Available online at www.sciencedirect.com

SCIENCE DIRECT®

Pattern Recognition Letters 26 (2005) 311–317

Pattern Recognition Letters

www.elsevier.com/locate/patrec

Two motion-blurred images are better than one

Alex Rav-Acha, Shmuel Peleg *

School of Computer Science and Engineering, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel

Received 24 June 2004; received in revised form 19 October 2004
Available online 9 December 2004

Abstract

Motion blur is a smearing of the image due to a long aperture time. We show that when two motion-blurred images are available, having different blur directions, image restoration can be improved substantially. In particular, the direction of the motion blur and the PSF (Point Spread Function) of the blur can be computed robustly.

Keywords: Motion-blur; Blind deconvolution; Image deblurring; Multi-image restoration

In memorium

This paper is dedicated to Prof. Azriel Rosenfeld, my Ph.D. advisor, to whom I owe my academic career, for better or worse. In spite of the time that has passed since my graduation, I often ask myself how Azriel would do the tasks I am facing. In particular, in my most important duty as an advisor to my students, I try to be as helpful and dedicated to them as Azriel was to me. Co-chairing the 12th ICPR in 1994 was no different, and I tried to do the job as efficiently as Azriel would have done it. And since Azriel liked to write his papers short, I hope that this paper is as short as Azriel would have liked.

(Shmuel Peleg)

1. Introduction

Restoration of images degraded by motion blur, assuming that the blur function is shift invariant and is known, has been studied extensively (Andrews and Hunt, 1977; Jansson, 1997). Early methods restored a single motion-blurred image without prior knowledge of the blur function (blind deconvolution) in cases that the blur function can be characterized by a regular pattern of zeros in the frequency domain, such as in a uniform motion blur (Stockham et al., 1975; Fabian and Malah, 1991). More

* Corresponding author.
E-mail address: peleg@cs.huji.ac.il (S. Peleg).

0167-8655/$ - see front matter © 2004 Published by Elsevier B.V.
doi:10.1016/j.patrec.2004.10.017

recent blind deconvolution methods deal with a wider range of blurs, but use strong assumptions on the image model. Common assumptions are that the image is spatially isotropic (Yitzhaky et al., 1998), or can be modeled as an autoregressive process (Reeves and Mersereau, 1992). Blind deconvolution can also be done by assuming smoothness on both the image and the blur function (Chan and Wong, 1998). A summary and analysis of many methods for "blind deconvolution" can be found in a paper published by Kundur and Hatzinakos (1996). A general motion blur PSF (Point Spread Function) can be recovered from various devices (Liu and Gamal, 2001; Ben-Ezra and Nayar, 2003).

Rather than handling the motion-blur frame by frame, the blur function can be estimated from a sequence of images (Tull and Katsaggelos, 1996a,b; Bascle et al., 1996; Patti et al., 1997). Most of these methods use the multi-image framework to infer the motion blur from the motion analysis. That is, the direction of blur is the same as the direction of motion, and the blur is proportional to the magnitude of motion.

The relation between the inter-frame motion and the motion blur is simple for sequences taken from a stable camera, when the within-frame motion is the same as the inter-frame motion. However, in other cases, such as a sequence taken by a trembling hand, consecutive images may have entirely different blur functions. In particular, the direction of motion blur can be different from one image to another.

We show how multiple images can be used for blind deconvolution when the motion blur functions differ from image to image. We exploit this difference of directions to recover the PSFs of the blur. Moreover, after the motion blur functions have been recovered, using multiple images improves the image restoration which is generally an ill-posed problem.

In a related study published by Shekarforoush and Chellappa (1999), the image restoration algorithm includes an estimation of the PSF from two images. However, it assumes a pure translation between the images, and uses the location of singularities in the frequency domain which are not stable.

In this paper we do not assume a pure image translation but recover a more general 2D parametric motion, e.g. a homography. The blur function is modeled by a one-dimensional PSF. Since aperture time is usually shorter than the frame rate, a one-dimensional translation can approximate the blur, while the inter-frame motion is modeled by a 2D parametric motion.

In a preliminary version of this work (Rav-Acha and Peleg, 2000), the same problem was addressed, but the proposed algorithm was different in many aspects. In particular, instead of recovering the inverse motion blur functions, we now recover the blur functions themselves, and only later restore the original image. This approach has several advantages:

- The blur functions usually have much smaller supports than their inverses, enabling the restoration of wider blurs.
- Both images are used (together with their recovered PSFs) to restore the original image. This results in a better restoration and robustness to noise.
- Regularization is incorporated in the algorithm, providing improved treatment of noise. It is easier to incorporate regularization when the recovery of the blur functions and the image restoration are done separately.

The recovery of the motion blur PSFs and the image restoration given these PSFs are described in Sections 2 and 3. In Sections 4 and 5 we show how to align the images and how to recover the directions of the motion blur PSFs. Finally, examples are presented which demonstrate the effectiveness of the proposed method.

2. Computing the motion blur functions

At this stage we will assume that the directions of the motion blurs are known, and that the two blurred images are aligned. Our methods for image alignment and for recovering the blur directions are presented in Sections 4 and 5.

Let g_i denote the observed image, degraded by motion blur with a one dimensional PSF $m_i =$

$(m_i(1),\ldots,m_i(K))$ at an angle α_i. Let f be the original image. We assume that g_i relates to f in the following way:

$$g_i(x,y)=\sum_{k=0}^{K-1} m_i(k)\cdot f(x+k\cos(\alpha_i),y+k\sin(\alpha_i)) \tag{1}$$

For simplicity, we will denote such a convolution in a certain angle by $g_i=f\overset{\alpha_i}{*}m_i$.

This model is valid when the blur function is one dimensional and shift-invariant. Otherwise, the image can be divided into regions having approximately the same blur. For a discrete image f, interpolation is used to describe gray levels at fractional locations.

The blurred input images g_1 and g_2 relate to the "ideal" image f blurred by the kernels m_1 and m_2 by:

$$g_1=f\overset{\alpha_1}{*}m_1,\quad g_2=f\overset{\alpha_2}{*}m_2 \tag{2}$$

Since convolution is commutative, applying the blur function of the first image on the second image and applying the blur function of the second image on the first image yield the same result. We use this observation to obtain a linear equation per each pixel (excluding the boundaries):

$$\left(g_1\overset{\alpha_2}{*}m_2\right)(x,y)=\left(g_2\overset{\alpha_1}{*}m_1\right)(x,y)$$

Recovering the motion blur PSFs m_1 and m_2 is done by minimizing the following error function on the region of analysis R:

$$E(m_1,m_2)=\sum_{x,y\in R}\left[\left(g_1\overset{\alpha_2}{*}m_2\right)(x,y)-\left(g_2\overset{\alpha_1}{*}m_1\right)(x,y)\right]^2 \tag{3}$$

Calculating the derivatives of this error function with respect to m_1 and m_2 yields a set of linear equations with K_1+K_2 unknowns, where K_1 and K_2 are the supports of the PSFs of the blurred images g_1 and g_2. We usually use PSFs of sizes 15–30 pixels (depending on the image size).

To avoid the trivial solution and to increase the stability of the solution, we add the constraint that both PSFs sum to 1, to ensure that the blur functions are energy preserving. Other constraints can also be added, such as the smoothness of the blur functions. By solving the simple linear system we recover the PSFs of the motion blurs. These PSFs are used in the image restoration stage described next.

3. Image deblurring

Deblurring images with a known blur function is commonly done using the Wiener filter. However, better results are usually obtained with iterative methods in the spatial domain. We use an iterative deconvolution that can easily be applied for more than one image, similar to the IBP method (Irani and Peleg, 1991). The deblurred image $\hat{f}$ will minimize the following error function:

$$E=\sum_{i=1}^{2}\left\|g_i-\hat{f}\overset{\alpha_i}{*}m_i\right\|^2 \tag{4}$$

Based on the gradient decent method we get the following iterative scheme:

- Initialize the restored-image $\hat{f}$ to be the average of the input images $\hat{f}\leftarrow\frac{1}{2}(g_1+g_2)$.
- Iteratively update $\hat{f}$ using the gradient of E:

  $$\hat{f}\leftarrow\hat{f}+\frac{1}{N}\sum_{i=1}^{2}m_i^T\overset{\alpha_i}{*}\left(g_i-\hat{f}\overset{\alpha_i}{*}m_i\right)$$

  where m_i^T is the flipped version of m_i.

When only a single image is used in the deconvolution, it is very important to add a regularization term to avoid noisy results, as deconvolution is an ill-posed problem. Given two images or more, the restoration problem becomes well-posed, yet a small regularization term can improve the quality of the restored images. We have added the regularization term: $L=(\|f_x\|_p)^p+(\|f_y\|_p)^p$ to the error function in Eq. (4), where f_x and f_y are the derivatives of f in the horizontal and vertical directions, and $(\|f\|_p)^p=\sum_{x,y\in R}|f(x,y)|^p$. The recursive updating formula is then changed to:

$$\hat{f}\leftarrow\hat{f}+\frac{1}{N}\sum_{i}m_i^T\overset{\alpha_i}{*}\left(g_i-\hat{f}\overset{\alpha_i}{*}m_i\right)-\lambda\left.\frac{\partial L}{\partial f}\right|_{\hat{f}}$$

The last term denotes the derivative of the regularization function L in the point $\hat{f}$, and is given by:

$$\mathrm{d}x^{\mathrm{T}} * \left[p\hat{f}_x^{\,p-1} \cdot \mathrm{sign}\left(\hat{f}_x\right)\right] + \mathrm{d}y^{\mathrm{T}} * \left[p\hat{f}_y^{\,p-1} \cdot \mathrm{sign}\left(\hat{f}_y\right)\right]$$

Where $\mathrm{d}x^{\mathrm{T}}$ and $\mathrm{d}y^{\mathrm{T}}$ are the flipped versions of the derivative kernels, and the operations in the rectangular brackets are all point-wise operations. In all our tests we have used $\lambda = 0.004$ and $p = 1.25$. We used this norm as the L_2 norm over smoothes the image, and the L_1 norm produces stair-casing artifacts. Other regularization methods may be used, e.g. the Total Variation regularization (Chan and Wong, 1998).

Using both images is theoretically superior to a deconvolution of a single image since it can restore frequencies which were lost by the motion blur degradation. An example is given in Fig. 1. This demonstrates the limits of any method which is based on a single frame—even when the motion blur PSF is recovered successfully, the blurred image still lacks all the information needed for image restoration.

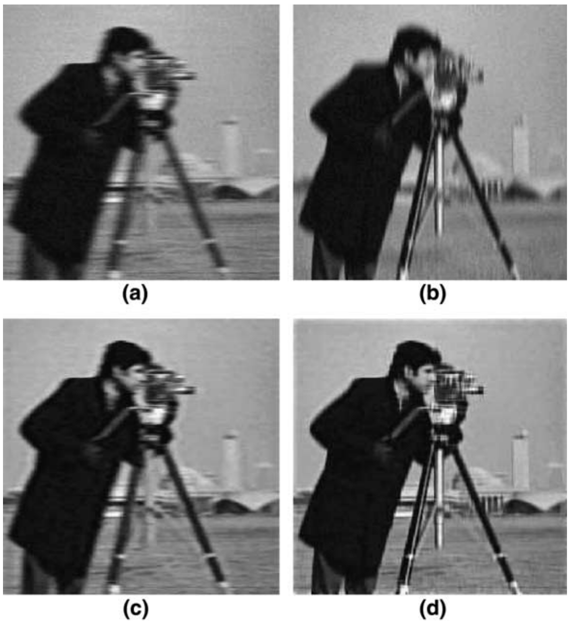

Fig. 1. The benefit of using multiple images for deconvolution. An image was degraded synthetically with uniform horizontal (a) and vertical (b) motion blur, and with some Gaussian noise. (c) The result of deconvolution using only image (a). (d) Restoration using simultaneously both (a) and (b).

Naturally, the benefit of using multiple images is most prominent for blur functions which have many zeros in their frequency response, such as uniform blur, and less prominent for invertible PSFs.

Based on these observations it is clear that the restoration is improved as the angle between the directions of the blur functions become orthogonal. In this case, the degradations are mostly independent.

4. Image alignment

Having a pair of blurred images g_1 and g_2 which may not be aligned, we follow the previous notation and write:

$$g_1 = f \overset{\alpha_1}{*} m_1, \quad g_2 = T(f) \overset{\alpha_2}{*} m_2 \tag{5}$$

Where T is a 2D image transformation aligning the two images.

Using multiple images for deblurring requires alignment between them. Accurate image translation is not needed since the translation can be viewed as a part of the motion blur PSF, convolved with a shifted delta impulse. We only need to compensate for large translations, and for affine transformations. We found that most images have a dominant phase (i.e.—the blur function is not entirely uniform) and conventional multi-resolution image alignment methods (Bergen et al., 1992) can be used successfully. For highly degraded images, a regularization term which favors small motions can be added to enforce convergence Fig. 2.

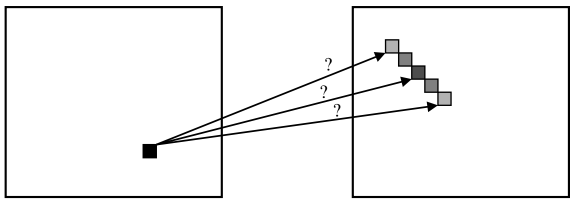

Fig. 2. With motion blur, the correspondence between images is fuzzy. It can be described by the convolution matrix that transforms the left image into the right image.

5. The direction of motion blurs

Most methods for recovering the direction of motion blur assume that the image is isotropic or can be modeled by an auto-regressive process. For example, it was suggested (Yitzhaky et al., 1998) that the spectrum energy of the image derivative in the direction of the motion blur is smaller than in other directions. The isotropic assumption is problematic for many scenes, such as in urban areas.

Instead, we use the basic scheme described in Section 2 (solving for the motion blur PSFs) to recover the blur directions. The directions are determined using an exhaustive search over the angles of the blur functions of both images: For each possible value of α_1 and α_2, we solve a set of linear equations, and compute the residual error defined in Eq. (3). This error is minimized for the actual blur angles of both images. This scheme is computationally intensive, so we slightly modify it:

- We allow smaller supports for the motion blur functions than the ones used to actually recover the PSFs. we used PSFs which were slighlty larger than half size of the ones used for the recovery. This modification both improves the run time and prevents the false positives obtained from inverting the role of the blurs (In this case, the solution of the linear system will give the inverse PSFs, which require much larger supports).
- We use a multi-resolution framework: A Gaussian pyramid is constructed for both images. An initial estimate of the blur functions is calculated using the smoothed and sub-sampled images, and is then refined using the high-resolution images. Working on a sub-sampled version of images is much faster, since both the sizes of the images and the supports of their blur functions are reduced by a factor of 2 for each level of the pyramid.

6. Experimental results

We would like to check the influence of the angle between the directions of the two motion blurs, on the power of the proposed method. Fig. 3 depicts the restoration error as a function of the angle between the directions of the motion

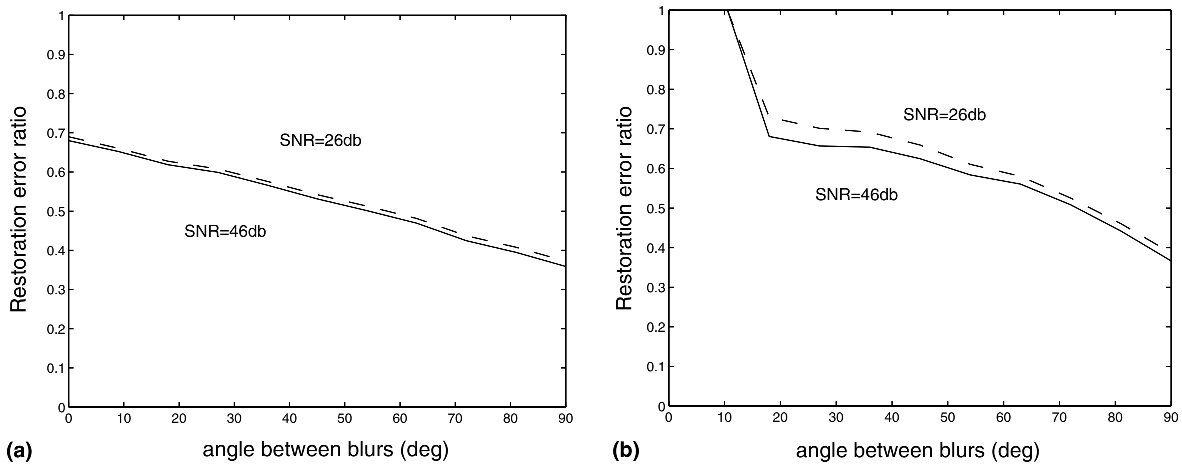

Fig. 3. An image was blurred in various directions, and contaminated with white noise (SNR = 26 db, 46 db). Then, the original image was restored from different pairs of images. The SSD (sum of squared differences) error between the original and the restored image was computed. For normalization reasons, the error function used here is the ratio between the SSD error for the restored image, and the SSD error for the blurred one (Thus, a ratio of one means no improvement). This error is plotted as a function of the angle between the directions of blurs. (a) The motion blur functions are known (both directions and PSFs). (b) The motion blur functions are computed by our algorithm.

blurs in the two input images. For this analysis, an image was blurred in various directions, and further contaminated with white noise. It can be seen from Fig. 3(a) that better restoration is achieved as the directions of motion blurs are closer to perpendicular. This is consistent with the idea that more information on the original image is preserved as the blur functions are more orthogonal.

The restoration error does not change when the motion blur PSFs are unknown (and are recovered using the proposed method) but their directions are known. On the other hand, if the motion blur directions are not given, the algorithm becomes unstable for small angles. Nevertheless, it is clear that the current algorithm handles successfully a wide range of angles, not only perpendicular blurs.

Figs. 4 and 5 demonstrate the process on real images taken by a video camera. Fig. 4 shows a frontal view of a hotel wall. Fig. 4(a) and (b) were taken with fast pan and tilt of the camera. The motion blur induced from the pan or tilt can fit our motion blur model (shift invariant blur) since the focal length was large. The restored image is shown in Fig. 4(c). Fig. 5 shows a blind restoration of a printed document. The enhancement of the image is evident from both the edges of the digits and the letters.

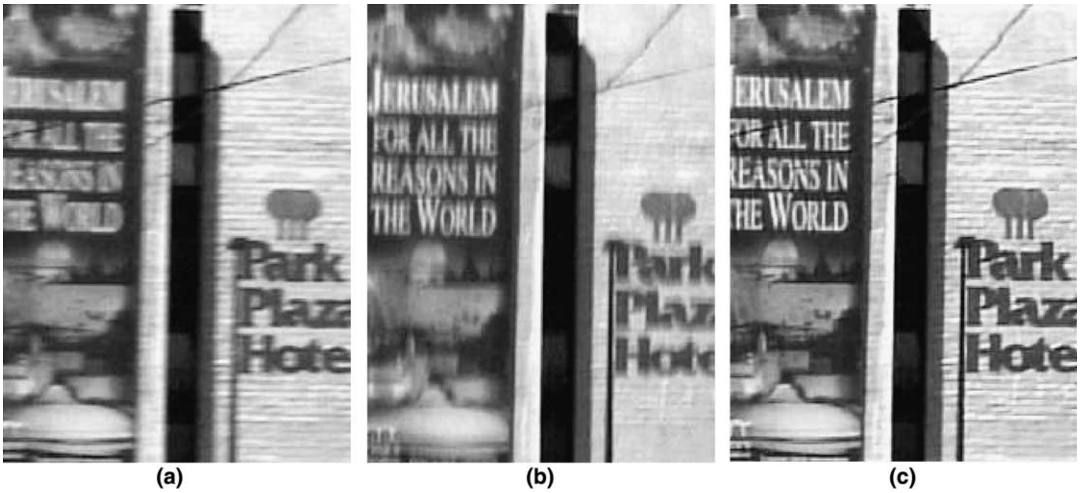

Fig. 4. An example of recovering two out-doors blurred images. (a) and (b) were degraded by horizontal and vertical motion blur, respectively, due to the fast panning and tilting of the hand. (c) The restored image.

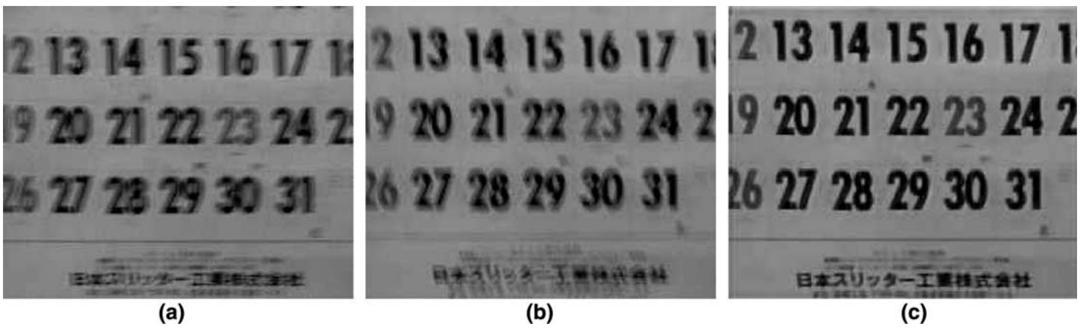

Fig. 5. An example of recovering two real blurred images. (a) and (b) were degraded by motion blur in the directions ≈−10° and ≈65° from the x axis respectively. (c) The restored image.

7. Concluding remarks

Two images of the same scene, having motion blur in different directions, prove to preserve large amount of information of the original scene. A simple and yet effective method for recovering this information is presented. This method does not require a prior knowledge regarding to the blur PSF or even its direction, and does not assume any relation between the image displacement and the motion blur. Such assumptions are a necessity for most existing methods for image restoration.

Due to the motion blur, the motion parameters are pre-computed, up to a small translation. Then the PSFs of the two images are recovered simultaneously, and the image is restored using an iterative scheme.

The strength of the algorithm (together with its weakness) is its reliance on multiple images blurred in different directions. As a results, this method cannot be used, for example, for sequences taken from a driving car, where the motion blur degradation is always in the same direction. On the other hand, it is very effective for images blurred due to hand tremble, where most assumptions about the relations between motion and motion blur fail.

The proposed algorithm can be generalized to handle a sequence with an arbitrary number of images. For example, three images of the same scene, blurred in directions that are in 60° one from the other, can be better enhanced simultaneously, rather than using each pair of the three.

References

Andrews, H., Hunt, B., 1977. Digital Image Restoration. Prentice Hall.

Bascle, B., Blake, A., Zisserman, A., April 1996. Motion deblurring and super-resolution from an image sequence. In: European Conference on Computer Vision. Vol. 2. Cambridge UK, pp. 573–582.

Ben-Ezra, M., Nayar, S.K., June 2003. Motion deblurring using hybrid imaging. In: IEEE Conference on Computer Vision and Pattern Recognition. Madison, Winsconsin, pp. 454–460.

Bergen, J., Anandan, P., Hanna, K., Hingorani, R., May 1992. Hierarchical model-based motion estimation. In: Second European Conference on Computer Vision. Santa Margherita Ligure, Italy, pp. 237–252.

Chan, T., Wong, C., 1998. Total variation blind deconvolution. IEEE Trans. Image Process. 7 (3), 370–375, March.

Fabian, R., Malah, D., 1991. Robust identification of motion and out-of-focus blur parameters from blurred and noisy images. Comput. Vision, Graphics, Image Process. 53 (5), 403–412, September.

Irani, M., Peleg, S., 1991. Improving resolution by image registration. CVGIP: Graph. Models Image Proces. 53, 231–239, May.

Jansson, P.A., 1997. Deconvolution of Image and Spectra, second ed. Academic-Press.

Kundur, D., Hatzinakos, D., 1996. Blind image deconvolution. IEEE Signal Process. Mag. 13, 43–64, May.

Liu, X., Gamal, A., 2001. Simultaneous image formation and motion blur restoration via multiple capture. In: International Conference on ASSP. Salt Lake City, Utah pp. 1841–1844.

Patti, A., Sezan, M., Tekalp, M., 1997. Superresolution video reconstruction with arbitrary sampling lattices and nonzero aperture time. IEEE Trans. Image Process. 6 (8), 1064–1078, August.

Rav-Acha, A., Peleg, S., December 2000. Restoration of multiple images with motion blur in different directions. In: Workshop on the Application of Computer Vision. Palm Springs pp. 22–28.

Reeves, S., Mersereau, R., 1992. Blur identification by the method of generalized cross-validation. IEEE Trans. Image Process. 1, 301–311, July.

Shekarforoush, H., Chellappa, R., 1999. Data-driven multichannel superresolution with application to video sequences. J. the Optical Soc. Am. A 16 (3), 481–492, March.

Stockham, T., Cannon, T., Ingebretsen, R., 1975. Blind deconvolution through digital signal processing. Proc. of the IEEE 63, 678–692.

Tull, D., Katsaggelos, A., 1996a. Iterative restoration of fast-moving objects in dynamic image sequences. Opt. Engng 35 (12), 3460–3469, December.

Tull, D., Katsaggelos, A., September 1996b. Regularized blur-assisted displacement field estimation. In: IEEE International Conferance on Image Processing. Lausanne, Switzerland pp. 85–88.

Yitzhaky, Y., Mor, I., Lantzman, A., Kopeika, N., 1998. Direct method for restoration of motion blurred images. J. Opt. Soc. Am. A 15 (6), 1512–1519, June.