

Online Blind Deconvolution for Astronomical Imaging

Stefan Harmeling, Michael Hirsch, Suvrit Sra, and Bernhard Schölkopf
Max Planck Institute for Biological Cybernetics

firstname.lastname@tuebingen.mpg.de

Abstract

Atmospheric turbulences blur astronomical images taken by earth-based telescopes. Taking many short-time exposures in such a situation provides noisy images of the same object, where each noisy image has a different blur. Commonly astronomers apply a technique called “Lucky Imaging” that selects a few of the recorded frames that fulfill certain criteria, such as reaching a certain peak intensity (“Strehl ratio”). The selected frames are then averaged to obtain a better image. In this paper we introduce and analyze a new method that exploits all the frames and generates an improved image in an online fashion. Our initial experiments with controlled artificial data and real-world astronomical datasets yields promising results.

1. Introduction

When we observe an astronomical object from the surface of the earth its emitted or reflected light has to invariably pass through the atmosphere, which yields a blurry observed image. Deblurring images of an observed celestial body is therefore a fundamental problem in Astronomy. This problem is compounded by the fact that the blur is not only unknown, but is also continually changing in time as well as spatially due to refraction-index fluctuations caused by atmospheric turbulence.

It is well known, that exposure times on a time scale where the turbulences are stationary (i.e. shorter than the tenth of a second) yield images that contain high-frequency information of the celestial object under investigation [5]. This fact is exploited in Speckle Imaging, which is a collection of various techniques and algorithms for recovering high frequency information encapsulated in short-time exposure images. Due to the stochastic nature of atmospheric turbulence, Speckle Imaging techniques have to take a large number of images into account to actually regain diffraction-limited information.

In Speckle Interferometric Imaging short-exposure pictures are processed in Fourier domain. Labeyrie [10] showed that diffraction-limited information could be re-

trieved by averaging the auto-correlation function. Two years later, Knox and Thompson presented an algorithm for phase reconstruction in [8]. For averaging usually several thousand images are taken into account.

Through new developments in CCD technology which provide superior signal-to-noise ratio (Electron Multiplying CCD cameras), in recent years so-called Lucky Imaging methods have become popular. Usually a large number of images (> 10000) is collected, of which only a very small percentage is utilised to yield one high quality image. Different registration and super-resolution methods have been proposed for combination of the selected images [6, 18].

To our knowledge, up to now all proposed approaches that address the problem of high frequency information retrieval from a sequence of images are post-processing (or batch-mode) methods. In this paper we introduce an efficient online scheme that models the multi-frame problem more realistically. We assume that the images are obtained in a streaming fashion, and that they are potentially too numerous to save or process efficiently, even if we could save all of them. While this may appear artificial at first glance, the rate of data acquisition in astronomy has been increasing tremendously over the years, to the point where data recorded in current sky surveys will take take years to be analyzed.

Several different online methods could be designed to model and tackle this multi-frame setting, and we introduce a particular one. Our approach to solving the problem has the following main advantages:

- It has low resource requirements because the images are tackled in a streaming fashion; this obviates the need to store the images, and results in a more efficient algorithm (in terms of run-time and memory usage).
- Typical online methods usually require the setting of a learning rate. Our method is based on multiplicative updates to the image estimate that effectively bypasses the need to set a learning rate, which simplifies the algorithm and its implementation.
- Our method is very simple to implement, and it can be implemented in the imaging hardware itself to provide dynamic deblurring, if desired.

Before discussing the formal problem setting and our algorithm (Sections 2 and 3) we first provide a short summary of additional related work in Section 1.1 below. Later, in Section 4 we provide several experiments on both artificial and real data to illustrate the benefits of our method. Our initial results are encouraging and point to several avenues of future work, which we summarize in Section 5.

1.1. Related Work

The general problem of image restoration in astronomy has a long history; we refer the reader to a recent survey for an introduction and useful references [11]. We limit our discussion below only to work related to the multi-frame deblurring problem, as that is the focus of our paper.

All the standard multi-frame methods known to us operate in a batch-mode. For example, [2] used an ordered-subset/expectation maximization (OS-EM) algorithm that is a block-iterative version of the highly popular Richardson-Lucy algorithm [14] applied to all input images simultaneously. Several other authors have also considered this multi-frame deblurring problem for Astronomy [16, 20, 21, 1], all of them functioning in a batch-mode.

Vio et al. [20] compute an approximate single image representing (in a least-squares sense) all the input images and then work on the resultant image only. Their choice has two main problems: first, it requires all the input image simultaneously, and second, the least-squares image can have negative entries. In contrast, our method deals with input images one by one, and non-negativity is enforced naturally without any *ad-hoc* procedures. Furthermore, our method does not assume that all the observed images are polluted by the same noise—an assumption fundamental to the single-image method of [20].

Intimately related to our problem is the general blind-deconvolution problem for images [9], where one usually has just one or two images [13]. If one observes an entire sequence of differently blurred images of the same underlying image, one could hope to do better deblurring.

Algorithmically, we perform a sequence of quadratic loss minimizations, motivated by two important reasons, viz., algorithmic simplicity and the observation that the least-squares objective yields more robust results than the traditional Poisson model for this multi-image problem. We observed the latter claim to hold true in our experiments, a fact further corroborated by [19]. However, our method can be easily adapted to use a Poisson model if needed (we will briefly remark on this later). We would also like to note that even though there exist several papers dealing with online convex optimization [7, 22], it seems difficult to leverage their theoretical guarantees as our problem is non-convex.

2. Problem Formulation

In this section we describe the problem setting (§2.1) and introduce a model formulating it (§2.2).

2.1. The setting

We assume that at each time point $t = 1, 2, \dots$ our telescope records a blurry image \mathbf{y}_t that is obtained from a common underlying image \mathbf{x} by blurring it with some unknown blur kernel \mathbf{f}_t and adding noise \mathbf{n}_t . The blurred images $\{\mathbf{y}_1, \mathbf{y}_2, \dots\}$ arrive in a streaming fashion, and to model reality we assume that we have only a limited buffer of size W (*window-size*) for saving them. A method is *online* if it has access to only one image at a time, i.e., $W = 1$, while if the image stream is terminated and all the observed (say T) images are available, it is called a *batch* method ($W = T$). Naturally, when $1 < W < T$ one has a *hybrid* procedure.

The aim of blind deconvolution for this setting is to recover the true underlying image \mathbf{x} , given the input stream of blurred images $\{\mathbf{y}_t\}_{t \geq 1}$. Estimates for the individual blur kernels \mathbf{f}_t are obtained as a byproduct.

2.2. Formalization

For notational simplicity we describe our method in terms of one-dimensional images represented as column vectors. However, our algorithm and results can be easily reformulated in terms of two-dimensional images. Indeed, all our experiments are actually performed in the two-dimensional (matrix) setting.

Let each recorded image \mathbf{y}_t be a vector of length l_y , and the unknown true image \mathbf{x} have length l_x . The atmospheric blurring of \mathbf{x} is modelled as a *non-circular* convolution with a non-negative blur kernel (point spread function (PSF)) \mathbf{f}_t ; we denote this convolution by $\mathbf{f}_t \star \mathbf{x}$, and \mathbf{f}_t can be represented as a column vector of length $l_f = l_x - l_y + 1$. Note that for streaming data coming from a telescope, each observed image \mathbf{y}_t will be convolved by a different and unknown PSF \mathbf{f}_t .

The non-circular convolution can be written equivalently (with appropriate zero-padding) as a matrix-vector multiplication, i.e., $\mathbf{f}_t \star \mathbf{x} = \mathbf{F}_t \mathbf{x}$, in which the $l_y \times l_x$ matrix \mathbf{F}_t depends solely on \mathbf{f}_t . We will use \mathbf{f}_t and \mathbf{F}_t interchangeably in our discussion. Note that for an efficient implementation, the special form of \mathbf{F}_t allows for a fast computation of the matrix-vector product $\mathbf{F}_t \mathbf{x}$ via the Fast Fourier Transform (FFT). We denote individual components of the vector $\mathbf{F}_t \mathbf{x}$ by $(\mathbf{F}_t \mathbf{x})_j$. Often we will write fractions of vectors that should be understood as component-wise fractions.

2.2.1 Noise model and loss function

For high-intensity astronomy data one often models [19] the relationship between \mathbf{y}_t , \mathbf{f}_t , and \mathbf{x} as

$$\mathbf{y}_t = \mathbf{f}_t \star \mathbf{x} + \mathbf{n}_t, \quad (1)$$

where the noise \mathbf{n}_t is assumed to be Gaussian with mean $\mathbf{0}$ and diagonal covariance matrix $\sigma^2 \mathbf{I}$. Assuming a fixed but unknown σ one can deblur by minimizing the loss

$$L_t(\mathbf{x}, \mathbf{f}_t; \mathbf{y}_t) = \|\mathbf{y}_t - \mathbf{F}_t \mathbf{x}\|_2^2 \quad (2)$$

that must be minimized subject to non-negativity restrictions on \mathbf{x} and \mathbf{f}_t . For low intensity data, the Poisson model is considered more suitable,

$$\mathbf{y}_t \sim \text{Poisson}(\mathbf{f}_t \star \mathbf{x}), \quad (3)$$

though as noted by [19] and also observed in our experience, the least-squares model is more robust. Hence, we will focus our attention on (1) and (2) alone. However, we note that our online optimization scheme generalizes trivially to handle (3), in case it is really desired.

2.2.2 The optimization problem

If one had access to all the images $\{\mathbf{y}_t\}_{t \geq 1}$, then ideally one would like to minimize the total overall loss

$$\min_{\mathbf{x}, \mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_T} \sum_{t=1}^T L_t(\mathbf{x}, \mathbf{f}_t; \mathbf{y}_t), \quad (4)$$

where $L_t(\mathbf{x}, \mathbf{f}_t; \mathbf{y}_t)$ is the loss incurred at the t^{th} step as given by (2). However, in our problem setting (§2.1) two main concerns must be addressed:

1. We observe the images \mathbf{y}_t in a streaming fashion with a limited buffer of size W , while minimizing the overall loss (4) requires saving *all* the \mathbf{y}_t .
2. The overall loss (4) is *non-convex* due to the joint dependence between \mathbf{x} and \mathbf{f}_t via the convolution $\mathbf{f}_t \star \mathbf{x}$.

Below (§3) we develop an online algorithm that approximately tackles both these difficulties. We describe the method for the purely online case of $W = 1$, noting that it can be easily generalized to deal with the hybrid scenario with $W > 1$.

3. Algorithm

From each image \mathbf{y}_t that we observe we need to estimate the PSF \mathbf{f}_t , as well as the underlying image \mathbf{x} . Let \mathbf{x}_{t-1} be the current estimate of the true image. Now we obtain a new observation \mathbf{y}_t and we wish to update our estimate of the true image, while at the same time estimating the PSF \mathbf{f}_t . To this end, we can minimize the loss incurred at time-step t . Formally stated, we seek to solve the following non-negatively constrained problem

$$\min_{\mathbf{f}_t \geq 0, \mathbf{x} \geq 0} L_t(\mathbf{x}, \mathbf{f}_t; \mathbf{y}_t) = \|\mathbf{y}_t - \mathbf{F}_t \mathbf{x}\|_2^2, \quad (5)$$

to obtain the updated estimate \mathbf{x}_t and the PSF \mathbf{f}_t .

Thus, by solving only Problem (5) at each step t we obtain an online procedure that addresses the first concern mentioned in §2.2.2. However, the second issue still remains as even Problem (5) is non-convex, whereby in general a globally optimal solution for it cannot be found efficiently. Fortunately, the objective function is sufficiently well-behaved as it is individually convex in each variable if the other is held fixed. This crucial property lies at the heart of a simple alternating minimization or descent scheme for solving (5). Such a scheme constructs a sequence of iterates $\{\mathbf{x}_t^k, \mathbf{f}_t^k\}$ ensuring descent, i.e.,

$$L_t(\mathbf{x}_t^{k+1}, \mathbf{f}_t^{k+1}; \mathbf{y}_t) \leq L_t(\mathbf{x}_t^k, \mathbf{f}_t^k; \mathbf{y}_t), \quad (6)$$

and under some weak assumptions it can be shown [12] to converge to a stationary point of (5).

However, performing alternating descent too well comes at a price, namely the overfitting of \mathbf{y}_t . To avoid overfitting, one can perform just a few iterations of alternating descent, and as our experiments will show one or two iterations suffice for good results. The question to answer now is how to perform such iterations so that the descent condition (6) is satisfied. We propose the following two steps:

1. $\mathbf{f}_t = \operatorname{argmin}_{\mathbf{f} \geq 0} \|\mathbf{y}_t - \mathbf{F} \mathbf{x}_{t-1}\|^2$
2. $\mathbf{x}_t = \operatorname{argmin}_{\mathbf{x} \geq 0} g_t(\mathbf{x}, \mathbf{x}_{t-1})$,

where $g_t(\mathbf{x}; \tilde{\mathbf{x}})$ is an ‘‘auxiliary’’ function (similar to the one used by the ISRA algorithm [4]) that must fulfill

$$\forall \mathbf{x}, \tilde{\mathbf{x}} : g_t(\mathbf{x}; \tilde{\mathbf{x}}) \geq g_t(\mathbf{x}; \mathbf{x}) = L_t(\mathbf{x}, \mathbf{f}_t; \mathbf{y}_t).$$

This condition ensures that our choice of \mathbf{x}_t in Step 2 can only decrease our loss function since by definition

$$g_t(\mathbf{x}_{t-1}; \mathbf{x}_{t-1}) \geq g_t(\mathbf{x}_t, \mathbf{x}_{t-1}) \geq g_t(\mathbf{x}_t; \mathbf{x}_t)$$

It is easily seen that for the quadratic loss such a function is given by

$$g_t(\mathbf{x}, \tilde{\mathbf{x}}) = \mathbf{y}_t^T \mathbf{y}_t - 2\mathbf{y}_t^T \mathbf{F}_t \mathbf{x} + \tilde{\mathbf{x}}^T \mathbf{F}_t^T \mathbf{F}_t \left(\frac{\mathbf{x} \odot \mathbf{x}}{\tilde{\mathbf{x}}} \right).$$

where \odot denotes elementwise product and division of two vectors is understood componentwise.

Step 1 is a non-negative least-squares (NNLS) problem and it can be solved efficiently using the LFBGS-B algorithm of [3]. The solution to Step 2 is obtained in closed form by solving $\nabla_{\mathbf{x}} g_t(\mathbf{x}_t, \mathbf{x}_{t-1}) = 0$ for \mathbf{x}_t , which yields the following *multiplicative update* (same as the ISRA [4] update),

$$\mathbf{x}_t = \mathbf{x}_{t-1} \odot \frac{\mathbf{F}_t^T \mathbf{y}_t}{\mathbf{F}_t^T \mathbf{F}_t \mathbf{x}_{t-1}}. \quad (7)$$

Note that this update respects the non-negativity constraint on \mathbf{x} as all involved quantities are non-negative.

At this point the reader might wonder why we solve for \mathbf{f}_t using NNLS, while for \mathbf{x} we merely perform descent using the technique of auxiliary functions. The reasoning is as follows: The unknown image \mathbf{x} is common to all the observed images, while the PSFs differ from image to image. Hence, \mathbf{x} and \mathbf{f}_t need to be treated differently. To that end, given the current estimate of the image \mathbf{x}_{t-1} and the new input image \mathbf{y}_t , we minimize the loss L_t w.r.t. \mathbf{f}_t , while for \mathbf{x} we merely descend on L_t . Doing so essentially ensures that \mathbf{f}_t is fit well¹, and that \mathbf{x}_t is not too different from \mathbf{x}_{t-1} .

After performing Steps 1 and 2 it is clear that we have *improved* the estimates of \mathbf{f}_t and \mathbf{x} , because Step 1 minimizes over \mathbf{f} , while in Step 2 the auxiliary function $g_t(\mathbf{x}; \mathbf{x}_{t-1})$ ensures we make a descent. Thus, the descent guarantee (6) holds, and takes the form

$$L_t(\mathbf{x}_t, \mathbf{f}_t; \mathbf{y}_t) \leq L_t(\mathbf{x}_{t-1}, \mathbf{f}; \mathbf{y}_t),$$

for arbitrary \mathbf{x}_{t-1} and \mathbf{f} . At a glance we illustrate the pseudo-code of our method as Algorithm 1.

```

Input: Stream of images  $\mathbf{y}_t$  for  $t \geq 1$ 
Output: Reconstructed image  $\mathbf{x}$ 
 $t \leftarrow 1$ ;
 $\mathbf{x}_t \leftarrow \mathbf{y}_t$  padded with zeros or with mean;
while another image  $\mathbf{y}_{t+1}$  available do
   $t \leftarrow t + 1$ ;
   $\mathbf{x}_t \leftarrow \mathbf{x}_{t-1}$ ;
  for a few steps do
     $\mathbf{f}_t = \operatorname{argmin}_{\mathbf{f} \geq 0} \|\mathbf{y}_t - \mathbf{F}\mathbf{x}_t\|^2$ ;
     $\mathbf{x}_t \leftarrow \mathbf{x}_t \odot \frac{\mathbf{F}_t^T \mathbf{y}_t}{\mathbf{F}_t^T (\mathbf{F}_t \mathbf{x}_t)}$ ;
  end
end
return last estimate  $\mathbf{x}_t$ 

```

Algorithm 1: Online Blind Deconvolution

Digression (Poisson): We mention in passing that for the Poisson case, our approach takes the form:

1. $\mathbf{f}_t = \operatorname{argmin}_{\mathbf{f} \geq 0} \operatorname{KL}(\mathbf{y}_t; \mathbf{f} \star \mathbf{x}_{t-1})$
2. $\mathbf{x}_t = \operatorname{argmin}_{\mathbf{x} \geq 0} g_t(\mathbf{x}, \mathbf{x}_{t-1})$,

where KL denotes the generalized Kullback-Leibler divergence and $g(\mathbf{x}; \tilde{\mathbf{x}})$ is an appropriate ‘‘auxiliary’’ function for the associated loss. Step 1 can be again solved using the LFBFGS-B method, while the EM derivation of [17] yields

$$g_t(\mathbf{x}, \tilde{\mathbf{x}}) = \mathbf{1}^T \mathbf{F}_t \mathbf{x} - \mathbf{y}_t^T \mathbf{\Lambda} \log(\mathbf{x}) + c,$$

where c is a constant independent of \mathbf{x} , and $\mathbf{\Lambda}$ is the matrix $\mathbf{\Lambda} = [\lambda_{jk}] = [(\mathbf{F}_t)_{jk} \tilde{\mathbf{x}}_k / (\mathbf{F}_t \tilde{\mathbf{x}})_j]$. Algorithm 1 can be modified appropriately to accommodate these changes.

¹For cases where \mathbf{f}_t does not differ too much from \mathbf{f}_{t-1} , it could be updated in a regularized fashion or even via a multiplicative update similar to that for \mathbf{x}_t .

4. Experiments

Having introduced our method, we first test it in a controlled setup where the ground truth is known to study how noise influences its performance. Then, we apply our method to real astronomy image sequences of a binary star system and of the planet Mars; both datasets were obtained from ground-based telescopes and were subject to atmospheric turbulences.

4.1. Controlled experiments to test influence of noise

We analyze below how the performance of our method is affected by varying noise levels. As the true image \mathbf{x} we chose the well-known Lena image that was subsampled to 128×128 , and scaled between zero and one. We added noise as per (1), i.e., $\mathbf{y}_t = \mathbf{f}_t \star \mathbf{x} + \mathbf{n}_t$, where \mathbf{n}_t was normally distributed with mean $\mathbf{0}$ and diagonal covariance matrix $\sigma^2 \mathbf{I}$. The PSFs \mathbf{f}_t were generated as Gaussian blurs with randomly chosen covariances with varying rotation and scaling along the axes. The additive noise was sampled independently for each blurry image $\mathbf{f}_t \star \mathbf{x}$ with varying standard deviation $\sigma \in \{0.0000, 0.0100, 0.0167, 0.0278, 0.0464, 0.0774, 0.1292\}$. The randomly generated PSFs \mathbf{f}_t were the same for all seven tested noise levels.

The initial examples (\mathbf{y}_1) of our dataset are shown in the first row of Fig. 1 (with increasing σ from left to right). The bottom row shows the corresponding estimates \mathbf{x}_t after having seen the blurry images $\mathbf{y}_1, \dots, \mathbf{y}_{100}$ for the different noise levels. In all these seven runs we see that the sharpness is improved. Further, as expected, the noisier the input blurry images, the noisier is the resulting reconstruction.

The top panel of Fig. 2 shows how the relative error develops as the seven runs proceed (plotted with noise level increasing from bottom to top). Naturally, the achievable performance depends on the amount of noise. Since all curves are decreasing, we see that the relative error is reduced for the plotted noise levels, which means that the initial estimate $\mathbf{x}_1 = \mathbf{y}_1$ is improved as the algorithm progresses.

The bottom panel of Fig. 2 shows the loss L_t at time t . It is noteworthy that whenever the loss was large, i.e., if the currently observed image \mathbf{y}_t could not be fit well by $\mathbf{f}_t \star \mathbf{x}_{t-1}$, the relative error decreased substantially: by the update formula, large errors lead to large updates. We have also observed that large errors often occur for observed images \mathbf{y}_t that are particularly sharp. Such images can be called ‘‘lucky’’ as they provide high frequency details that benefit the algorithm.

4.2. Application to astronomical data

Before we describe our results on real astronomical data, a few words about the physical setup are in order.

Diffraction-index fluctuations caused by atmospheric turbulence significantly degrade image quality (called ‘‘seeing’’ in astronomy) and result in a time and spatially



Figure 1. Top row: observed y_1 with increasing noise level from left to right. Bottom row: corresponding recovered images x_{100} .

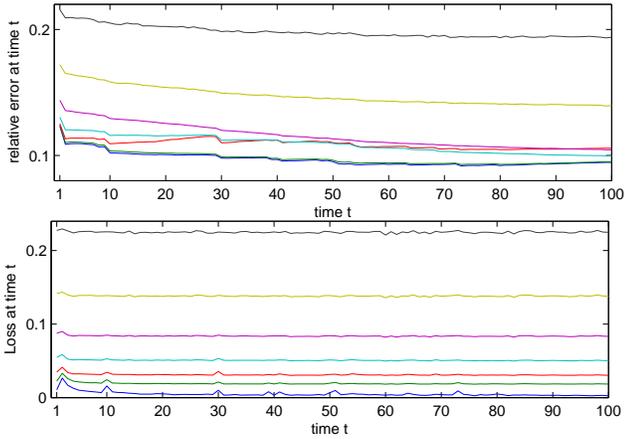


Figure 2. Top panel: relative error at time t for increasing noise standard deviations from bottom to top. Bottom panel: The corresponding instantaneous losses at time t .

varying PSF. In the Kolmogorov model, the effect of turbulent air layers is included in so-called phase screens that can be modelled as Gaussian random fields obeying certain second order statistics. These phase screens are stationary on a time scale of a few milliseconds and depend on the temporary seeing conditions. Thus, for images with short enough exposures the turbulence is effectively frozen. Long-exposures that essentially perform a time average inevitably wipe out high frequency information, while short-exposures encapsulate information up to the diffraction-limited upper frequency bound.

Because of their large distance from our solar system, observed stars may be considered as point sources. The image and diffraction pattern of a single star recorded by an earth-based observatory corresponds to the combined PSF of the telescope and the atmosphere. The region within which this PSF is constant is called an *isoplanatic patch*. In both of the following datasets we assume that the whole image is covered by a single isoplanatic patch, i.e., we assume the PSF to be invariant in a given frame.

The first dataset is an image sequence of the binary star

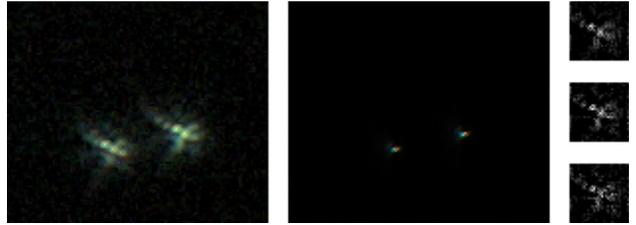


Figure 3. Epsilon Lyrae 1: typical observed image y_{300} (left), reconstruction x_{300} after iterations (middle), estimated PSFs f_{300} for each color channel.

system Epsilon Lyrae 1 of the constellation Lyra. This sequence consists of 333 frames, each 132×112 pixels in size. Fig. 5 shows in two columns the first 40 frames. Each row shows from left to right the original image y_t , the image estimated by our algorithm, and the averaged image $t^{-1} \sum_{i=1}^t y_i$. We chose PSFs of size 30×30 pixels. Fig. 3 shows the result of our algorithm after 300 iterations along with the estimated PSFs for each color channel. Note how noisy the observed image y_t is (left), while our estimate is almost noise free (middle). Furthermore we see that both stars have almost identical diffraction patterns which moreover look very similar to the estimated PSFs (shown on the right). This finding justifies our assumption about a constant PSF for the whole image, which is further reconfirmed by the fact that the stars of Epsilon Lyrae 1 are separated by only $2''$ (two arc seconds). Thus our assumption that they lie in the same isoplanatic patch is highly realistic for Epsilon Lyrae 1.

Next we evaluate our method on an extended celestial object. We run it on an image sequence of Mars. The original recording consists of 1450 frames, each 160×170 pixels in size. Fig. 6 shows the temporal evolution of our algorithm in two columns, where in each row from left to right the input frame, the estimated, and the averaged images are shown. After just a few iterations our estimated image reveals more detailed structure than what can be seen in any single frame of the original sequence. Unfortunately, some ringing artifacts on the right half of Mars also become



Figure 4. Mars: typical observed frame y_{36} (left), estimated image x_{36} (middle), an image reconstructed by REGISTAX (right).

visible. At present, we can not tell whether these effects are solely caused by overfitting, or whether the assumption of a PSF which is invariant all over Mars (which had an angular size of around $19''$ in May 2001 when our dataset was recorded²) is somewhat unrealistic and thus limiting the achievable performance. Nonetheless, Fig. 4 compares our estimated image after 36 iterations (center) with a typical frame from the image sequence (left) and a reconstruction obtained by REGISTAX (right), a widely used lucky imaging software. For reconstruction REGISTAX combined the best 30 frames out of *all* 1450 frames by image registration. No post-processing or finishing was performed in any case. In spite of the presence of some ringing artifacts, in some image regions our proposed algorithm reveals finer details than REGISTAX.

5. Conclusions and Future Work

In this paper we studied blind deconvolution for astronomy images, introducing an efficient online algorithm for it. Controlled experiments and applications to real astronomy pictures showed that our proposed online methods works, and that it alleviates some of the distortions caused by atmospheric turbulences to astronomical images.

Several interesting aspects remain yet to be analyzed. Currently our method does not rely on additional assumptions like image priors or PSF priors. However, incorporating such prior knowledge is a promising extension. For instance, regularization schemes can model the fact that the PSFs vary slowly over a sequence of very short exposures. Generalizing our approach to more difficult image distortions, such as spatially variant PSFs or nonlinear PSFs will open up new challenging directions of research.

References

- [1] B. Anconelli, M. Bertero, P. Boccacci, M. Carbillet, and H. Lanteri. Restoration of interferometric images - III. Efficient Richardson-Lucy methods for LINC-NIRVANA data reduction. *Astron. & Astrophys.*, 430(2), 2005. 2
- [2] M. Bertero and P. Boccacci. Image restoration methods for the Large Binocular Telescope. *Astron. & Astrophys.*, 147, 2000. 2

²According to [15], the isoplanatic angle is usually smaller.

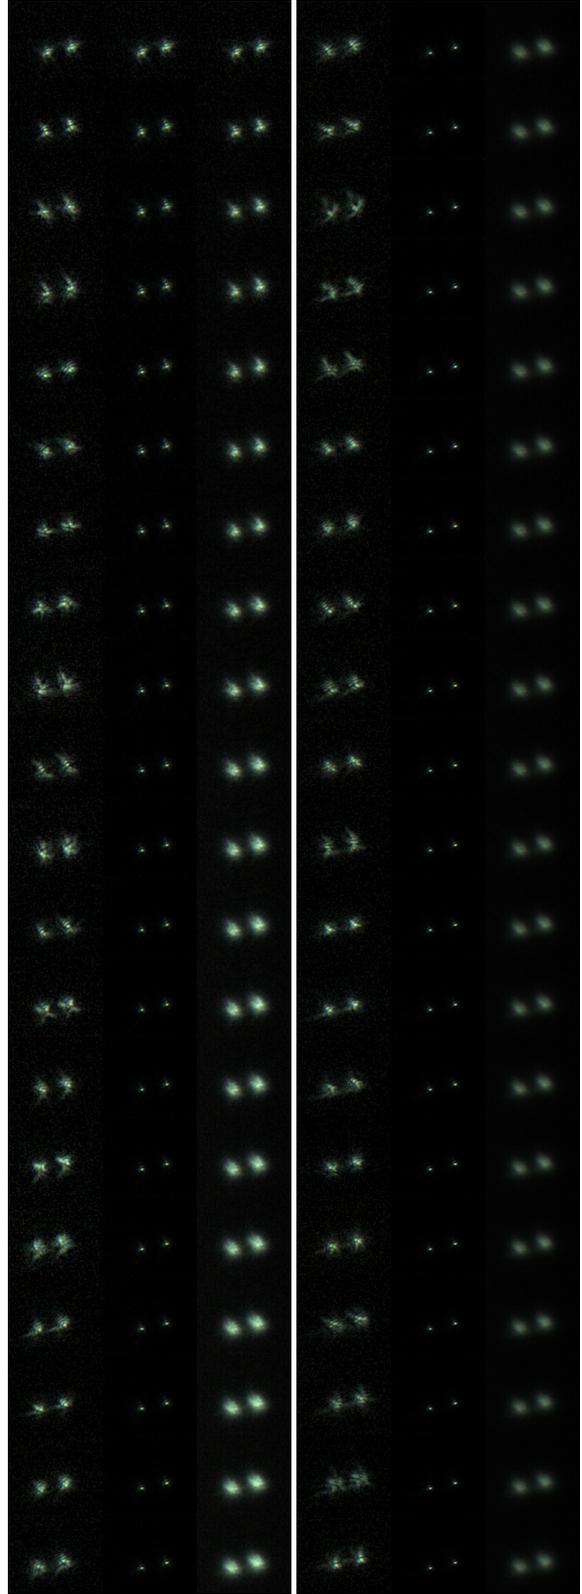


Figure 5. Temporal evolution of our algorithm applied to an image sequence of the double star system Epsilon Lyrae 1. The first 40 iterations are shown. In each row from left to right: original frame, estimated image, averaged image.

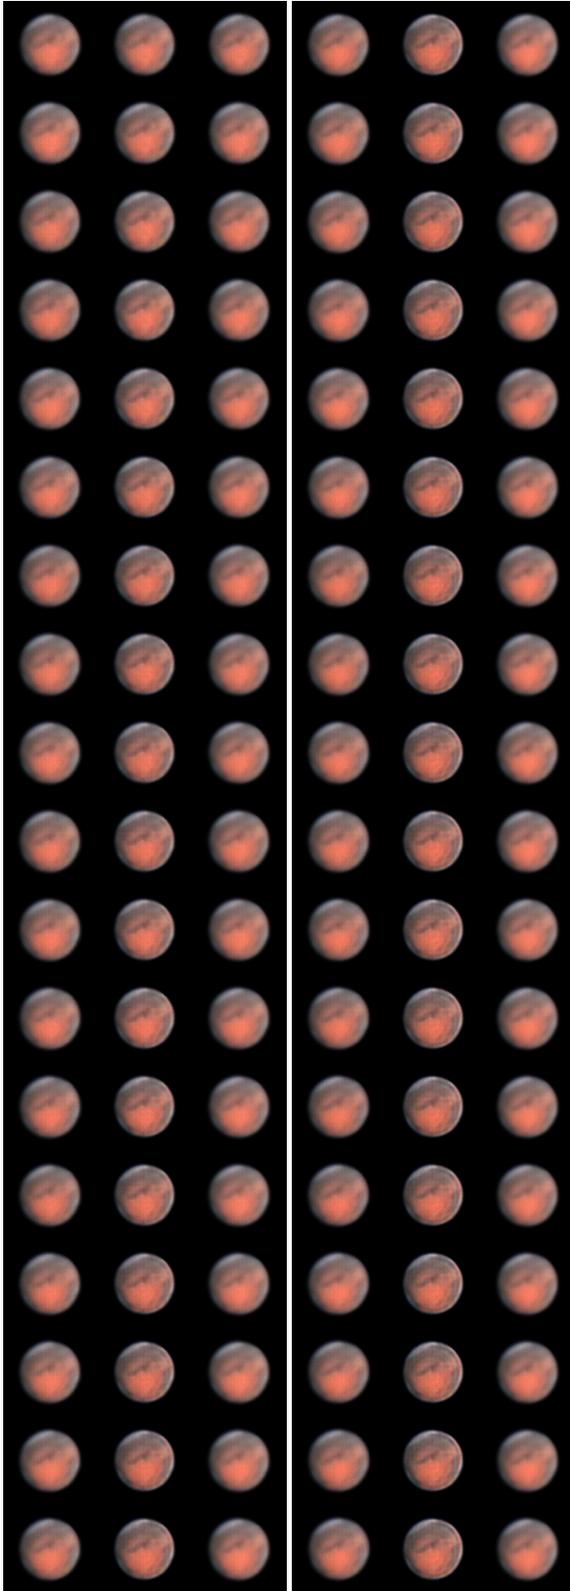


Figure 6. Temporal evolution of our algorithm applied to an image sequence of Mars. The first 36 iterations are shown. In each row from left to right: original frame, estimated image, averaged image.

- [3] R. Byrd, P. Lu, J. Nocedal, and C. Zhu. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM J. Scientific Comp.*, 16(5):1190–1208, 1995. 3
- [4] M. E. Daube-Witherspoon and G. Muehllehner. An Iterative Image Space Reconstruction Algorithm Suitable for Volume ECT. *IEEE Tran. Medical Imaging*, 5(2):61–66, 1986. 3
- [5] D. L. Fried. Probability of getting a lucky short-exposure image through turbulence. *J. Optical Soc. Amer.*, 86(1651-1657), 1978. 1
- [6] A. S. Fruchter and R. N. Hook. Drizzle: A Method for the Linear Reconstruction of Undersampled Images. *Publications Astron. Soc. of the Pacific*, 114:144–152, 2002. 1
- [7] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3), 2007. 2
- [8] K. T. Knox and B. J. Thompson. Recovery of images from atmospherically degraded short-exposure photographs. *Astrophysical J.*, 193:L45–L48, 1974. 1
- [9] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Mag.*, 13(3):43–64, 1996. 2
- [10] A. Labeyrie. Attainment of Diffraction Limited Resolution in Large Telescopes by Fourier Analysing Speckle Patterns in Star Images. *Astron. & Astrophys.*, 6, 1970. 1
- [11] R. Molina, J. N. nez, F. J. Cortijo, and J. Mateos. Image Restoration in Astronomy: A Bayesian Perspective. *IEEE Signal Proc. Mag.*, 2001. 2
- [12] D. Nettleton. Convergence Properties of the EM Algorithm in Constrained Parameter Spaces. *The Canadian J. Statistics*, 27(3):639–648, 1999. 3
- [13] A. Rav-Acha and S. Peleg. Two motion-blurred images are better than one. *Pattern Recognition Lett.*, 26(3), 2005. 2
- [14] W. H. Richardson. Bayesian-based iterative method of image restoration. *J. Optical Soc. Amer.*, 62:55–59, 1972. 2
- [15] F. Roddier, J. M. Gilli, and G. Lund. On the origin of speckle boiling and its effects in stellar speckle interferometry. *Journal of Optics*, 13:263–271, 1982. 6
- [16] T. J. Schulz. Multiframe blind deconvolution of astronomical images. *J. Optical Soc. Amer.*, 10:1064–1073, 1993. 2
- [17] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Tran. Medical Imaging*, 1:113–122, 1982. 4
- [18] R. N. Tubbs. *Lucky Exposures: Diffraction limited astronomical imaging through the atmosphere*. PhD thesis, Cambridge Univ., 2003. 1
- [19] R. Vio, J. Bardsley, and W. Wamsteker. Least-squares methods with Poissonian noise: Analysis and comparison with the Richardson-Lucy algorithm. *Astron. & Astrophys.*, 436, 2005. 2, 3
- [20] R. Vio, J. Nagy, and L. T. and. A simple but efficient algorithm for multiple image deblurring. *Astron. & Astrophys.*, 416(1), 2004. 2
- [21] R. Vio and W. Wamsteker. Use of the single-image approach in multiple-image deblurring problems with poissonian noise. *Astron. & Astrophys.*, 439, 2005. 2
- [22] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Int. Conf. Machine Learning (ICML)*, 2003. 2