Max–Planck–Institut für biologische Kybernetik
Max Planck Institute for Biological Cybernetics

MAX-PLANCK-GESELLSCHAFT

Technical Report No. 187

An Incremental GEM Framework for Multiframe Blind Deconvolution, Super-Resolution, and Saturation Correction[1]

Stefan Harmeling[2], Suvrit Sra[2], Michael Hirsch[2], Bernhard Schölkopf[2]

November 2009

[1] This paper has been submitted to Int. Conf. on Computational Photography (ICCP), 2010.
[2] Department Schölkopf, email: harmeling;sra;hirsch;bs@tuebingen.mpg.de

An Incremental GEM Framework for Multiframe Blind Deconvolution, Super-Resolution, and Saturation Correction

Stefan Harmeling, Suvrit Sra, Michael Hirsch, Bernhard Schölkopf
MPI for Biological Cybernetics, Spemannstraße 38, 72076 Tübingen, Germany
firstname.lastname@mpg.tuebingen.de

Abstract

We develop an incremental generalized expectation maximization (GEM) framework to model the multiframe blind deconvolution problem. A simplistic version of this problem was recently studied by Harmeling et al. *[4]. We solve a more realistic version of this problem which includes the following major features: (i) super-resolution ability* despite *noise and unknown blurring; (ii) saturation-correction, i.e., handling of overexposed pixels that can otherwise confound the image processing; and (iii) simultaneous handling of color channels. These features are seamlessly integrated into our incremental GEM framework to yield simple but efficient multiframe blind deconvolution algorithms. We present technical details concerning critical steps of our algorithms, especially to highlight how all operations can be written using matrix-vector multiplications. We apply our algorithm to real-world images from astronomy and super resolution tasks. Our experimental results show that our methods yield improved resolution and deconvolution at the same time.*

1. Introduction

Blind deconvolution of blurred images is a well-known difficult task. In this paper we study a particular blind-deconvolution setup: *multiframe blind deconvolution* (MFBD). Here one seeks to recover an underlying "true" image from a sequence of blurred (and noisy) observations. By having access to multiple blurred samples, one can hope to obtain a better reconstruction than would be possible from just a single image. The reconstruction task is however not trivial because each observation is differently blurred and noisy. Recently, in the context of Computational Photography Harmeling *et al*. [4] introduced a simple and efficient approach for MFBD. However, their setup lacks several crucial features that must be considered for deconvolution in a realistic multiframe scenario. We expand on this claim below.

1.1. Motivation & Contributions

For concreteness, we also focus (like [4]) on MFBD in astronomy, though clearly astronomy is not the only possible application domain for our algorithms. To ensure a controlled but realistic setup we caputured our own images of the sky; doing so models the typical difficulties astronomers might face when working with multiple frames of the same underlying object.

[4] *et al*. broadly had the following approach to MFBD: (i) at each time-step obtain a blurry (and noisy) image frame as input; (ii) subjected this frame to blind deconvolution to estimate the blur kernel and the "true" image; and (iii) update an estimate of the true image and *discard* the observation. This simple approach immediately raises the question whether we can better exploit the observations without dramatically impacting storage and computational costs?

We show that one can indeed exploit the entire sequence to improve reconstruction quality. Having an entire sequence at our disposal suggests that we should be able to *super-resolve* [7] the reconstructed image. But one major hurdle still remains: the individual frames are both blurred and noisy which impedes super-resolution. We however design an algorithm that combines super-resolution with blind deconvolution, and obtain improved reconstructions.

A problem that particularly plagues astronomy images is that of *saturation* of pixels due to overexposure. These pixels receive so many photons that they reach the peak intensity permitted by the hardware. Such saturated pixels make it difficult to reliably deblur, denoise, or to even run an inference procedure such as estimating the true brightness (magnitude) of a star. We note that another problem that arises with long exposures is that one loses the higher-frequency structure (finer details) of the image due to "averaging."—a problem that super-resolution can alleviate.

We present a simple but effective approach to tackle the adverse effects of overexposure. Saturation-correction also benefits super-resolution, allowing it to more reliably reconstruct long-time exposures. The MFBD paper [4] in contrast considered only short-term exposures, and ignored the influence of saturation.

The third (and last) main enhancement that we show is how to handle images with different color channels.

To seamlessly integrate super-resolution, saturation-correction, and color channels into MFBD, we provide a theoretical model that views the whole setup using incremental generalized expectation maximization (GEM) [6]. This viewpoint also lends theoretical justification to the heuristic procedure of [4], while opening doors to further theoretical analysis (e.g., convergence).

1.2. Related Work

Beyond the immediately relevant work of Harmeling *et al*. [4], there are several other papers that study multiframe deblurring. Some references are: [9, 12, 13, 1, 5]. All of these papers approach the multiframe problem in a non-incremental fashion as opposed to our incremental choice. When the number of frames grows, such non-incremental approaches rapidly become computationally prohibitive, and surprisingly, as observed by [4], possibly even inferior (this observation can be now justified using our incremental GEM interpretation; see [6] for more intuition). Additional relevant astronomy papers are summarized in [4].

In the non-astronomy setting, the idea of using more than one frame to correct for motion (or even other blurring) has recently gained attention [8, 3], though similar previous ideas also exist as noted by Cai *et al.* [3]. We observe that even though the literature on blind-deconvolution and super-resolution is vast, only a counted few papers discuss the setting where both are done simultaneously.

Directly relevant is the work of Šroubek *et al.* [10, 11] who consider simultaneous (non-incremental) super-resolution and blind deconvolution, but depend on image and blur priors. Šroubek *et al*. themselves note (in Section 4.1 of [11]) that their method becomes unstable for larger super resolution factors than 2.5. Our model exploits the abundance of data by *not* assuming any image or blur kernel priors, thereby leading to a simpler algorithm with almost no parameter tuning. Our method is applicable to much higher super resolution factors than 2.5. In fact in the experiments section we will resolve datasets that Šroubek *et al.* [11] considered up to a factor of eight and compare it against their reported results.

1.3. Summary of Remainder

The remainder of this paper is organized as follows. Section 2 presents an incremental GEM view of the problem. Subsequently in Section 3 and its subsections, we present our extensions to the multiframe blind deconvolution setup, including super-resolution, correcting for overexposed pixels, and simultaneous handling of several color channels. In Section 4 we describe some of the technical details that are crucial for implementing the algorithm. Finally, Section 5 presents the results of our algorithms on real-world data.

2. Blind deconvolution as incremental GEM

For notational simplicity we describe everything in terms of one-dimensional images and point spread functions (PSFs). Our exposition easily generalizes to two-dimensional images and PSFs, and we detail this later in Section 4, especially because in an actual implementation we work with 2D images.

We denote the "true" image by x, each observed (input) image by y_t, and each PSF by f_t. Throughout the paper we use $f * x$ to represent convolution. All results apply to circular and non-circular convolutions.

First we derive the algorithm of Harmeling *et al.* [4] as an instance of incremental GEM. For this derivation we view the image sequence $y_1, \ldots, y_T$ as observed random variables, while the PSFs are seen as latent variables. The sought-after image x is a parameter of a factorizing probabilistic model,

$$p(y_1, \ldots, y_T|x) = \prod_{t=1}^{T} \int p(y_t, f_t|x) df_t. \qquad (1)$$

Given the observed frames $y_1, \ldots, y_T$, our goal is to find the maximum likelihood estimate of parameter x. As usual, the log-likelihood $\ell(x) = \log p(y_1, \ldots, y_T|x)$ can be bounded from below using Jensen's inequality,

$$\begin{aligned} \ell(x) &= \sum_{t=1}^{T} \log \int p(y_t, f_t|x) df_t & (2) \\ &\geq \sum_{t=1}^{T} \int q_t(f_t) \log \frac{p(y_t, f_t|x)}{q_t(f_t)} df_t & (3) \\ &= \sum_{t=1}^{T} \mathcal{F}_t(q_t, x) & (4) \end{aligned}$$

where each q_t is an arbitrary distribution of f_t. The incremental variant of EM picks *some* image y_t and estimates q_t and x_t via:

E-step: find q_t that maximizes $\mathcal{F}_t(q_t, x_{t-1})$,
M-step: find x_t that maximizes $\mathcal{F}_t(q_t, x_t)$.

Instead of searching for arbitrary distributions q_t we restrict ourselves to delta peaks which we parametrize with f_t. This is merely the "winner-take-all" variant of incremental EM [6]. This choice for f_t greatly simplifies the E- and M-step to a maximization of $p(y_t, f_t|x)$. Assuming further a flat prior on f_t, the E- and M-step maximize $p(y_t|f_t, x)$ with respect to f_t and x, respectively.

Modelling each observed frame y_t as a convolution of the underlying image (parametrized by) x with some unknown PSF f_t plus Gaussian noise, the density $p(y_t|f_t, x)$ can be chosen to be a Gaussian with mean $f_t * x$ and some diagonal covariance,

$$p(y_t|f_t, x) \propto \exp\bigl(-\tfrac{1}{2}\bigl\|y_t - f_t * x\bigr\|^2\bigr). \tag{5}$$

This choice naturally leads to E- and M-steps that are (nonnegative) least squares problems.

The M-step requires further discussion. When updating the parameter x, it is also sensible to perform only a single (or a few steps) of the minimization of $\|y_t - f_t * x\|^2$ with respect to x. Do so is particularly useful if we consider noncircular convolution, for which x is longer than y_t. Thus, instead of a full minimization with respect to x, we prefer to perform merely an ascent in the M-step, which is nothing but GEM.

Since we are dealing with images we further assume that x and $f_1, \ldots, f_T$ are nonnegative. Thus, the incremental GEM approach with “winner-take-all” distributions for f_t results in the following update steps, which are executed one by one on all observed images $y_1, \ldots, y_T$:

E-step: find $f_t \geq 0$ that minimizes

$$\|y_t - f_t * x_{t-1}\|^2, \tag{6}$$

M-step: find $x_t \geq 0$ such that

$$\|y_t - f_t * x_t\|^2 \leq \|y_t - f_t * x_{t-1}\|^2. \tag{7}$$

These are exactly the two alternating steps that [4] proposed, but motivated only heuristically.

3. Improved Multiframe Blind Deconvolution

The basic algorithm described above is able to recover the underlying, clear image given a sequence of blurry images with similar resolution [4]. We discuss now how to obtain super-resolved results and how to deal with saturated pixels. We also discuss how images with different color channels can be handled.

3.1. Super-resolution

Since we are interested in a *single* image x but have *several* observations y_t, despite blurring there is hope that we could infer a super-resolved image x, provided we incorporate change of resolution into the forward model. For resolution change we define the resizing matrix,

$$D_m^n = (I_m \otimes 1_n^T)(I_n \otimes 1_m)/m, \tag{8}$$

where I_m is the $m \times m$ identity matrix, 1_m is an m dimensional column vector of ones, and $\otimes$ denotes the Kronecker product. D_m^n transforms a vector v of length n into a vector of length m. Note that the sum of v's entries $1_n^T v = 1_m^T D_m^n v$ is preserved (formally verified by applying Eq. (20) twice). This is a favorable property for images, as the number of photons observed should not depend on the resolution. Note that even if the sizes m and n are not multiples of each other, D_m^n will interpolate appropriately.

To avoid double indexing let $n = l_y$ be the length of y. For k times super-resolution we choose x and f large enough such that the vector $f * x$ has length kn. Then we replace the density of y_t given x and f_t by,

$$p(y_t|f_t, x) \propto \exp\bigl(-\tfrac{1}{2}\|y_t - D_n^{kn}(f_t * x)\|^2\bigr) \tag{9}$$

which leads to the following update steps:

E-step: find $f_t \geq 0$ that minimizes

$$\|y_t - D_n^{kn}(f_t * x_{t-1})\|^2, \tag{10}$$

M-step: find $x_t \geq 0$ such that

$$\|y_t - D_n^{kn}(f_t * x_t)\|^2 \leq \|y_t - D_n^{kn}(f_t * x_{t-1})\|^2. \tag{11}$$

Note that the positive scaling factor k is not restricted to be integral.

3.2. Overexposed pixels

As previously motivated, saturation impacts image restoration considerably. Even more so in astronomical imaging where we would like to capture faint stars together with stars that are orders of magnitude brighter. In such a setup, overexposed pixels occur frequently and pose problems. A realistic deconvolution method should be able to deal with pixels that are saturated, i.e., those that hit (or come close to) the maximal possible pixel value.

One way to deal with these pixels is to identify them, but not let them impact the objective function much, by ignoring them. However, since each frame y_t can have different pixels attaining saturation (different frames are aligned differently), we have to check at each iteration which pixels in the current y_t are saturated. To ignore these pixels we define a weighting matrix,

$$C_t = \text{Diag}\bigl([y_t < 65535]\bigr) \tag{12}$$

assuming that the value of a saturated pixel is 65535, and using the Iverson brackets component-wise, i.e., $[y_t < 65535]$ is a vector of ones and zeros depending on whether the corresponding component of y_t is less than 65535. Then, we can write the updates ignoring the saturated pixels simply by replacing the Euclidean norm with a weighted norm $\|v\|_C^2 = v^T C v$. The update steps now become:

E-step: find $f_t \geq 0$ that minimizes

$$\|y_t - f_t * x_{t-1}\|^2_{C_t}, \tag{13}$$

M-step: find $x_t \geq 0$ such that

$$\|y_t - f_t * x_t\|^2_{C_t} \leq \|y_t - f_t * x_{t-1}\|^2_{C_t}. \tag{14}$$

Note that this approach is equivalent to removing the saturated pixels from the probabilistic model.

One might ask whether we can really recover pixels in x that are saturated in most of the frames y_t? The answer is yes, and can be understood as follows. The photons corresponding to such a pixel in x have been spread across a whole set of pixels in each frame y_t because of the PSF f_t. Thus, if not all these pixels in y_t are saturated, the true value for the corresponding pixel in x will be recovered.

Using the same ideas as above, we can also ignore dead pixels that do not record any actual intensity, but rather return arbitrary values (due to the camera electronics). Most of these pixels can be identified by finding outliers in the components of the residual vector $y_t - f_t * x_t$. Setting their corresponding entries in C_t to zero also makes our method robust against dead pixels.

3.3. Simultaneous Estimation of Color Channels

Color images can be seen as at stack of grey images, each for a different color channel, e.g., $y_t^{(R)}, y_t^{(G)}, y_t^{(B)}$ for red, green, and blue, respectively. The simplest option to deal with colors is to apply our method to each channel separately and then to combine the results by image registration. However, a simpler solution exists because sometimes all colors channels share the same PSF f. In such cases, we can model y_t as

$$p(y_t|f_t, x) \propto \exp\Big(-\tfrac{1}{2}\left\| \begin{bmatrix} y_t^{(R)} \\ y_t^{(G)} \\ y_t^{(B)} \end{bmatrix} - \begin{bmatrix} f_t * x^{(R)} \\ f_t * x^{(G)} \\ f_t * x^{(B)} \end{bmatrix} \right\|^2\Big). \tag{15}$$

Translating this model assumption to an E- and M-step is straightforward. Section 4 shows how convolution of multiple color channels can also be seen as a matrix-vector multiplication operation, which greatly eases implementation.

4. Implementation details

Although in the theoretical part we only consider vectors, one-dimensional convolutions, and vector-norms, all results naturally generalize to two-dimensional images. However, efficiently implementing these algorithms for two-dimensional images requires careful consideration, and some technical details which this section aims to provide.

4.1. Convolution as Matrix-Vector Multiplication

We introduced $f * x$ as the convolution, which could be either circular or non-circular. Both can also be written in two ways as matrix-vector multiplication,

$$f * x = Fx = Xf \tag{16}$$

depending whether our interest is in x or f. Matrices F and X can be specified as

$$F = I_y^T W^{-1} \operatorname{Diag}(W I_f f) W \tag{17}$$

$$X = I_y^T W^{-1} \operatorname{Diag}(W x) W I_f \tag{18}$$

where for the circular convolution we drop I_y^T. Matrix W is the discrete Fourier transform matrix, i.e. Wx is the Fourier transform of x. $\operatorname{Diag}(v)$ denotes the diagonal matrix with vector v along its diagonal. I_f and I_y are zero-padding matrices that ensure that $I_f f$ and $I_y y$ have the same length as x. Note that I_f appends zeros while I_y prepends zeros.

For two-dimensional images and PSFs we have to consider two-dimensional Fourier transforms that can be written as left- and right-multiplications with W, which can also be represented as a single matrix-vector multiplication using the vectorization operator and Kronecker product,

$$\operatorname{vec}(W x W) = (W \otimes W) \operatorname{vec}(x), \tag{19}$$

noting that W is symmetric and Kronecker products fulfill

$$\operatorname{vec}(A\, B\, C^T) = (C \otimes A) \operatorname{vec}(B). \tag{20}$$

Zero-padding for two-dimensional images can be written in an similar way.

4.2. Resizing matrices for images

Resizing two-dimensional images can also be implemented by left- and right-multiplications: let x be an $m \times n$ image, then $D_p^m\, x\, (D_q^n)^T$ is an image of size $p \times q$. Using Eq. (20) we can write this operation as a matrix-vector multiplication as well,

$$\operatorname{vec}\big(D_p^m\, x\, (D_q^n)^T\big) = \big(D_q^n \otimes D_p^m\big) \operatorname{vec}(x). \tag{21}$$

4.3. Convolution of color image

Above we have shown how to write $f*x$ as matrix-vector multiplications Fx and Xf. This can be generalized to images with different color channels:

$$\begin{bmatrix} f_t * x^{(R)} \\ f_t * x^{(G)} \\ f_t * x^{(B)} \end{bmatrix} = \begin{bmatrix} F & 0 & 0 \\ 0 & F & 0 \\ 0 & 0 & F \end{bmatrix} \begin{bmatrix} x^{(R)} \\ x^{(G)} \\ x^{(B)} \end{bmatrix} = \begin{bmatrix} X^{(R)} \\ X^{(G)} \\ X^{(B)} \end{bmatrix} f_t \tag{22}$$

These representations allow us to apply the whole machinery to colored images.

	E-step		M-step	
Algorithm	A	C_t	A	C_t
Basic	X_{t-1}	I	F_t	I
Super-resolved	$D_n^{kn} X_{t-1}$	I	$D_n^{kn} F_t$	I
Saturated	X_{t-1}	$\mathrm{Diag}([y_t < 65535])$	F_t	$\mathrm{Diag}([y_t < 65535])$

Table 1. How to formulate the presented extensions with the quadratic cost function in Eq. (23).

4.4. Multiplicative updates

The suggested algorithms do not depend on how the alternating steps are implemented. [4] used LBFGS-B [2] for their E-step, and EM-style multiplicative updates for their M-step (similar updates can also be used for the E-step if desired). An important observation here is to *not* even solve for the PSF exactly, as was done by [4]. Doing so greatly impacts not only efficiency, but sometimes also experimental results. Performing early-stopping can act like a regularizer and can prevent overfitting.

More generally, minimizing quadratic cost functions under non-negativity constraints,

$$\min_{z \geq 0} \|y - Az\|_{C_t}^2 \tag{23}$$

with $\|v\|_C^2 = v^T C v$, can be implemented via multiplicative updates, e.g.,

$$z_{t+1} = z_t \odot \frac{A^T C_t y}{A^T C_t A z_t}, \tag{24}$$

where one starts from a strictly positive initial value $z_0 > 0$. Table 1 summarizes how all of the presented approaches can be implemented with multiplicative updates with appropriate choices of A and C_t. We rely crucially on the fact that the convolution $f * x$ can be represented either as matrix-vector multiplications Fx or Xf.

5. Results on real images

The first experiments compare our method on some super resolution tasks against the method of [11]. Then we show results on astronomical images, which were recorded using an off-the-shelf 12-inch f/10 MEADE LX200ACF Schmidt-Cassegrain telescope. For the latter we used two different cameras. The short exposure images were taken using an AVT PIKE F-032B uncooled CCD camera. The globular cluster M13 was imaged using a Canon EOS 5D digital single lens reflex (DSLR).

5.1. Super-resolution examples

The method which is most closely related to ours is the blind super-resolution approach of Šroubek *et al.* [11] against which we will compare in the following. For this we applied our method to some of the datasets of S. Farsiu and P. Milanfar[1] and show Šroubek *et al.*'s results reported in [11]. For brevity we consider only the "text" dataset (20 frames of size 57 × 49) and the "disk" dataset (20 frames of size 57 × 49).

The top rows of Figures 1 and 2 show typical input frames of the corresponding image sequences. The bottom rows show the result of our method for blind deconvolution with increasing super resolution factor (for one, two, four, and eight). We used the lower resolved images as initializations for the next scale factor. The obtained resolution with a factor of eight is clearly superior to the result of [11] who limited the factor to two in order to avoid stability issues. Note that our result with a factor of two compares favorably to Šroubek's as well.

5.2. Orion Trapezium

In our first experiment with astronomical images, we used the fast AVT PIKE camera to record a short video (191 frames acquired at 120 fps) of the Trapezium in the constellation Orion. It is formed by four stars ranging in brightness from magnitude 5 to magnitude 8, with angular separations around 10" to 20". The exposure of the individual frames is sufficiently short to "freeze" the atmospheric turbulence and thus retain the high-frequency information which is present in the atmospheric PSF—see Figure 3 for sample frames. Combining a stack of such frames to a single high-resolution image is the goal of the field of *lucky imaging*; usually this is achieved by registering and averaging over a subset of "lucky" frames where the turbulence is relatively small (and the "seeing" is thus good).

The first row of Figure 4 shows from left to right (a) an enlarged unprocessed frame, (b) the deconvolution results obtained by the basic algorithm of [4], (c) the result using the proposed method to handle saturation, and (d) the results if we additionally apply the proposed method for four times super-resolution.

An important application in astronomy is the measurement of the brightness of stars and other celestial objects (Photometry). To this end, a linear sensor response is required (for our purposes, the used CCD sensor may be assumed linear). The intensity counts can then be translated

[1] Available from http://www.soe.ucsc.edu/~milanfar/software/sr-datasets.html.

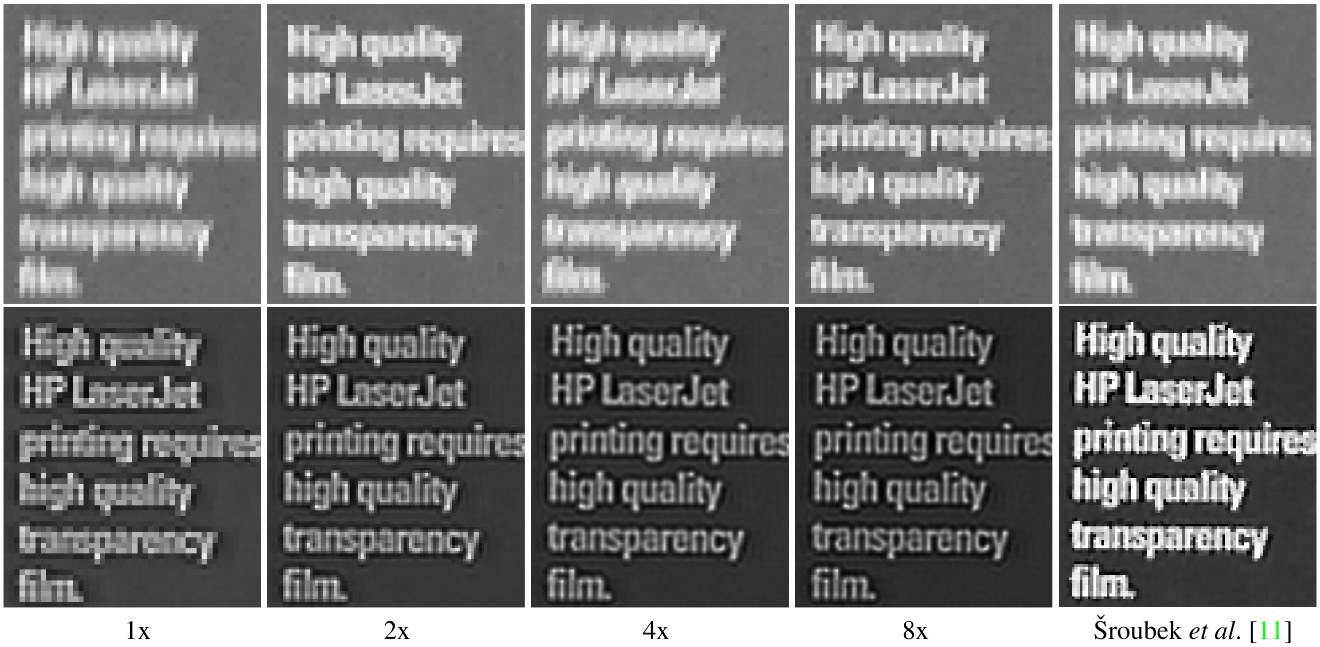

Figure 1. Text data: typical example frames (top row), results of our method for blind deconvolution with increasing super-resolution factor compared with Šroubek's results in [11] (bottom row from left to right).

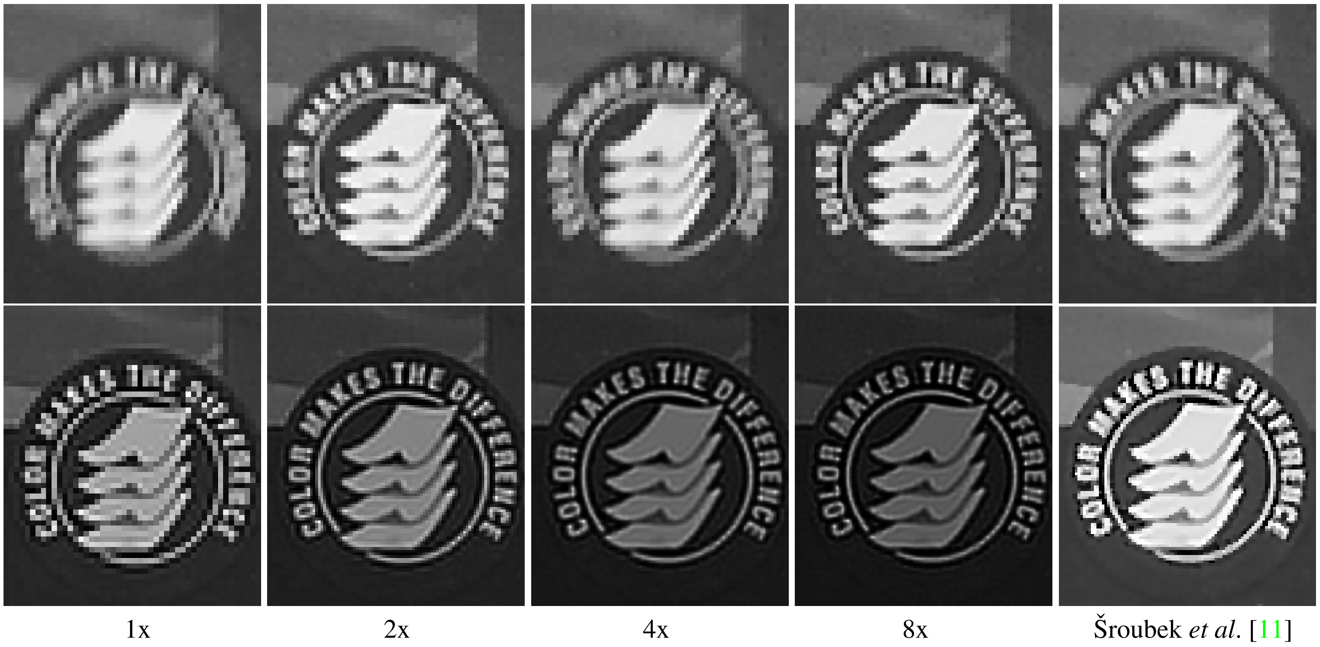

Figure 2. Disk data: typical example frames (top row), results of our method for blind deconvolution with increasing super-resolution factor compared with Šroubek's results in [11] (bottom row from left to right).

into stellar magnitudes.[2] Clearly, this is not directly possible for stars that saturate the CCD (i.e., where so many photons are recorded that the full well-capacities of the pixels are exceeded). However, we can use the proposed method for deconvolution with saturation correction and reconstruct the photon counts (image intensities) that we would have recorded had the pixel not been saturated; then we convert these into stellar magnitudes. For the Trapezium stars we obtain the encouraging results shown in Table 2.

[2] E.g., http://en.wikipedia.org/wiki/Apparent_magnitude.

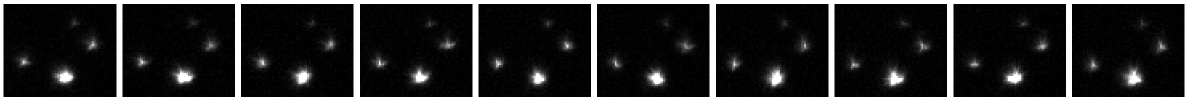

Figure 3. Orion Trapezium Cluster: example sequence of observed frames, $y_1, \ldots, y_{10}$.

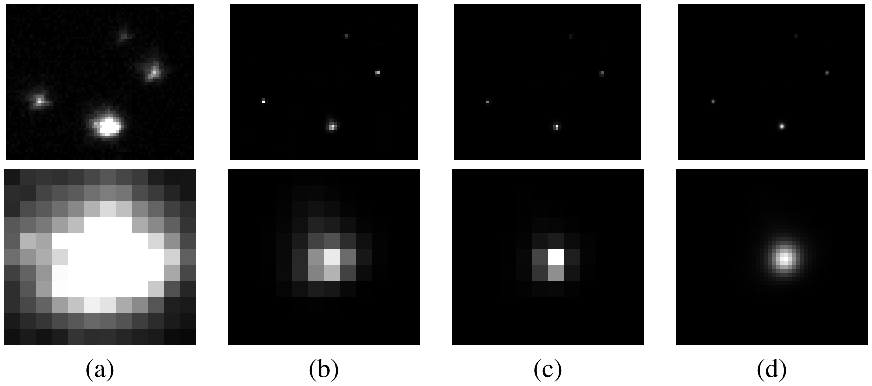

Figure 4. Orion Trapezium Cluster: (a) the first observed frame, (b) x_{191} for basic algorithm [4], (c) x_{191} for saturation corrected, and (d) x_{191} for saturation corrected and four times super-resolved. Top row shows the overall trapezium; bottom row shows the brightest star enlarged. The bottom row should show large squared pixels (this should work in Acrobat Reader).

Star	A	B	C	D
True magnitude	6.7 - 7.5	8.0 - 8.5	5.1	6.7
Est. mag. (w. sat. cor.)	6.4	8.0	5.2	6.0
Est. mag. (w/o sat. cor.)	6.8	8.0	6.5	6.4

Table 2. True star brightnesses (note that stars A and B have variable brightness), and values estimated after deconvolution, normalizing the brightness of star B to 8.0 and computed under the assumption of zero offset. The latter assumption is unrealistic, rendering the absolute values inaccurate; nevertheless it is reassuring that the proposed method for saturation correction leads to the correct *ordering* of brightness values, with star C being the brightest.

5.3. Globular cluster M13

M13 is a globular cluster in the constellation Hercules, around 25,000 light years away, with an apparent size of around 20'. It contains several 100,000 stars, the brightest of which has an apparent magnitude of 12. Such faint stars cannot be imaged using our equipment using short exposures; however, long exposures with budget equipment typically incur tracking errors, caused by telescope mounts that do not perfectly compensate for the rotation of the earth. In our case (focal length 3m, exposure times of 60sec per frame), this induced a significant motion blur in the images that we attempted to remove using the same algorithm used above on short exposures.

The top row of Figure 5 displays a long exposure with motion blur (left panel and the twice super-resolved result of our algorithm (right) applied to 26 motion degraded frames. In the bottom row we clearly see details in our reconstructed image (right) which where hidden in the recorded frames (left).

6. Conclusion

We proposed an incremental generalized expectation maximization framework for the multiframe blind deconvolution problem. We showed how simple linear algebra modifications allow us to incorporate super-resolution as well as saturatation correction for overexposed pixels, effectively increasing the dynamic range of the sensor. We compared our method against the results of [11] on some super resolution benchmark problems and demonstrated its practical applicability on astronomical images, including very short exposures, where the convolution is caused by atmospheric turbulence, as well as long exposures, where the convolution is induced by effective motion blur arising from mechanical inaccuracies in the telescope mount.

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. & Astrophy.*, 430(2), 2005. 2

[2] 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. 5

[3] J.-F. Cai, H. Ji, C. Liu, and Z. Shen. Blind motion deblurring using multiple images. *Journal of Computational Physics*, 2009. 2

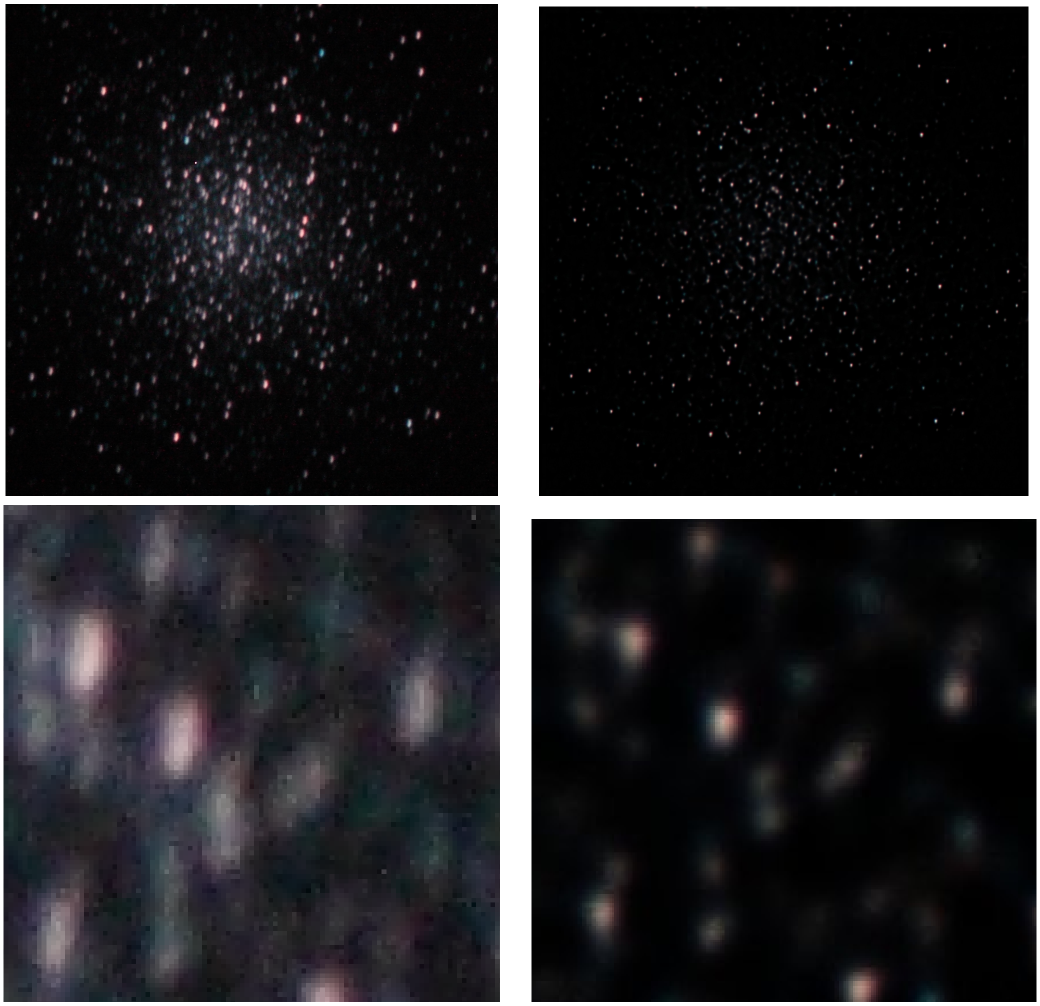

Figure 5. Globular cluster M13: (left) example observed frame, (right) result of saturation corrected, super-resolved multi-frame blind deconvolution; (top) overview, (bottom) detail.

[4] S. Harmeling, M. Hirsch, S. Sra, and B. Schölkopf. Online blind image deconvolution for astronomy. In *Proceedings of the IEEE Conference on Computational Photography*, 2009. 1, 2, 3, 5, 7

[5] N. F. Law and D. T. Nguyen. Multiple frame projection based blind deconvolution. *Electronics Letters*, 31(20), 1995. 2

[6] R. M. Neal and G. E. Hinton. A View of the EM Algorithm that Justifies Incremental, Sparse, and Other Variants. In M. I. Jordan, editor, *Learning in Graphical Models*. MIT press, 1998. 2

[7] S. C. Park, M. K. Park, and M. G. Kang. Super-resolution image reconstruction: a technical overview. *IEEE Sig. Proc. Mag.*, 20(3), 2003. 1

[8] A. Rav-Acha and S. Peleg. Two motion-blurred images are better than one. *Pattern Recognition Lett.*, 26(3), 2005. 2

[9] T. J. Schulz. Multiframe blind deconvolution of astronomical images. *J. Optical Soc. Amer.*, 10:1064–1073, 1993. 2

[10] F. Šroubek, G. Cristobál, and J. Flusser. A Unified Approach to Superresolution and Multichannel Blind Deconvolution. *IEEE Tran. Imag. Proc.*, 2007. 2

[11] F. Šroubek, G. Cristobál, and J. Flusser. Simultaneous superresolution and blind deconvolution. *Journal of Physics: Conference Series*, 124, 2008. 2, 5, 6, 7

[12] R. Vio, J. Nagy, and L. T. and. A simple but efficient algorithm for multiple image deblurring. *Astron. & Astrophy.*, 416(1), 2004. 2

[13] R. Vio and W. Wamsteker. Use of the single-image approach in multiple-image deblurring problems with poissonian noise. *Astron. & Astrophy.*, 439, 2005. 2