariance measurements with the (guessed) pattern, subsequent applications of the estimator may weight useful instances of the star pattern lower than necessary, and vice versa. To make the method adapt to the quality of the current patch estimate, we add the term σ_g^2 to the r.h.s. of (16) when first performing the estimation of the full pixel pattern locations, defined as

$$\sigma_g^2 = \frac{1 - \left(\max_{\boldsymbol{x}} \cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1}\right)\right)^2}{K}. \tag{17}$$

This value corresponds exactly to the quadratic error in our guess if (1) there is an instance of the exact pattern in the image, and (2) that instance is the closest one in the image, in a covariance factor sense, to our guess. The use of this extra term permits a less shape-selective weighting of candidate patterns when the quality of the estimated pattern is still relatively low.

3.4 Initial and subsequent pattern guesses

In our experiments, for the first block we used as a guess a Gaussian PSF, with a spatial dispersion $\sigma = 2$, using a 19×19 patch support. For subsequent blocks, processed in a horizontal raster scan order, we use the average of the estimated PSFs for previously estimated blocks above and on the left (or only one of them, for the first row or the first column).

3.5 Summary of the local PSF estimation method

In Figure 2 (left) we have depicted a scheme summarizing our local PSF estimation method. We also provide an algorithmic description below:

Algorithm 1. Local PSF estimation in star fields

1: Obtain an initial guess $\hat{\boldsymbol{p}}_0$
2: **for all** blocks in the image **do**
3: Estimate σ_n^2 and σ_g^2
4: **for all** $\boldsymbol{x}$ in the block **do**

Compute σ_ε^2, a, μ, and c
Compute covariance factor among patch and pattern $\cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}) - \mu(\boldsymbol{x})\mathbf{1}\right)$
Compute $A(\boldsymbol{x})$ using (14)
end for
$X \leftarrow \{\}$
loop
$\hat{\boldsymbol{x}}_i \leftarrow \arg\max_{\boldsymbol{x}'} A(\boldsymbol{x}')$ s.t. $\boldsymbol{x}' \notin \mathcal{N}(\boldsymbol{x}), \forall \boldsymbol{x} \in X$
if $A(\hat{\boldsymbol{x}}_i) < T$ or $\hat{\boldsymbol{x}}_i$ has no solution **then**
$M \leftarrow |X|$
break
else
if $\cos\angle\left(\boldsymbol{p}_0, \boldsymbol{y}(\boldsymbol{x}_i) - \mu_i\mathbf{1}\right) > c_{min}$ **then**
$X \leftarrow X \cup \{\hat{\boldsymbol{x}}_i\}$
end if
end if
end loop
repeat
Re-estimate $\hat{\boldsymbol{p}}_0$ using (12)
for all $\hat{\boldsymbol{x}}_i \in X$ **do**
Improve $\hat{\boldsymbol{x}}_i$ by local search using (15)
end for
until convergence
Update the guess $\hat{\boldsymbol{p}}_0$ for the next block
end for

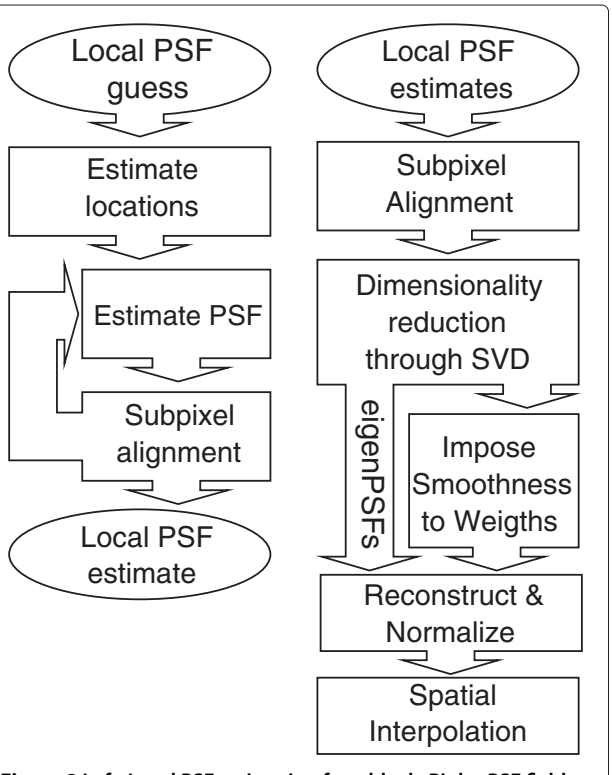

Figure 2 Left: Local PSF estimation for a block; Right: PSF field estimation from a grid of local PSF estimates.

4 Estimating smooth PSF fields from local estimates using SVD

When we observe a set of patterns under noise, relevant features can be extracted by means of finding the orthogonal basis that optimally (in a least-squares sense) compresses that particular set of patterns as a weighted sum of orthogonal components. The SVD has been in widespread use for many years to solve this problem.[e] In a high dimensional space, such as the one used to represent rich unconstrained patterns, like our PSFs, the noise effect on the obtained dominant terms is small, because, if it is white, it distributes its energy equally among all the orthogonal components. The terms above the noise level, on the other hand, may carry a large proportion of the total pattern energy. Therefore, those terms provide an approximation to the original (uncorrupted) set of patterns, when the components below the noise level are filtered out before reverting the linear transform.

This kind of linear spectral filtering can be effective to reliably estimate the dominant *eigenPSFs* of the PSF field, but it may not be enough to characterize the field by itself, when there are strong spurious fluctuations in the estimated patterns. These fluctuations may be caused by imperfections of the algorithm and/or the model (in our case, we do some incorrect assumptions, like that the PSFs do not overlap), or due to the presence of "outliers" (image contents that fit neither the dominant pattern nor the background statistics). Then, using prior information about the typical features of the pattern field is essential to increase robustness. In our problem, the local PSF estimates are far from being equally reliable (e.g., because of inhomogeneities of the local density of stars, different brightness, background effects, etc.). Using the prior knowledge of the PSF field being smooth helps us to discount the outliers' effect on the estimation, and, thus, to improve the results.

4.1 Estimating a smooth PSF field from a grid of local estimates

As the method of the previous section does not ensure alignment of the local estimates with respect to each other, they need to be aligned prior to being subjected to dimensionality reduction, as illustrated in Figure 2 (right).

4.1.1 Local PSF sub-pixel alignment

Note that the sub-pixel alignment performed here is conceptually slightly different from the intra-block alignment of patterns, as we do not want to optimize a similarity measures with respect to a certain reference pattern here, but rather maximize the similarity of each pattern with respect to the others. A sensible solution is to find the subpixel shifts that minimize the Euclidean error to the projection of each pattern onto the first eigenvector of the set, which is recomputed at each iteration (using a hierarchical

iterative approach using cubic splines, as before). It is straightforward to prove that, following this procedure we will reach a stationary solution for the spatial shifts which maximize (locally, at least) the first eigenvalue of the covariance matrix of the set of patterns, as it effectively minimizes the sum of squares of the components along orthogonal directions. Similarly to the intra-block alignment case, we have experienced a very significant decrease in the energy of the components orthogonal to the first eigenvector (around 4–5 dB) by doing this inter-block alignment. We observe a spatial drift of the estimated local pattern across image blocks. This undesirable effect does not affect the model fitting negatively.

4.1.2 Dimensionality reduction and nonlinear filtering of the spatial weights

After the alignment, we proceed to perform a linear dimensionality reduction of the set, using the SVD. Although it is not difficult to obtain a suitable number of singular vectors/values automatically, we have chosen the number simply by visual inspection of the eigenvalues profile. We obtain a set of eigenvectors (*eigenPSFs*), plus the corresponding optimal spatial weights to represent the grid of extracted local patterns as a linear combination of them. Then, the "outliers" can easily be detected and suppressed by iteratively detecting the spatial weights that largely deviate with respect to their neighbor values, and substituting them by the local mean of their neighbors. Once the spatial weights are regularized, we revert the orthogonal transform and obtain a spatial grid of regularized patterns. The ground level of these must be shifted back to zero and their volume normalized, in order to force both positivity and unit volume, as required for PSFs. We use a mode estimation to estimate the ground level of each pattern.

4.1.3 PSF field interpolation to intermediate spatial locations

We may want to obtain PSF estimates for each pixel of the image; for example, if we want to use the PSF field estimation to perform deblurring on the same or similarly captured images. Given a low-dimensional description of a vector at different image locations, such as obtained through the SVD, it is easy to achieve that by interpolating the scalar coefficient corresponding to each basis vector, provided the vector field is smooth. However, as we enforce the PSF constraints (positivity and unit volume) after reverting the transform, the resulting set of local PSF estimates are no longer confined to a low-dimensional linear subspace. One solution to this is to interpolate before reintroducing the constraints, and then fixing the ground level and volume for each pixel location in the image. However, a more efficient manner consists of performing a second SVD on the previously estimated set of positive and volume normalized PSFs, and retaining enough singular components for an excellent approximation (let us say, more than 60–80 dB). After this second dimensionality reduction, we can perform a vector interpolation very efficiently by simply interpolating the obtained weights as scalar functions, using cubic splines. This method ensures the right form of our PSF field estimate: a number of *eigenPSFs*, plus their corresponding smooth spatial weights (one weight for each eigenPSF–pixel pair). We may then apply the deformable filtering techniques explained in Part I to restore the observation using this estimate.

4.2 Relationship to the image estimation method proposed in Part I

The dimensionality reduction approach on the (aligned) PSFs is a shared approach of Parts I and II of this double article. In both cases, the local PSF is modeled as a linearly deformable kernel with a reduced number of dimensions. Whereas the main motivation for using deformable kernels in Part I was computational efficiency (as the linear deformable filtering can be implemented as a sum of masked convolutions), our main motivation in this second part is robustness in the characterization of a PSF field from which we only have partial and noisy information. More importantly, both techniques are connected in practice by the use one would give to them for image SV blind restoration: first estimating the PSF field (Part II), then using that estimate for applying the SV image restoration estimation described in Part I. We also believe that, by developing the PSF field estimation and the image estimation in the context of a single framework, we add value to each of the described techniques.

5 Experiments: results and discussion

5.1 PSF field estimation in simulated star fields

In order to obtain an objective reference to evaluate our PSF field estimation results, we have simulated both PSF fields and star field images. We degraded the latter by noise after the SV blur.

5.1.1 Simulating SV blurred star fields

We consider three different Poisson noise levels, corresponding to having maximum values for 8 bits (from 0 to 255), 12 bits (0 to 4095), and 16 bits (0 to 65535) in the simulated star fields. Also, we add zero mean Gaussian noise of unit variance. We have implemented a simulation for the star field giving roughly the same (relative, observed from the Earth) magnitude statistics as the observed stars, which approximately fit the law

$$\log n(m) = \alpha m + C,$$

n being the number of stars of a given apparent magnitude m. According to actual statistics of observable stars [21],

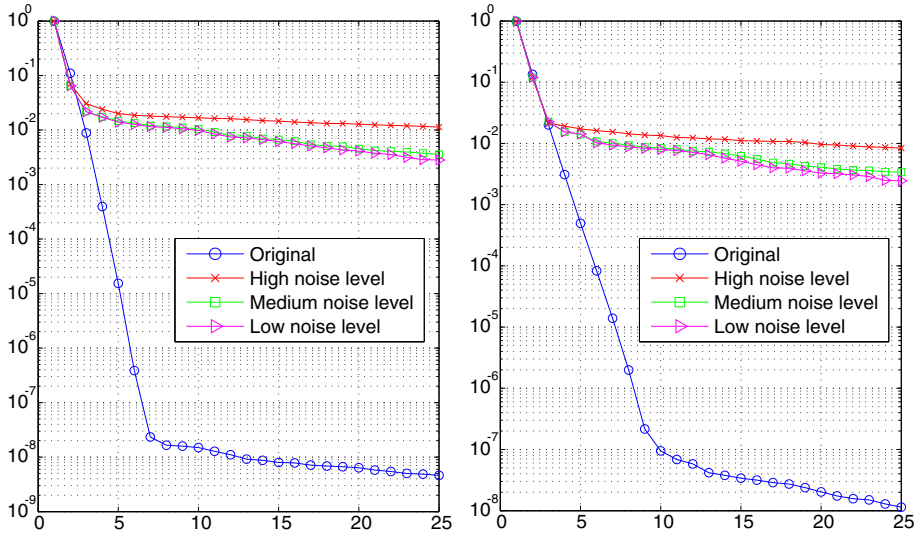

Figure 3 First set of eigenvalues in the PSF field of BLUR1 and BLUR2. Original and estimated for 3 different Poisson noise levels (density: 40.000 stars).

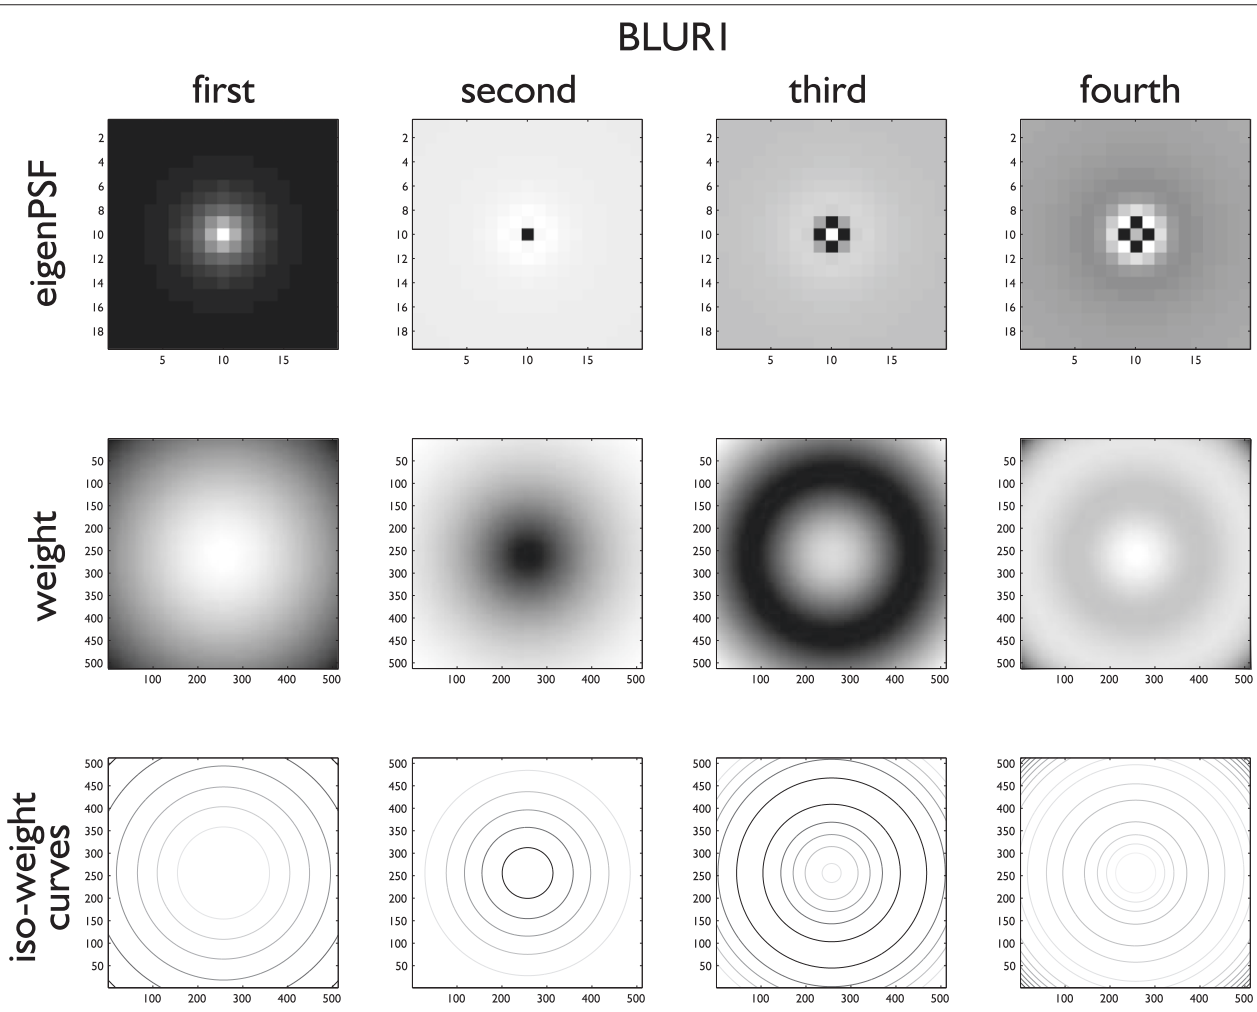

Figure 4 Original first BLUR1 *eigenPSFs* and their associated weights.

the number of stars with apparent magnitude between m and $m+1$ (at least for the studied interval $0 \leq m \leq 10$) is around three times larger than the number of stars with apparent magnitude between $m-1$ and m. Considering that an increase by one unit of magnitude corresponds to a decrease by 2.15 times the apparent brightness, this yields $\alpha = \log(3)/\log(2.15) \approx 1.43$. For simulating apparent brightness samples, we use the standard technique of integrating and inverting the pdf (a power law in this case[f]) with this coefficient, and then sampling from an uniform density in $[\epsilon, 1]$ and applying the obtained function to the uniform samples. In this case, the resulting function is still a power law with $\beta = -1/(\alpha - 1) \approx -2.3$. This yields an output brightness range of $[1, \epsilon^{-2.3}]$, which is normalized afterwards to the desired maximum value. We have chosen $\epsilon = 0.01$.

We have implemented two synthetic PSF fields, BLUR1 and BLUR2, which are quite similar to the ones used in Part I. BLUR1 is defined by:

$$h_{\mathrm{BLUR1}}(\mathbf{x};\mathbf{x}^0) \propto 1/(k(x^0,y^0)^2 + x^2 + y^2) - 0.95 \min\{1/(k(x^0,y^0)^2 + x^2 + y^2)\},$$

with x and y ranging now from -9 to 9. Except for keeping a constant spatial support and the offset correction, this PSF field can approximately be interpreted as a spatial magnification of the PSF from the center (where it is sharpest), with the zoom factor being $k(x^0,y^0) = 2^{-1/2}\sqrt{1 + ((x^0)^2 + (y^0)^2)/(L^2/6)}$, being L the side of the pattern (assumed squared). The second PSF field, BLUR2, is a Gaussian function whose width (square root of variance along the horizontal axis) is fixed and equal to $\sigma_x = 1.6$, whereas its height changes exponentially along the vertical axis by $\sigma_y(x^0,y^0) = 1.6 \times 2^{y^0 - L/2}$. Note that these two PSF fields are defined through uni-parametric

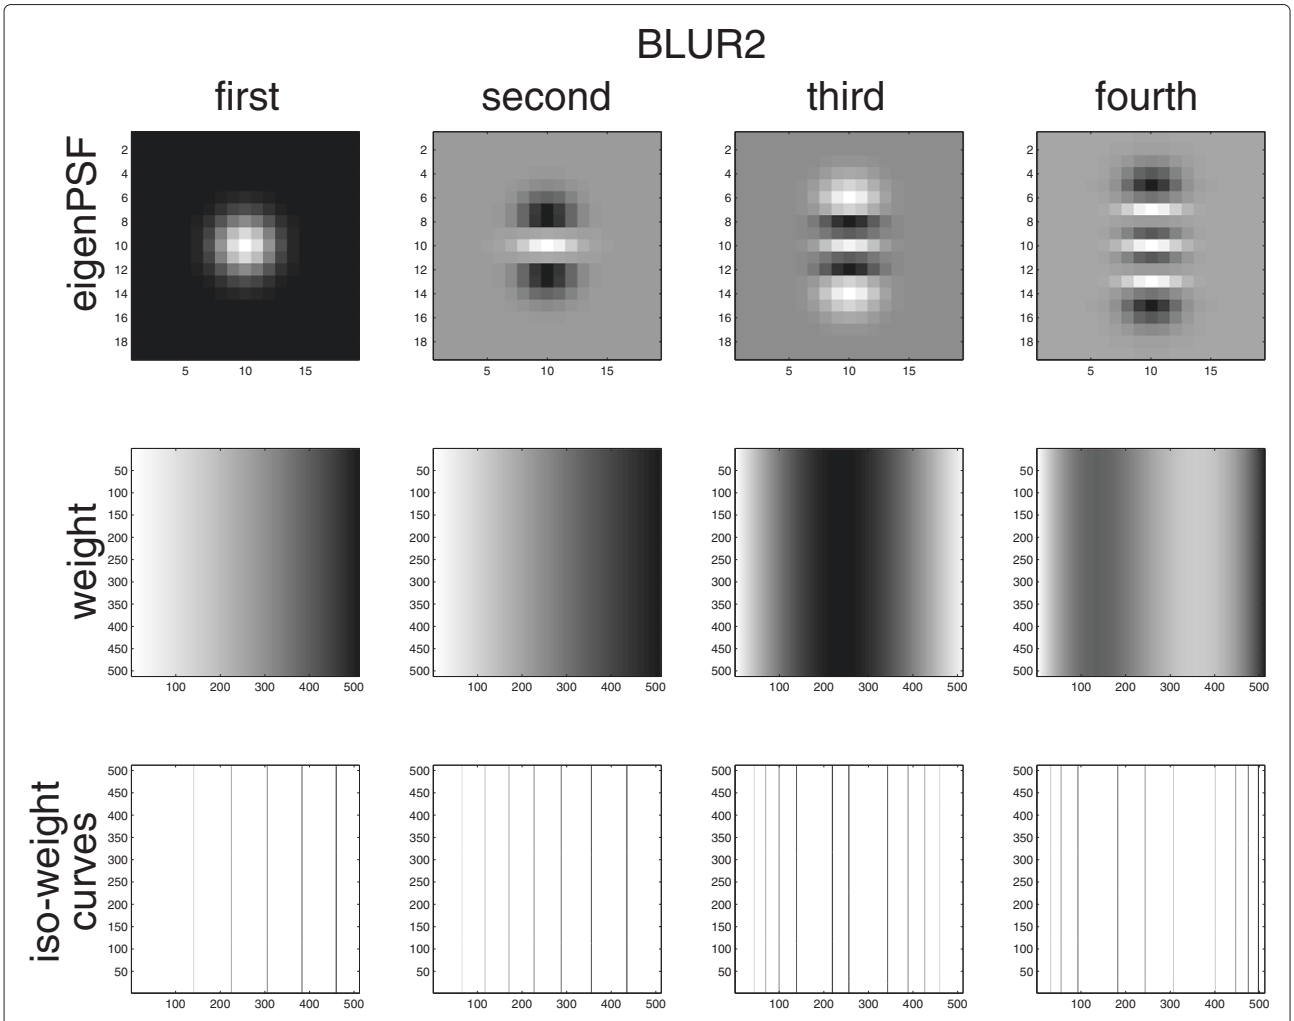

Figure 5 Original first BLUR2 *eigenPSFs* and their associated weights.

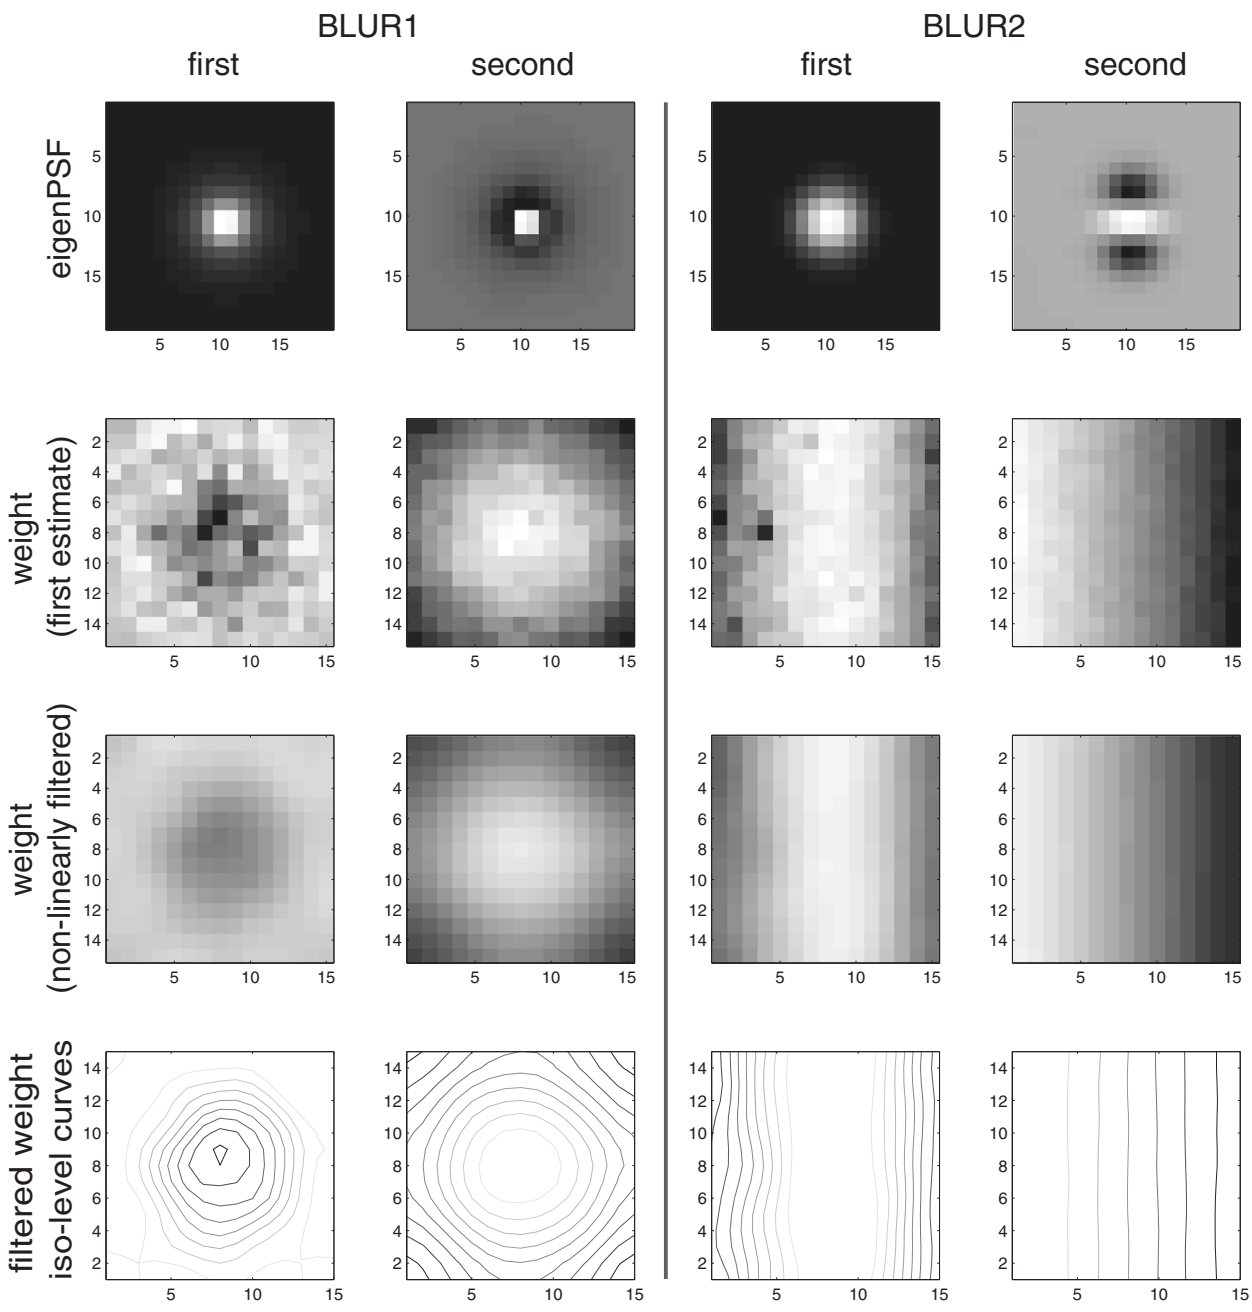

Figure 6 Estimated first BLUR1 and BLUR2 eigenPSFs from simulations with low Poisson noise level (maximum level of 65535, 16 bits), and their associated weights over the image.

PSF functions, with the parameter being a smooth function of the image coordinates. Thus, they are smooth uni-parametric PSF fields.

A degraded star field simulation is obtained by (1) choosing an image size (2048 × 2048 in our case) and a number of stars for the image ($2 \times 10^4, 4 \times 10^4, 1.6 \times 10^5$ or 6.4×10^5), locating each at non-integer random coordinates, and then associating to each of them a random brightness according to the power law explained above; (2) applying the synthetic PSF field (BLUR1 or BLUR2); (3) normalizing the resulting maximum to the desired value (255, 4095, or 65535, for 8, 12, or 16 bits, respectively), (4) simulating Poisson statistics on the previously noise-free image, and (5) adding white Gaussian noise (zero mean, unit variance), rounded to integer values (plus a constant, to avoid negative values).

Table 1 Accuracy of the PSF field estimation (Signal-to-noise ratio, in decibels)

	BLUR1				BLUR2			
Poisson noise level	**LD**	**MD**	**HD**	**UHD**	**LD**	**MD**	**HD**	**UHD**
8 bits	6.95	7.02	7.12	6.64	10.01	9.68	10.15	8.90
12 bits	7.72	6.99	6.88	6.63	10.35	10.20	10.16	8.79
16 bits	7.48	7.14	7.16	6.76	10.33	10.38	10.44	9.47

LD has used a pooling region of 512^2 pixels, MD and HD 256^2 pixels, and UHD 128^2 pixels.

5.1.2 Results for simulated blurred star fields

Figure 3 shows the eigenvalues profile of the two synthetic PSF fields (BLUR1 and BLUR2), and compares them to the profiles obtained by using our method on the grid of initial PSF estimates to the simulated images, for the three noise levels considered (for comparison purposes, global energy has been normalized in all cases). It is very interesting to see: (1) the rapid decay of the ranked eigenvalues; (2) the estimated profiles fitting the theoretical curves for the first two or three eigenvalues, the rest staying effectively beneath observation noise; (3) the influence of different Poisson count levels being quite small for the chosen rank (8, 12, and 16 bits). Observing these eigenvalues profiles, it seems legitimate to use just two dimensions for the reconstruction, in both blur cases (BLUR1 and BLUR2).

Figures 4 and 5 show the structures obtained by applying the SVD to the original PSF fields BLUR1 and BLUR2, respectively. It is worth noting that the weights associated to the extracted *eigenPSFs* for each PSF field all have the same iso-level structure (circular for BLUR1, and vertical for BLUR2), indicating a PSF field with a single degree of freedom, as expected. In addition, weighting functions are smooth, as expected for smooth PSF fields. The basic shape of the first two *eigenPSFs* and their iso-level structures are quite well reproduced in the PSF field estimation from the simulated observations, as shown in Figure 6.

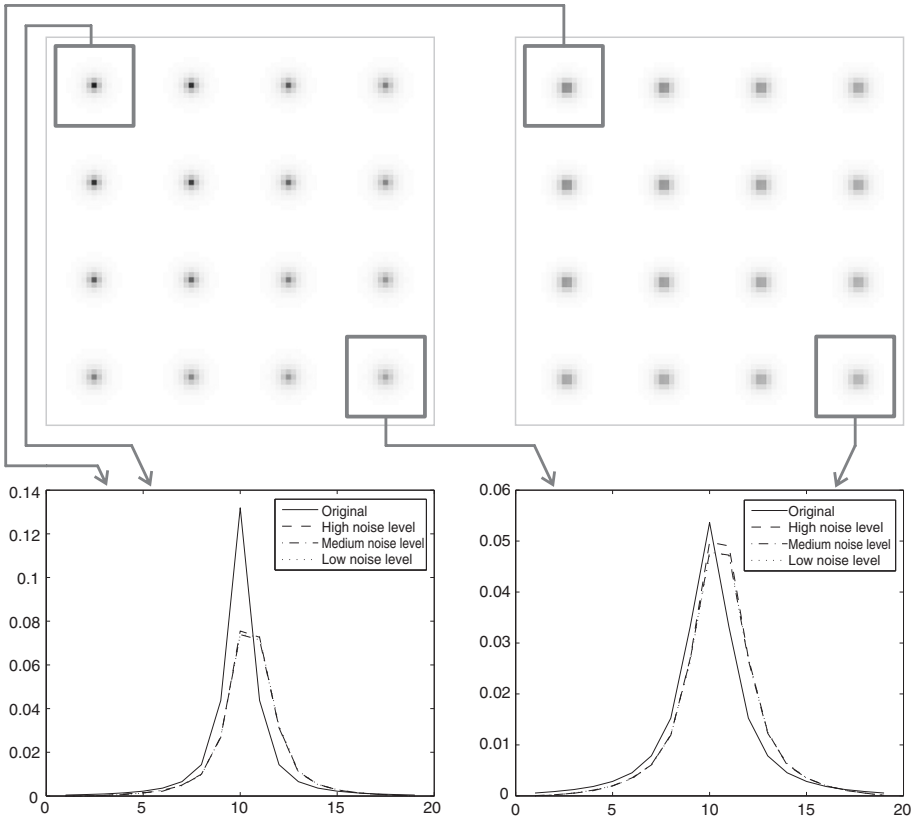

Figure 7 Sampled PSFs with pair indexes from the lower right quadrant of BLUR1 PSF field with medium star density (40,000 stars, pooling region of 256×256 pixels with half block overlapping). Top left: Original PSFs; top right: estimated PSFs in a 12-bit image (medium noise level). In the bottom, both linear plots show a central cross section of selected PSFs for all considered noise levels.

Here, one can fully appreciate the importance of non-linearly removing the outliers, i.e., local weight values departing from the generally spatially smooth behavior.

Table 1 shows a quantitative comparison of the estimated PSF field to the original. It shows Signal-to-noise ratio values, in decibels, of PSF field estimation in BLUR 1 and BLUR 2 with 4 Poisson noise levels and 3 different star densities: low density LD (20,000 simulated stars), medium density MD (40,000 simulated stars), high density HD (160,000 stars), and ultra high density UHD (640,000 stars). While the obtained figures are relatively modest, it is worth noting the high robustness of the method against noise and changes of the star density. We only notice a slight performance decrease when the density becomes too high, which causes a high amount of PSF overlap. Figures 7, 8 and 9 further illustrate this robustness, by comparing the obtained results to samples of the original PSF fields, for different Poisson levels and star densities. They also indicate that the blur is slightly overestimated. However, more importantly, a big part of the error comes from obtaining kernels with different relative locations. This is not strictly an error in the estimate, as there is an intrinsic ambiguity for the pattern location in terms of its internal coordinates $\boldsymbol{x}$ in the model, but it still affects the numerical results negatively.[g] Disregarding this spurious shift effect, it seems that the most important source of error in these results comes from the *blurring* of the local pattern field caused by obtaining our local PSF estimates on relative large blocks. More sophisticated choices for the shapes of the spatial regions upon which making our local estimates, instead of square blocks, should improve the classical bias-variance trade-off (bigger regions produce higher bias and lower variance, and vice versa) present at any estimation problem on non-stationary random fields, as pointed out in the conclusions.

5.2 PSF field estimation in real astronomical images

We apply our algorithm to three raw images, which were taken with the historic Cassegrain 0.9-m (36-inch) telescope at the National Science Foundation's Kitt Peak National Observatory: Crescent Nebula (Caldwell27), Swan Nebula (M17), and Ring Nebula (M57). All these images are coded in FITS format. They are publicly available for academic and research purposes in the

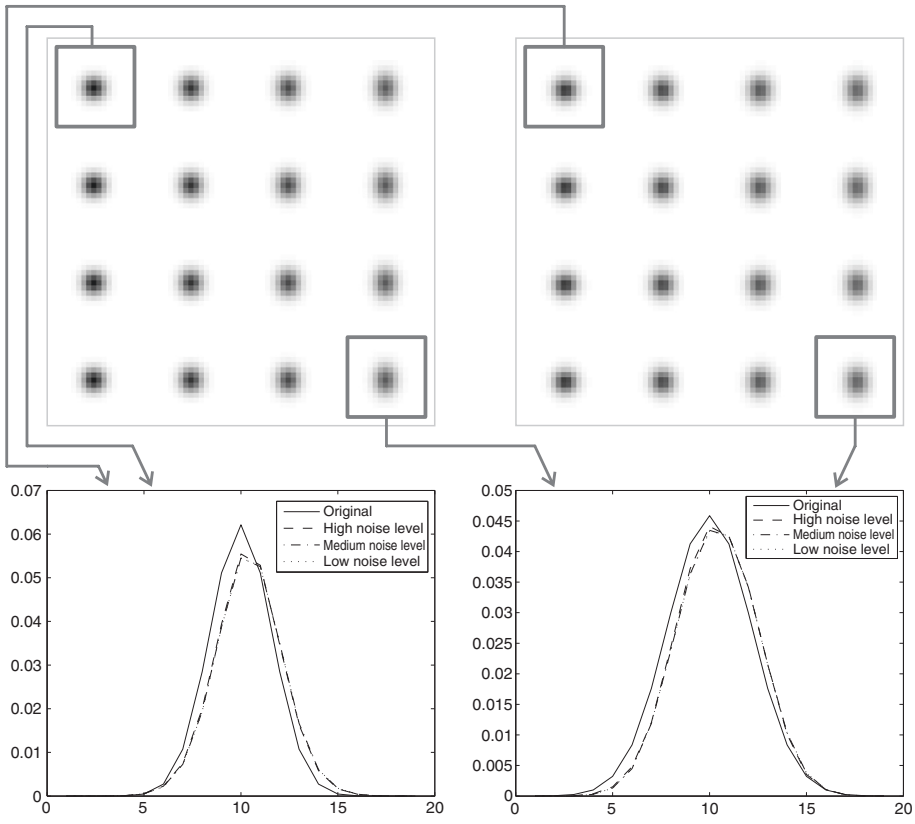

Figure 8 Sampled PSFs with pair indexes from the lower right quadrant of BLUR2 PSF field with medium star density (40,000 stars, pooling region of 256×256 pixels with half block overlapping). Top left: Original PSFs; top right: estimated PSFs in a 12-bit image (medium noise level). In the bottom, both linear plots show a central cross section of selected PSFs for all considered noise levels.

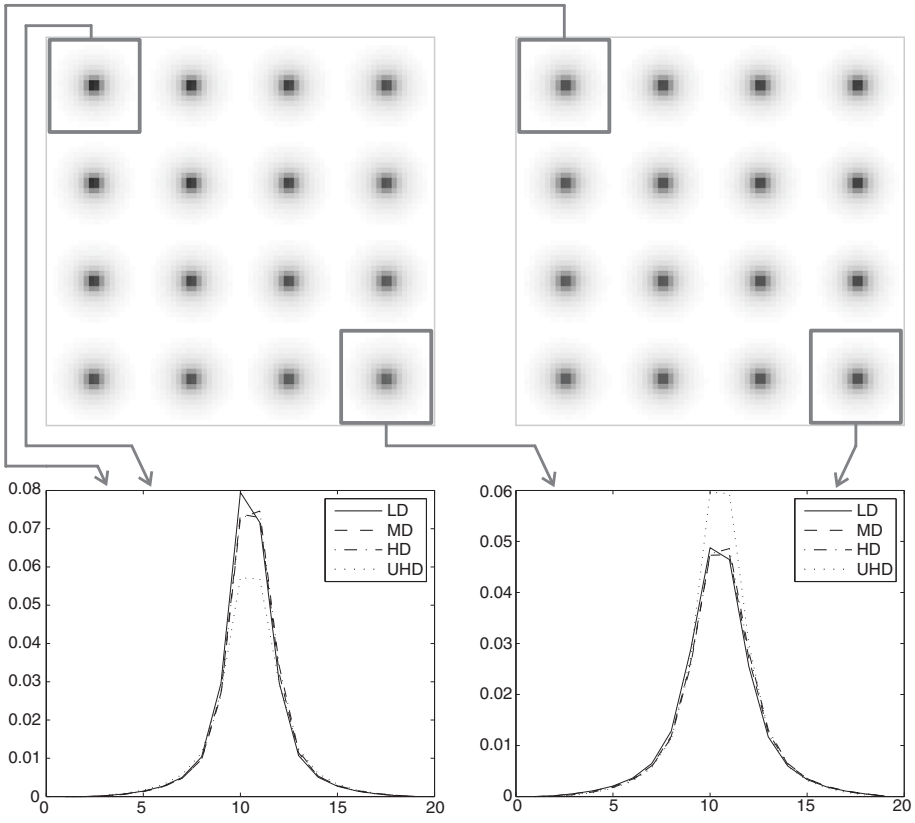

Figure 9 Sampled PSFs with pair indexes from the lower right quadrant of BLUR1 PSF field with low level of noise and different star densities. Top left: medium density (40,000 stars) with pooling region of 256 × 256 pixels; top right: high density with smaller pooling region (128 × 128 pixels). In the bottom, both linear plots show a central cross section of selected PSFs for all considered density studied cases.

observatory educational website http://www.noao.edu/education/arbse/arpd/ia.

The apparent magnitude scale, which we apply in the previous simulations, is adapted to the sensitivity of the human visual system. For these experiments, we have used V-filtered raw images, because the formal photoelectric V peak (around 540 nm) is closest to the peak of the human eye's (dark-adapted) detection efficiency. The images have not been preprocessed in order to show the robustness of the proposed algorithm.

The images are star fields of 2048 × 2048 pixels, and have ranges approximately corresponding to 16 bits. Figure 10 shows the eigenvalue profiles for the initial grid of estimated PSF fields, which advise us to use three significant (above noise level) components for the Caldwell27 image, and only two for M17 and M57. In Figure 11 we show a 15 × 15 PSF panel summarizing the result of the PSF field estimation, for Caldwell27. We observe that the PSF varies smoothly across the image, becoming more concentrated (less blurred) as we get closer to the bottom right corner. By looking at the three significant *eigenPSFs* (Figure 12, top left) we see how they are modulating the spatial extension of the PSF, which basically follows an off-symmetric close-to-circular pattern. There is something remarkable about the behavior of the associated weights: to a first approximation, they

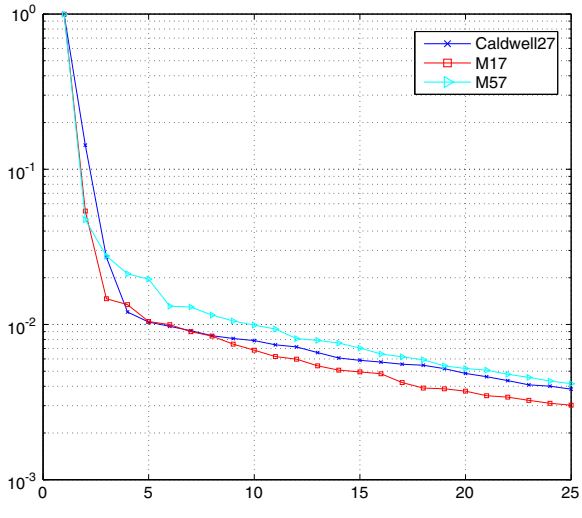

Figure 10 First set of eigenvalues in real astronomical image cases under study: Caldwell27, M17 and M57.

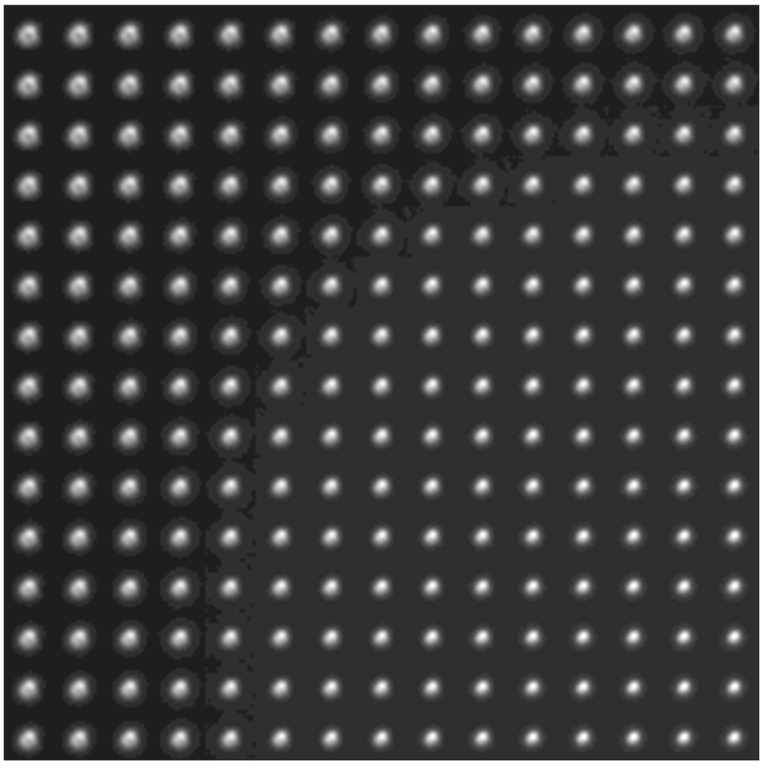

Figure 11 Estimated PSF field for Caldwell27, sampled in a 15×15 grid.

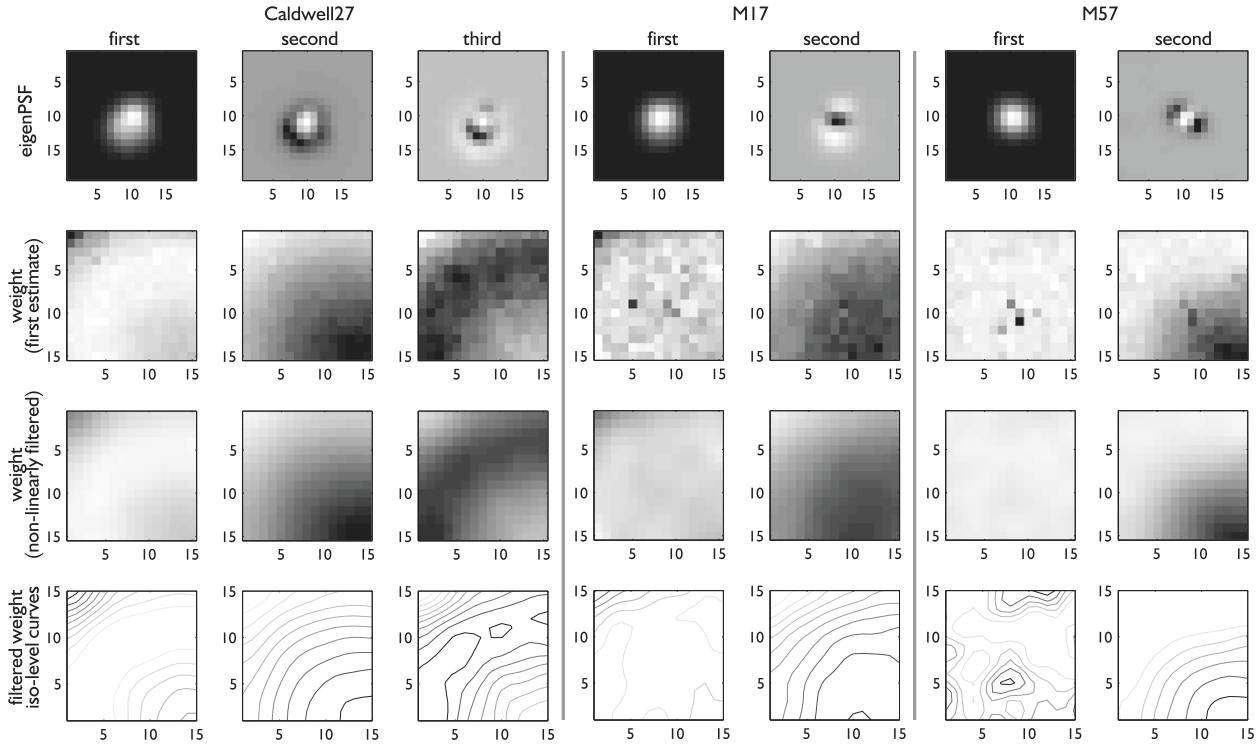

Figure 12 Estimated first *eigenPSFs* from real astronomical image cases under study: Caldwell27, M17 and M57, and their associated weights over the image.

share the same iso-level curves. This means that, in this real example, the modeled PSF field is not only smooth, but also (approximately) uni-parametric. This surprising result is further illustrated in Figure 13, where the three computed weights have been visualized as dots in a 3D space. The result clearly shows a one dimensional structure (a curve), reflecting a single effective degree of freedom in the PSF field.

Interestingly enough, the other two studied images have rather different PSF fields, which translate into different *eigenPSF* shapes, especially for the second component (see Figure 12, top, center, and right). Nevertheless, the processed weights follow remarkably the same approximated pattern of iso-level weights for this significant second component. Note that the lack of a clear iso-level structure of the weights associated to the first *eigenPSF* is not significant in those examples, as the first component has a fairly constant weight all over the image. Such a pattern appears in other images (data not shown here) from the same telescope. This pattern is not just due to an artifact of the estimation method. We know that, because we have seen a radically different (and, by comparison to the known reference, correct) behavior in the simulations. Our provisional explanation is that such iso-level structure is related to the approximately fixed PSF field component caused by the telescope structure (which uses an off-center sensor with a parabolic mirror, in this case), whereas significant differences in the overall PSF field would be caused by the atmospheric PSF at the moment and sky location of the capture. The observed PSF field can be modeled as the result of convolving atmospheric PSF field (approximately spatially invariant, for the small angle subtended, but variable for each image capture) and the telescope PSF field component (strongly spatially variant, in this case, but approximately constant for each image on the set). However, we do not have a solid explanation yet for the approximately uni-parametric behavior of the PSF fields.

5.2.1 *Image blind restoration in real astronomical images*

Finally, we apply the SV restoration method from Part I using the estimated PSF fields to compensate for the blur. Because the dominant noise source in these images is not Gaussian, but Poissonian, we use a variant of the L0-AbS method [22]. Regarding all other aspects, the image estimation method and its implementation is the same as the one explained in Part I. The resulting images gain a lot of sharpness, which allows to resolve very close stars that appear as a single cloud in the observations (as shown in Figure 14). The spatial "averaging" effect due to the block-wise estimation of the local PSFs may cause a slight overestimation of the blur, which also appears in the simulation results. That would explain some over-sharpening effects such as the low-contrast halo shown in the bottom right subfigure.

6 Conclusions and future work

We propose an efficient and robust method for PSF field estimation on star fields. The method is based on careful detection and sub-pixel alignment of the stars

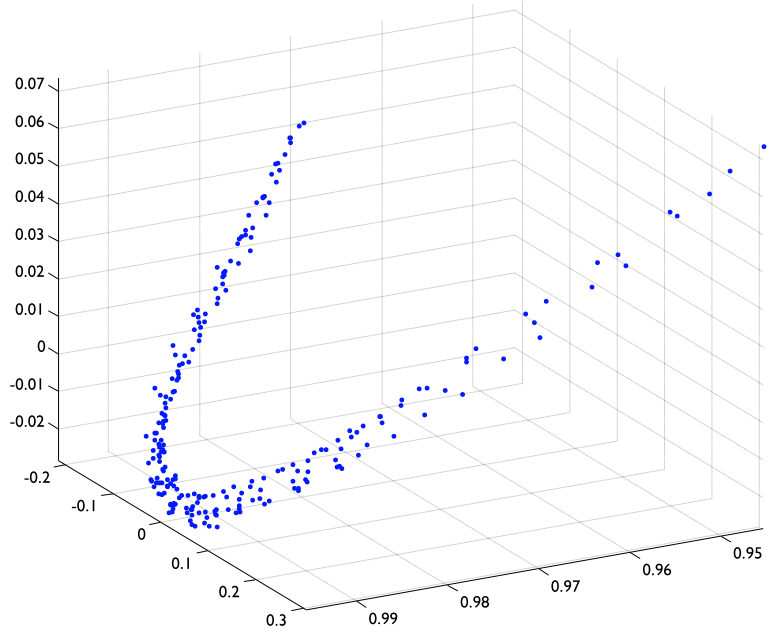

Figure 13 Estimated weights of the three main *eigenPSFs* computed for Caldwell27, composing a one dimensional structure in the space (225 dots coming from sampling in a 15×15 grid).

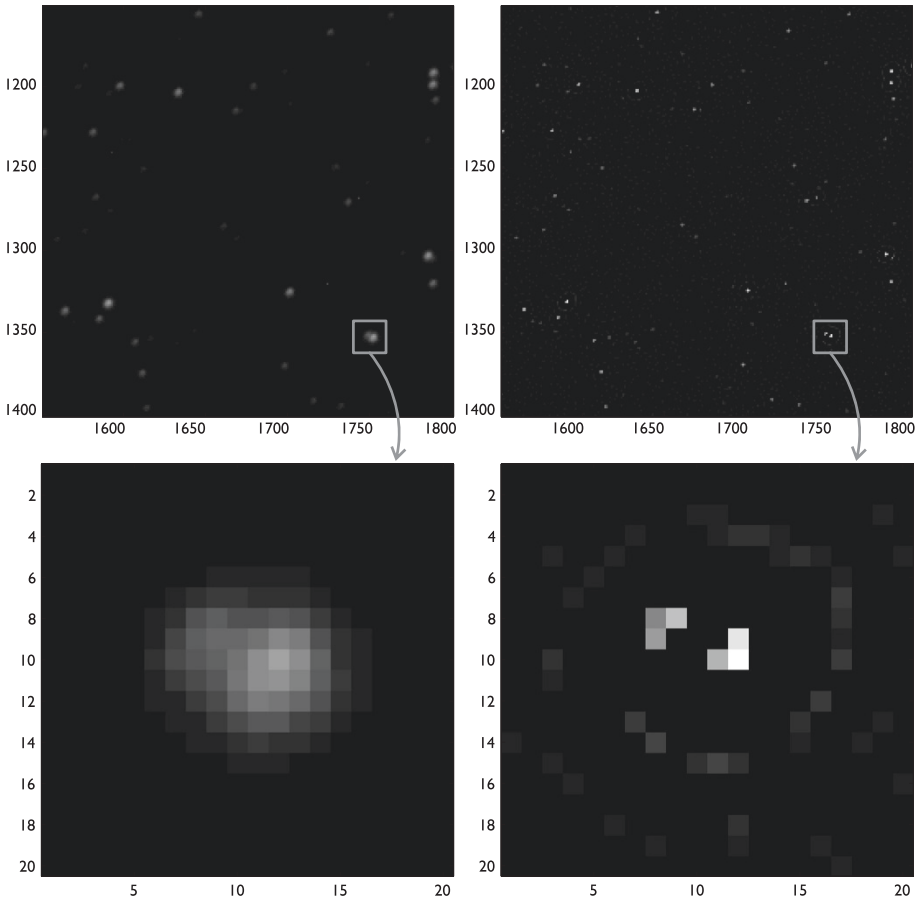

Figure 14 Original (left) and restored (right) portions of Caldwell27 case study image. One detail has been enlarged to show how close astronomical objects are resolved by the algorithm. Gray-level scale in the image has been nonlinearly modified because of the large dynamic range.

according to a simple statistical local model and the ML principle, followed by a combination of dimensionality reduction and nonlinear filtering of outliers on the weights of the *eigenPSFs*. We have demonstrated that the proposed method is able to robustly capture the structure of the spatially variant blur of smooth PSF fields, in spite of the simplifying assumption that the local PSFs do not overlap. This is demonstrated by image and PSF field simulations, using different Poisson counts and star densities.

Beyond the introduction of the new local model, the article offers interesting results obtained from applying the method to real images (from a medium quality, Earth-based, telescope image dataset), for which the method has revealed significant information about the PSF field structure. We have also demonstrated the applicability of this technique in conjunction with the SV restoration described in our companion paper (Part I), to jointly build a blind SV restoration method assuming Poisson statistics.

We believe that the repeated-pattern local image model presented here has a big potential to be extended and improved in several directions, even for tasks different than PSF field estimation (e.g., texture characterization, or other forms of intelligent combinations of image analysis, processing and synthesis). Perhaps the most promising and natural among these possibilities could come from formalizing the pattern explicitly as a low-dimensional *deformable kernel* in a global image model, instead of making the local approximation of considering it constant on certain regions, and then building a smooth field from the resulting local estimates, as we did here. It may be both very interesting and potentially useful to explore, at least for the cases where a simple and robust structure of the PSF field is detected (like the approximately uni-parametric behavior observed in the real astronomical images studied here), alternative estimation strategies for exploiting the spatial structure of this non-local redundancy. This should result in reducing the bias of the estimates without increasing the associated variance. Finally, another natural evolution of this type of models could arrive by substituting the useful, but limited concept of *linear* dimensionality reduction by a much

more powerful (and also more difficult to deal with) *nonlinear* dimensionality reduction concept (see, e.g., [23]) which aims at modeling the underlying manifold associated to a varying pattern, low dimensional only in a local sense. Evidence of low-dimensional curved manifolds, like the one shown in Figure 13 are very motivating to address this conceptual and technical challenge for practical purposes.

Endnotes

[a]To jointly attack the deformable PSF field estimation requires, even in the favorable case of star fields, a substantially more sophisticated approach. We are currently working on that, given the applicability of such an approach to a wider range of images and estimation problems.

[b]If the neighborhood size is chosen appropriately, it is a safe assumption that $p(\boldsymbol{x})$ goes to zero at the neighborhood boundaries (μ_p is the offset to this "ground level"). Consequently, it is easy to estimate μ_p from $\boldsymbol{p}_0$ and revert this decomposition, even though μ_p is lost in $\boldsymbol{p}_0$.

[c]Although the sample mean may depart significantly from the background level when there is a star in the neighborhood, we have tested that this effect, in this case, is small compared to that of the variance in the denominator.

[d]The interpolation is always computed on the original full-pixel samples to avoid accumulating error.

[e]As in the previous section, we are keeping the PSF patterns zero-mean and L_2-normalized. We, as other authors (see, e.g., [15]) have experienced that mean subtraction followed by L_2 normalization generally helps to produce a higher spectral concentration as compared to leaving the PSFs positive and normalized in volume.

[f]Although this is an improper pdf, as it is not integrable around zero, this does not prevent us from using it for this purpose.

[g]We intend to remove this ambiguity in future versions of this method.

Competing interests
The authors declare that they have no competing interests.

Acknowledgements
Astronomical real images used to illustrate the algorithm (Caldwell27, M17, and M57) are copyrighted by National Optical Astronomy Observatory/Association of Universities for Research in Astronomy/National Science Foundation. They have been used strictly for educational and research purposes, as NOAO/AURA Image Library Conditions of Use states. We also thank the anonymous reviewers, for their meaningful comments and constructive suggestions. David Miraut has been funded by grant CICYT TIN2010-21289-C02-01. Javier Portilla has been funded by grants CICYT TEC2009-13696 and CSIC PIE201050E021. All grants from the Spanish Education, and Science&Innovation Ministries. Johannes Ballé stay in Madrid has been partially funded by the CSIC.

Author details
[1]E.T.S. Ingenieria Informatica, Universidad Rey Juan Carlos, Madrid, Spain.
[2]Institut für Nachrichtentechnik, RWTH Aachen University, Aachen, Germany.
[3]Instituto de Óptica, Consejo Superior de Investigaciones Científicas, Madrid, Spain.

Received: 5 January 2012 Accepted: 16 July 2012
Published: 31 August 20112

References

1. MSC Almeida, LB Almeida, Blind and semi-blind deblurring of natural images, IEEE Trans. Image Process. **19**, 36–52 (2010)
2. A Gupta, N Joshi, C Lawrence Zitnick, M Cohen, B Curless, in *Computer Vision–ECCV 2010* Single image deblurring using motion density functions. (2010), pp. 171–184
3. YW Tai, H Du, MS Brown, S Lin, Correction of spatially varying image and video motion blur using a hybrid camera, IEEE Trans. Pattern Anal. Mach. Intell. **32**(6), 1012–1028 (2010)
4. N Joshi, R Szeliski, DJ Kriegman, in *IEEE Conference on Computer Vision and Pattern Recognition, 2008. CVPR 2008* PSF estimation using sharp edge prediction. (IEEE, 2008), pp. 1–8
5. E Kee, S Paris, S Chen, J Wang, in *IEEE International Conference on Computational Photography (ICCP 2011)* Modeling and removing spatially-varying optical blur. (IEEE, 2011), pp. 1–8
6. Q Shan, B Curless, T Kohno, in ed. by K Daniilidis, P Maragos, N Paragios, *Computer Vision - ECCV 2010,* vol. 6316 of *Lecture Notes in Computer Science* Seeing through obscure glass, (Springer, Berlin, 2010), pp. 364–378
7. GM Bernstein, M Jarvis, Shapes and shears, stars and smears: optimal measurements for weak lensing, Astronom. J. **123**(2), 583 (2002)
8. A Refregier, Shapelets—I. A method for image analysis, Month. Notices R. Astronom. Soc. **338**(1), 35–47 (2003)
9. L Denis, E Thiebaut, F Soulez, in *2011 18th IEEE International Conference on Image Processing (ICIP)* Fast model of space-variant blurring and its application to deconvolution in astronomy. (September 2011), pp. 2817–2820
10. MJ Jee, HC Ford, GD Illingworth, RL White, TJ Broadhurst, DA Coe, GR Meurer, A Van Der Wel, N Benítez, JP Blakeslee, et al., Discovery of a ringlike dark matter structure in the core of the galaxy cluster Cl 0024+ 17, Astrophys. J. **661**, 728 (2007)
11. P Perona, Deformable kernels for early vision, IEEE Trans. Pattern Anal. Mach. Intell. **17**(5), 488–499 (1995)
12. GH Golub, M Heath, G Wahba, Generalized cross-validation as a method for choosing a good ridge parameter, Technometrics. **21**, 215–223 (1979)
13. R Lupton, JE Gunn, Z Ivezic, GR Knapp, S Kent, in *Proceedings of Astronomical Data Analysis Software and Systems X,* vol. 238, 12–15 November 2000 The SDSS imaging pipelines. (Astronomical Society of the Pacific, Boston, MA, USA, 2001), pp. 269–279
14. MJ Jee, JP Blakeslee, M Sirianni, AR Martel, RL White, HC Ford, Principal component analysis of the time-and position-dependent point-spread function of the advanced camera for surveys, Publications Astronom. Soc. Pacific. **119**(862), 1403–1419 (2007)
15. MJ Jee, JA Tyson, Towards precision LSST weak-lensing measurement-I: impacts of atmospheric turbulence and optical aberration. arXiv:1011.1913v3 [astro-ph.IM], 2011
16. M Jarvis, B Jain, Principal component analysis of PSF variation in weak lensing surveys. arXiv:astro-ph/0412234v2, 2004
17. T Schrabback, J Hartlap, B Joachimi, M Kilbinger, P Simon, K Benabed, M Bradac, T Eifler, T Erben, CD Fassnacht, F William High, S Hilbert, H Hildebrandt, H Hoekstra, K Kuijken, P Marshall, Y Mellier, E Morganson, P Schneider, E Semboloni, L Van Waerbeke, M Velander, Evidence of the accelerated expansion of the universe from weak lensing tomography with COSMOS, Astronom. Astrophys. **516**, 1–26 (2010)
18. TR Lauer, Deconvolution with a spatially-variant PSF. arXiv:astro-ph/0208247v1, 2002
19. M Elad, MAT Figueiredo, Y Ma, On the role of sparse and redundant representations in image processing, IEEE Proc. (Special Issue on Applications of Sparse Representation and Compressive Sensing). **98**(6), 972–982 (2010)
20. DL Donoho, IM Johnstone, Ideal spatial adaptation by wavelet shrinkage, Biometrika. **81**, 425–455 (1994)
21. L Strous, Mr. Sunspot's Answer Book. National Solar Observatory/Sacramento Peak, Sunspot, NM, USA (data available through

the Wikipedia article: "Apparent magnitude", and in the web archive of the original source), 1996
22. E Gil-Rodrigo, J Portilla, D Miraut, R Suarez-Mesa, in *18th IEEE International Conference on Image Processing (ICIP)* Efficient joint Poisson-Gauss restoration using multi-frame L2-relaxed-L0 analysis-based sparsity. (September 2011), pp. 1385–1388
23. JA Lee, M Verleysen, *Nonlinear Dimensionality Reduction* (Series: *Information Science and Statistics*), Springer, New York, 2007

doi:10.1186/1687-6180-2012-193
Cite this article as: Miraut *et al.*: Efficient shift-variant image restoration using deformable filtering (Part II): PSF field estimation. *EURASIP Journal on Advances in Signal Processing* 2012 **2012**:193.