MULTIFRAME BLIND DECONVOLUTION, SUPER-RESOLUTION, AND SATURATION CORRECTION VIA INCREMENTAL EM

Stefan Harmeling, Suvrit Sra, Michael Hirsch, Bernhard Schölkopf

MPI for Biological Cybernetics, Spemannstraße 38, 72076 Tübingen, Germany

ABSTRACT

We formulate the multiframe blind deconvolution problem in an *incremental* expectation maximization (EM) framework. Beyond deconvolution, we show how to use the same framework to address: (i) *super-resolution* despite noise and unknown blurring; (ii) *saturation-correction* of overexposed pixels that confound image restoration. The abundance of data allows us to address both of these *without* using explicit image or blur priors. The end result is a simple but effective algorithm with no hyperparameters. We apply this algorithm to real-world images from astronomy and to super resolution tasks: for both, our algorithm yields increased resolution and deconvolved images simultaneously.

Index Terms— multiframe, blind deconvolution, super-resolution, saturation, incremental EM.

1. INTRODUCTION

The focus of this paper is *multiframe blind deconvolution* (MFBD): the problem of recovering the underlying "true" image from a sequence of blurry and noisy observations. This problem arises naturally in astronomy [1], for instance. One might think that having access to a sequence of observations alleviates the typical difficulties associated with blind deconvolution. But it is not so. MFBD is still very challenging because each observation is noisy and differently blurred.

Recently, in [1], we proposed an algorithm for MFBD which, though efficient, lacked several features crucial to realistic setups. This paper overcomes the deficiencies of [1], while still retaining its efficiency. As in [1] we do not employ any explicit image priors, a valid simplification because several frames are available, and the associated (linear) optimization is over-determined. Without an image prior, the method is not only simpler, but also avoids additional hyper parameters which otherwise would have to be tuned.

1.1. Summary of Contributions

The main contributions of this paper are summarized below.

- *Super-resolution*: We exploit the entire sequence of available images to artificially enhance resolution of the reconstructed image. The distinguishing feature is that we perform super-resolution and blind-deconvolution simultaneously.
- *Saturation correction*: We use multiple frames to tackle over-exposed pixels by essentially weighting them to not have an impact on the reconstruction.
- *Theoretical derivation*: We show how to derive the algorithm in [1] as an instance of incremental EM.

However, before describing further details, we put our contributions in perspective by briefly summarizing related work (for more references see [2]).

1.2. Related Work

Our preliminary work [1] presented an efficient method for MFBD, but this method lacked the ability to perform super-resolution or saturation correction. Moreover, it was only heuristically motivated, while this paper provides a formal motivation via an incremental EM framework. For the MFBD problem itself, there are numerous other papers—e.g., [3, 4, 5, 6, 7]. All of these approach the multiframe problem non-incrementally as opposed to our incremental method. When the number of input frames grows large, such non-incremental approaches rapidly become computationally (and storage-wise) prohibitive. Also related is a somewhat restricted form of MFBD, where usually two or merely a few frames are used [8, 9].

With regards to super-resolution we remark that even though the literature on super-resolution is vast, only very few papers discuss settings where both blind deconvolution and super-resolution are performed simultaneously. Here, 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 (Sec. 4.1 of [11]) that their method becomes unstable for super resolution factors larger than 2.5. In contrast, our model exploits the abundance of data by *not* assuming any image or blur kernel priors, while still robustly handling super resolution factors much larger than 2.5. In fact, in the experiments section we successfully resolve up to a factor of eight the images considered by Šroubek *et al*. [11].

2. BLIND DECONVOLUTION AS INCREMENTAL EM

For notational simplicity we describe our framework using one-dimensional images and point spread functions (PSFs). Our exposition generalizes straightforwardly to two-dimensional images and PSFs—see our technical report [2] for details.

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

First we derive the heuristic method of Harmeling *et al*. [1] 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,

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

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

Jensen's inequality, so that

$$\sum_{t=1}^{T} \log \int p(y_t, f_t|x) df_t \geq \sum_{t=1}^{T} \int q_t(f_t) \log \frac{p(y_t, f_t|x)}{q_t(f_t)} df_t, \quad (2)$$

where each q_t is an arbitrary distribution of f_t. We abbreviate the integrals on the right hand side as $\mathcal{F}_t(q_t, x)$. Then, 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 $\mathcal{F}_t(q_t, x_t)$.

If instead of searching for arbitrary distributions q_t over the PSFs we restrict ourselves to distributions that are concentrated at a single PSF, then we obtain the "winner-take-all" variant of incremental EM [12]. This choice greatly simplifies the E- and M-steps, reducing them to a maximization of $p(y_t, f_t|x)$; a further flat prior on f_t reduces this to maximizing $p(y_t|f_t, x)$ with respect to f_t and x.

Modeling 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,

$$p(y_t|f_t, x) \propto \exp\big(-\tfrac{1}{2}\|y_t - f_t * x\|^2\big). \quad (3)$$

This choice is computationally simple, and leads to E- and M-steps that are (nonnegative) least squares problems of the form

$$\min_{z \geq 0} \quad \|y - Az\|_C^2, \quad \text{where } \|v\|_C^2 = v^T C v. \quad (4)$$

When both A and C are also nonnegative, (4) can be solved via multiplicative updates (see [2]),

$$z_{t+1} = z_t \odot \frac{A^T C\, y}{A^T C A\, z_t}, \quad (5)$$

starting with positive $z_0 > 0$. Here $\odot$ and the fraction bar denote pixel-wise multiplication and division. Note, that we rely crucially on the fact that the convolution $f * x$ can be represented by both matrix-vector products: Fx and Xf.

The M-step requires further discussion. When updating the parameter x, it is beneficial to perform only a few steps towards minimizing $\|y_t - f_t * x\|^2$ with respect to x. Doing so is particularly useful if we consider non-circular convolution, for which x is longer than y_t and as such we have an under-determined system. Performing a few steps also has a regularizing effect.

Summarizing the above discussion we see that the "winner-take-all" variant yields the (generalized EM) updates:

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

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

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

$$\|y_t - f_t * x_t\|^2 \leq \|y_t - f_t * x_{t-1}\|^2, \quad (7)$$

which when applied to each input frame y_t, immediately yield a formal derivation for the two alternating steps proposed by [1].

3. IMPROVED MULTIFRAME BLIND DECONVOLUTION

The basic algorithm described above can recover the underlying, clear image given a sequence of blurry images [1]. Now we extend the basic method and show how to incorporate super-resolution and saturation correction. Our method is simple, but highly effective as it exploits the abundance of data without sacrificing efficiency.

3.1. Super-resolution

Since we are interested in a *single* image x but have *several* observations y_t, despite blurring we could potentially infer a super-resolved image—provided we incorporate change of resolution into the forward model. To this end we define the resizing matrix,

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

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. The matrix 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. This is a favorable property for images, as the number of photons observed should not depend on the resolution. Note that even if m and n are not multiples of each other, D_m^n will interpolate appropriately.

Let n be the length of y. For k times super-resolution we choose x and f large enough so that $f * x$ has length kn. Now we use the image model $p(y_t|f_t, x) \propto \exp\big(-\frac{1}{2}\|y_t - D_n^{kn}(f_t * x)\|^2\big)$, which leads to the modified update steps:

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

$$\|y_t - D_n^{kn}(f_t * x_{t-1})\|^2, \quad (9)$$

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

$$\|y_t - D_n^{kn}(f_t * x_t)\|^2 \leq \|y_t - D_n^{kn}(f_t * x_{t-1})\|^2. \quad (10)$$

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

3.2. Overexposed pixels

Saturated (overexposed) pixels can impact image restoration adversely, particularly so in astronomical imaging where we might want to capture faint stars together with stars that are orders of magnitude brighter. 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 reasonable way to deal with saturated pixels is to weight them out. Since each frame y_t can have different pixels that are saturated (different frames are aligned differently), we must check at each iteration which pixels are saturated. To ignore these pixels we define a weighting matrix,

$$C = \text{Diag}\big([y_t < \rho]\big) \quad (11)$$

where ρ denotes the maximum pixel intensity, and the Iverson brackets $[\cdot]$ apply component-wise. Then, we can write the updates ignoring the saturated pixels simply by replacing the Euclidean norm in (6) and (7) with the weighted norm $\|v\|_C^2 = v^T C v$. 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.

4. EXPERIMENTAL RESULTS

We show now results of two main experiments: (i) comparison against the super-resolution method of [11]; and (ii) MFBD results on astronomical images.

1x 2x 4x 8x Šroubek *et al*. [11]

Fig. 1. Text and disk data: typical example frames (top row), results of our method for blind deconvolution with increasing super-resolution factor compared with Šroubek's results taken from [11] (bottom row from left to right). Already at 2x our results appear to be better than Šroubek's. Note that Šroubek's pictures seem to be postprocessed as their background appears to be darker than that of the input sequence.

4.1. Super-resolution example

The method most closely related to ours is the state-of-the-art *blind* super-resolution approach of Šroubek *et al*. [11]. We compare their method against ours by showing results on some datasets of S. Farsiu and P. Milanfar[1]; we show Šroubek *et al*.'s results as 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 first and third rows of Figure 1 show typical input frames. The second and fourth rows show the result of our method with increasing super resolution factors (one, two, four, and eight times). Already with a factor of two, our results compare favorably with Šroubek's, which is surprising because our method does not depend on any detailed image or blur priors like Šroubek's method. The image obtained using a factor eight super-resolution is clearly superior to the result of [11], whose method could not super-resolve beyond a factor of 2 because of algorithmic instability.

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

4.2. Astronomical imaging

For our astronomical imaging experiment, we used the fast AVT PIKE camera to record a short video (191 frames acquired at 120 fps) of the Trapezium in the Orion constellation. This trapezium is formed by four stars ranging in brightness from magnitude 5 to magnitude 8, with angular separations around 10" to 20".

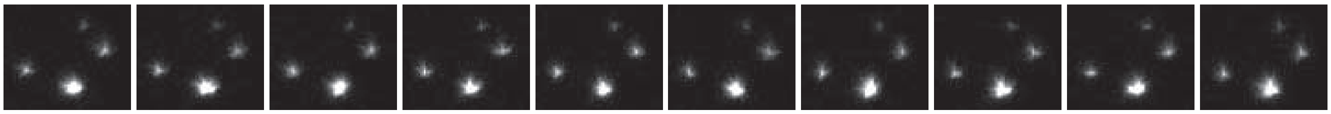

Fig. 2. Orion Trapezium Cluster: example sequence of observed frames, $y_1, \ldots, y_{10}$.

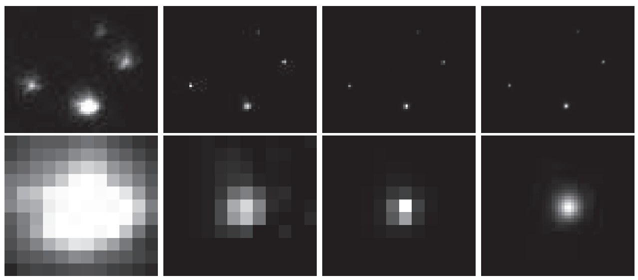

(a) (b) (c) (d)

Fig. 3. Orion Trapezium Cluster: (a) the first observed frame, (b) x_{191} for basic algorithm [1], (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 1. 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.

The individual frames recorded are blurred by atmospheric turbulence—see Figure 2 for sample frames. Beyond blurring, these images often have saturated pixels. We processed the sequence of 191 frames, and show sample results in Figure 3. Here the first row shows from left to right: (a) an enlarged unprocessed frame; (b) the deconvolution results obtained by the basic algorithm of [1]; (c) the result of our method that handles saturated pixels; and (d) the results if we additionally perform four times super-resolution.

Another application that we consider is the estimation of the brightness of stars (Photometry), for which a linear sensor response is required (for our purposes, the used CCD sensor may be assumed linear). The intensity counts can then be translated into stellar magnitudes.[2] Clearly, doing so is not directly possible for stars that saturate the CCD. However, we can use the proposed method for deconvolution with saturation correction to reconstruct the photon counts (image intensities) that we might have recorded were the pixel not saturated. We then convert these counts into stellar magnitudes. For the Trapezium stars we obtain encouraging results—see Table 1.

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

5. CONCLUSION

We presented an incremental expectation maximization method for the multiframe blind deconvolution problem. We showed how to efficiently incorporate important enhancements such as super-resolution and saturation correction, thereby effectively increasing the dynamic range of the sensor. We compared our method against the results of [11] on super resolution benchmark problems. We also showed practical applicability via experiments on astronomical images where saturation correction leads to better reconstructions.

6. REFERENCES

[1] 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.

[2] S. Harmeling, Suvrit Sra, Michael Hirsch, and B. Schölkopf, "An incremental GEM framework for multiframe blind deconvolution, super-resolution, and saturation correction," Tech. Rep. 187, Max Planck Institute for Biological Cybernetics, November 2009.

[3] T. J. Schulz, "Multiframe blind deconvolution of astronomical images.," *J. Optical Soc. Amer.*, vol. 10, pp. 1064–1073, 1993.

[4] R. Vio, J. Nagy, L. Tenorio, and W. Wamsteker, "A simple but efficient algorithm for multiple image deblurring," *Astron. & Astrophy.*, vol. 416, no. 1, 2004.

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

[6] 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.*, vol. 430, no. 2, 2005.

[7] N. F. Law and D. T. Nguyen, "Multiple frame projection based blind deconvolution," *Electronics Lett.*, vol. 31, no. 20, 1995.

[8] A. Rav-Acha and S. Peleg, "Two motion-blurred images are better than one," *Pattern Recognition Lett.*, vol. 26, no. 3, 2005.

[9] J.-F. Cai, H. Ji, C. Liu, and Z. Shen, "Blind motion deblurring using multiple images," *J. Comp. Phy*, 2009.

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

[11] F. Šroubek, G. Cristobál, and J. Flusser, "Simultaneous super-resolution and blind deconvolution," *Journal of Physics: Conference Series*, vol. 124, 2008.

[12] R. M. Neal and G. E. Hinton, "A View of the EM Algorithm that Justifies Incremental, Sparse, and Other Variants," in *Learning in Graphical Models*. MIT press, 1998.