ISBN 978-1-4020-8922-0

NATO Science for Peace and Security Series - B:
Physics and Biophysics

Unexploded Ordnance
Detection and Mitigation

Edited by
Jim Byrnes

Springer

| NATO OTAN | This publication is supported by: | The NATO Science for Peace and Security Programme |

Recent Advances in Space-Variant Deblurring and Image Stabilization

Michal Šorel[*,1], Filip Šroubek[*] and Jan Flusser[*]

Abstract The blur caused by camera motion is a serious problem in many areas of optical imaging such as remote sensing, aerial reconnaissance or digital photography. As a rule, this problem occurs when low ambient light conditions prevent an imaging system from using sufficiently short exposure times, resulting in a blurred image due to the relative motion between a scene and the imaging system. For example, the cameras attached to airplanes and helicopters are blurred by the forward motion of the aircraft and vibrations. Similarly when taking photographs by hand under dim lighting conditions, camera shake leads to objectionable blur. Producers of imaging systems introduce compensation mechanisms such as gyroscope gimbals in the case of aerial sensing or optical image stabilization systems in the case of digital cameras. These solutions partially remove the blur at the expense of higher cost, weight and energy consumption. Recent advances in image processing make it possible to remove the blur in software. This chapter reviews the image processing techniques we can use for this purpose, discusses the achievable performance and presents some promising results achieved by the authors.

Keywords: Camera shake, image stabilization, image registration, space-variant restoration, deblurring, blind deconvolution, point spread function, regularization

1 Introduction

The blur caused by sensor motion is a serious problem in a large number of applications from remote sensing to landmine detection to amateur photography. In general, this problem occurs if the time needed to capture an image is so long that the imaging system moves relative to the scene.

[*,1] Institute of Information Theory and Automation of the ASCR, Pod Vodarenskou vezi 4, Praha 8, CZ-18208, Czech Republic, e-mail: {sorel,sroubekf,flusser}@utia.cas.cz

An example application in landmine detection is the general survey of minefields in the aftermath of military conflicts using visible light or infrared cameras. The cameras attached to airplanes and helicopters are blurred by the forward motion of the aircraft and vibrations. While the vibrations can be dumped to some extent using gyroscope stabilizers, there is no simple way to do the same with the forward movement. A similar problem arises in the case of cameras attached to moving vehicles. For example, thermal infrared cameras attached to armoured vehicles can be used to detect anti-personnel and anti-tank mines on roads and tracks.

Similarly, when taking photographs under low light conditions, the camera needs a long exposure time to gather enough light to form the image, which leads to objectionable blur. To mitigate this problem, producers of digital cameras introduced two types of hardware solutions. The technically simpler one is to increase the sensitivity of a camera (ISO) by amplifying the signal from the sensor, which permits faster shutter speed. Unfortunately, especially in the case of compacts, this results in a decrease of image quality because of more noise. Optical image stabilization (OIS) systems, containing either a moving image sensor or an optical element to counteract camera motion, are technologically more demanding. They help to remove blur without increasing noise level but at the expense of higher cost, weight and energy consumption.

A system removing the blur in software would be an elegant solution to the problem. In this chapter we give an overview to possible approaches to this problem. The algorithms are explained in connection with photography but the results can be applied to other cases such as aerial reconnaissance and infrared imaging as well.

We start with an outline of approaches. Then, in Section 3 we describe a mathematical model of blurring. For each approach (Sections 4–7), we summarize its strong and weak points and present a typical state-of-the-art method. Section 8 summarizes results and indicates the potential of individual approaches.

2 Overview of Approaches

An obvious way to avoid camera motion blur is to take a sequence of underexposed images so that the exposure time is short enough to prevent blurring. After registration, the whole sequence can be summed to get the original sharp image with a reasonable noise level. In Section 4 we briefly discuss why this idea turns out to be impractical for more than a few images. In the rest of this chapter, we discuss situations where we already have a blurred image (or a sequence of images) and wish to remove the blur.

To simplify the problem, the blur is usually assumed to be homogenous in the whole image. In this case the blur can be modeled by convolution. That is why the reverse problem to find the sharp image is called *deconvolution*. If the PSF is not known, which is the case in most real situations, the problem is called *blind deconvolution*.

While non-blind deconvolution problems can be easily solved, solutions of blind deconvolution problems from a single image are highly ambiguous. To find a stable solution some additional knowledge is required. This case is treated in Section 5. The most common approach is regularization, applied both on the image and blur. Regularization terms mathematically describe *a priori* knowledge and play the same role as prior distributions in stochastic models. For the present, probably the best published blind deconvolution methods are those of Fergus et al. [2] and coming soon [9].

Another approach, extensively studied in past years, is to use multiple images capturing the same scene but blurred in a different way (Section 6). The camera takes two or more successive images and each exhibits different blurring due to the basically random motion of the photographer's hand or, for example, aircraft vibrations. Multiple images permit estimation of the blurs without any prior knowledge of their shape, which is hardly possible in single image blind deconvolution [10].

One particular multi-image setup attracted considerable attention only recently. Taking images with two different exposure times (long and short) results in a pair of images, in which one is sharp but underexposed and another is correctly exposed but blurred. Instead of the underexposed image we can equivalently take an image with high ISO. Both can be easily achieved in continuous shooting mode by exposure and ISO bracketing functions of DSLR cameras. For Canon compact cameras these functions can be written in the scripting language implemented within the scope of the CHDK project (http://chdk.wikia.com/wiki/CHDK).

To estimate the sharp image, two different ideas were proposed in the literature. The first adjusts the contrast of the underexposed image to match the histogram of the blurred one [7]. However, this technique is applicable only if the difference between exposure times is small. The second way [5, 11] uses the image pair to estimate the blur and then deconvolves the blurred image. This path was followed by [15], where the authors show an effective way to suppress ringing artifacts produced by Richardson-Lucy deconvolution. In Section 7 we give an example of an algorithm of this type proposed by the authors of this chapter. To be applicable even for wide angle lenses, we consider space-variant blur.

3 Blur Model

It is well known that homogenous blurring can be described by *convolution*

$$\mathbf{z} = \mathbf{u} * \mathbf{h} \, [x, y] = \int \mathbf{u}(x - s, y - t) \mathbf{h}(s, t) \, ds dt, \tag{1}$$

where $\mathbf{u}$ is an original image, $\mathbf{h}$ is called the *convolution kernel* or *point-spread function* (PSF) and $\mathbf{z}$ is the blurred image. In our case of camera motion blur the PSF is a plane curve given by an apparent motion of each pixel during the exposure.

If the focal length of the lens is short or camera motion contains a significant rotational component about the optical axis, this simple model is not valid. The blur

is then different in different parts of the image and is a complex function of camera motion and depth of scene [14]. We can see an example in Figure 5, where the image was divided into 49 (7×7) rectangles and convolution kernels were estimated within these subimages (by a method described in Section 7). Notice for example the difference between the upper left and right kernels.

Nevertheless, this spatially varying blur can be described by a more general linear operation

$$\mathbf{z} = \mathbf{u} *_v \mathbf{h}\,[x,y] = \int \mathbf{u}(x-s,y-t)\mathbf{h}(x-s,y-t;s,t)\,dsdt, \qquad (2)$$

where $\mathbf{h}$ is again called the *point-spread function* as in the case of convolution. Note that convolution is a special case, with the function h independent of coordinates x and y, that is $\mathbf{h}(x,y;s,t) = \mathbf{h}(s,t)$. We can look at (2) as convolution with a kernel that changes with its position in the image, and speak about *space-variant convolution*. The subscript v distinguishes from ordinary space-invariant convolution, denoted by asterisk.

Because the rotational component of camera motion is usually dominant, the blur is independent of depth and the PSF changes in a continuous gradual way. Therefore the blur can be considered locally constant and can be locally approximated by convolution. This property can be used to efficiently estimate even the space-variant PSF, as described in Section 7.

4 Summing of Underexposed Images

At first sight, the idea to sum a sequence of underexposed images seems to be very attractive. It is a well known property of shot (Poisson) noise that an image taken with an exposure time t has the same level of noise as the sum of N images each taken with time t/N. So, apparently, the only problem we must solve is to register images with sufficient precision. There exist many fast image registration methods and, without doubt, one of them could be used in this case. Registration is made easier also by the fact that the difference between images is not large as the images are taken quickly one after another.

Unfortunately, for the present, there is a serious problem that limits the use of this idea in practice. Images taken by present day digital cameras are huge and it takes a lot of time to read them out from sensor to camera memory. For consumer level DSLRs it typically takes about $1/3$ of second, for compacts even more. For example, imagine that we want to replace one $1/4$ s image by a sequence of 16 images taken with exposure time $1/60$ s, which corresponds to the use of ISO $1,600$ instead of ISO 100. Now the camera needs $16 \times 1/3$, or more than 5 s. For many situations this is simply too long.

To summarize, on one hand this approach is computationally simple and can potentially be implemented inside a camera. On the other hand, to be useful for really low lighting conditions, the read-out time will have to be significantly shortened.

In the rest of this chapter we will treat blurred images, which is less demanding with respect to read-out time and can actually be used with present day cameras. On the other hand, deblurring is computationally more time consuming and assumes postprocessing on the photographer's personal computer.

5 Single-Image Blind Deconvolution

There has been a considerable effort in the image processing community in the last three decades to find a reliable algorithm for single image blind deconvolution. For a long time, the problem seemed too difficult to be solved for complex blur kernels. Proposed algorithms usually worked only for special cases such as astronomical images with uniform (black) background. There was no reliable result applicable to natural scenes.

Only recently, in 2006, Rob Fergus et al. [2] proposed an interesting Bayesian method with very impressive results. Another method of this kind should appear at SIGGRAPH 2008 [9]. The authors claim even better results than [2] with much simpler and faster computation. In this chapter we briefly describe the method [2].

The method assumes a simple convolution model of blurring

$$\mathbf{z} = \mathbf{u} * \mathbf{h} + \mathbf{n}, \tag{3}$$

where $\mathbf{n}$ is an independent Gaussian zero mean noise.

The basic idea is to estimate the *a posteriori* probability distribution of the gradient of the original image and of the blur kernel

$$p(\mathbf{u}, \nabla\mathbf{h}|\nabla\mathbf{z}) = p(\nabla\mathbf{z}|\nabla\mathbf{u}, \mathbf{h})p(\nabla\mathbf{u})p(\mathbf{h}), \tag{4}$$

using knowledge of independent prior distributions of the image gradient $p(\nabla\mathbf{u})$ and of the kernel $p(\mathbf{h})$. The likelihood $p(\nabla\mathbf{z}|\nabla\mathbf{u}, \mathbf{h})$ is considered Gaussian with mean $\nabla\mathbf{u} * \mathbf{h}$ and an unknown variance. After estimation of the full posterior distribution $p(\mathbf{u}, \nabla\mathbf{h}|\nabla\mathbf{z})$, it computes the kernel with maximal marginal probability. Finally, the original image is restored by the classical Richardson-Lucy algorithm. This final phase could obviously be replaced by an arbitrary non-blind deconvolution method.

The algorithm is quite complex. It approximates the full posterior distribution by the product $p(\mathbf{u}|\nabla\mathbf{z})p(\nabla\mathbf{h}|\nabla\mathbf{z})$ in the sense of Kullback-Leibler distance, which can be efficiently computed by the variational scheme described in [6] for cartoon images. The image gradient prior is considered in the form of a Gaussian mixture. In a similar way, the prior on kernel values is expressed as a mixture of exponential distributions, which reflects the fact that most kernel values for motion blur are zero. Both types of priors are learned from a typical natural image.

Figure 1 shows an example of an image restored by this method. We can see that the convolution kernel is recovered surprisingly well. Some artifacts appear because there are no smoothing constraints in the algorithm. Another problem is the high number of artifacts produced by non-blind deconvolution in the final phase of the

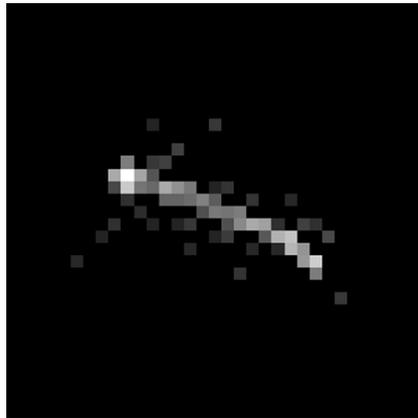(a) Blurred image, 800×600 pixels

(b) Deblurred image

(c) Estimated PSF, 35×35 elements

Fig. 1 Example results of single-image blind deconvolution provided by Fergus et al. [2].

algorithm. A typical example is the well known ringing effect. New papers [9,15,16] seem to deal with this problem successfully.

Bringing this all together, there are reliable methods for estimating the blur kernel and subsequent restoration from a single blurred image. The main problem is the need for user assistance to choose a suitable part of the image for kernel inference.

6 Multi-image Blind Deconvolution

In this approach we use multiple images (Fig. 2a) capturing the same scene but blurred in a different way. We can easily take such a sequence using continuous shooting modes of present day cameras. Multiple images permit one to estimate the blurs (Fig. 3) without any prior knowledge of their shape.

(a) Blurred input images, 1024×768 pixels

(b) Deconvolution from only first image   (c) Result of multi-image deconvolution

Fig. 2 Example results achieved by multi image blind deconvolution algorithm [10].

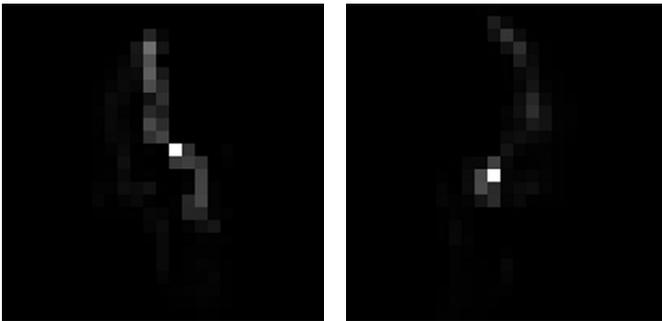**Fig. 3** Convolution kernels corresponding to images in Figure 2a.

Mathematically, the situation is described as convolution of the original image **u** with P convolution kernels $\mathbf{h}_p$

$$\mathbf{z}_p = \mathbf{u} * \mathbf{h}_p + \mathbf{n}_p, \quad p = 1, .., P. \tag{5}$$

In this section, we describe one of the best working multi-image deblurring algorithms [10].

As in the single image situation, the algorithm can be viewed as a MAP (maximum a posteriori) estimate of distributions of the sharp image and the blur kernels. It is equivalent to minimization of the functional

$$E(\mathbf{u}, \mathbf{h}_1, ..., \mathbf{h}_p) = \frac{1}{2} \sum_{p=1}^{P} \|\mathbf{u} * \mathbf{h}_p - \mathbf{z}_p\|^2 + \lambda_u Q(\mathbf{u}) + \lambda_h \sum_{i \neq j} R(\mathbf{h}_i, \mathbf{h}_j) \qquad (6)$$

with respect to the latent image $\mathbf{u}$ and blur kernels $\mathbf{h}_1, ..., \mathbf{h}_p$. The first term of (6), called the *error term*, is a measure of the difference between input blurred images $\mathbf{z}_p$ and the original image $\mathbf{u}$ blurred by kernels $\mathbf{h}_k$. The size of the difference is measured by the L_2 norm $\|.\|$. The inner part of the error term is nothing more than the matrix of errors at the individual points of image p, which should be close to zero for the correct image and kernel. Note that kernels $\mathbf{h}_p$ incorporate a possible shift of the camera between the images.

The role of regularization terms

$$Q(\mathbf{u}) = \int |\nabla \mathbf{u}| \qquad (7)$$

and

$$R(\mathbf{h}_i, \mathbf{h}_j) = \|\mathbf{z}_j * \mathbf{h}_i - \mathbf{z}_i * \mathbf{h}_j\| \qquad (8)$$

is to make the problem well-posed and incorporate prior knowledge about the solution [12].

Thus, $Q(\mathbf{u})$ is an image regularization term which can be chosen to properly represent the expected character of the image function. For the majority of images a good choice is total variation (7), where $\nabla \mathbf{u}$ denotes the gradient of $\mathbf{u}$. The size of the gradient is integrated over the whole area of the image. Very good anisotropic denoising properties of the total variation were shown by Rudin et al. [8]. A reason why total variation works so well for real images is that it favors piecewise constant functions. In real images object edges create sharp steps that appear as discontinuities in the intensity function. For a more detailed discussion of image regularization, see [1, 10, 13]. The kernel regularization term is a constraint useful for kernels of limited support.

The functional (6) is minimized by alternating minimization in the subspaces corresponding to the image and the blur kernels.

The main problem of the multi-image approach is speed. For this reason, it is practically impossible to generalize this approach to space-variant blur. As a result, this approach can be applied mainly for tele-lens photos if the rotational component of camera motion about the optical axis is negligible. In general, it usually works for the central section of an arbitrary blurred image.

7 Restoration from a Pair of Blurred and Noisy Images

The idea to use two images with two different exposure times appeared only recently [5, 11, 15]. Most algorithms of this group first estimate the blur from the image pair and then deconvolve the blurred image. The main problem of the deconvolution phase is suppression of ringing artifacts. A method of handling this problem for the Richardson-Lucy algorithm was proposed in [15, 16].

None of the aforementioned methods are general enough to be applicable to full uncropped photos. The reason is that the blur is not constant throughout the image, especially in the case of lenses with shorter focal length (<50 mm). In addition, it often happens that camera motion has a considerable rotational component about the optical axis and then the blur is space-variant, even for tele-lenses. Another effect modifying blurs is lens distortion. All these effects are accentuated in regions close to image borders. Therefore a space-variant approach is necessary for artifact-free results.

Space-variant restoration was already considered in astronomy and microscopy but there is almost no work applicable in photography. Only recently, in [14], is space-variant blur considered for a camera moving without rotation, but this assumption does not correspond to the real trajectory of a handheld camera.

In the following paragraphs we describe a state-of-the-art algorithm proposed by the authors. To avoid ringing effects we use a constrained least squares method with total variation regularization. To be applicable even for wide angle lenses, we consider space-variant blur.

7.1 Algorithm

For input the algorithm requires a pair of images, one of them blurred and another noisy but sharp. The algorithm works in three phases:

1. Robust image registration
2. Estimation of convolution kernels on a grid of windows followed by an adjustment at places where estimation failed
3. Restoration of the sharp image

In the first step, we need a robust registration procedure working with precision significantly better than the considered size of blur kernels. We can assume that the change of camera position is negligible with respect to scene distance (very short baseline) and consequently it can be approximated by a projective transform independent of scene depth. Experiments have also shown that misalignments due to lens distortion do not harm the algorithm because they are compensated by the shift of the corresponding part of the space-variant PSF. For the purpose of this algorithm, we apply the standard RANSAC [3, 4] approach to estimate the homography matrix. Then we transform the blurred image accordingly. The transformed image will be denoted by $\mathbf{z}^T$.

In the second step of the algorithm we make use of the fact that the blur can be locally approximated by convolution. We do not estimate the blur kernels in all pixels. Instead, we divide the image into rectangular windows (a 7×7 grid in our example in Figure 6) and estimate only a small set of kernels $\mathbf{h}_{i,j}$ ($i, j = 1.7$ in our example in Figure 5). The estimated kernels are assigned to centers of the windows where they were computed. In the rest of the image, the PSF $\mathbf{h}$ is approximated by bilinear interpolation from blur kernels in four adjacent windows.

Thus, we estimate blur kernels on a grid of windows, where the blur can be approximated by convolution

$$\mathbf{z}_{i,j}^T = (\mathbf{u}_{i,j} - \mathbf{n}_{i,j}) * \mathbf{h}_{i,j} = \mathbf{u}_{i,j} * \mathbf{h}_{i,j} - \mathbf{n}_{i,j} * \mathbf{h}_{i,j}, \tag{9}$$

where $\mathbf{z}_{i,j}^T$ is a section of the transformed blurred image $\mathbf{z}^T$, $\mathbf{u}_{i,j}$ the corresponding part of the noisy image, $\mathbf{h}_{i,j}$ the locally valid convolution kernel and $\mathbf{n}_{i,j}$ an independent Gaussian noise contained in the noisy image.

We estimate the solution of this problem in a least squares sense as

$$\mathbf{h}_{i,j} = \arg\min_{\mathbf{k}} \|\mathbf{u}_{i,j} * \mathbf{k} - \mathbf{z}_{i,j}^T\|^2 + \alpha \|\nabla\mathbf{k}\|^2, \qquad \mathbf{k}(s,t) \geq 0, \tag{10}$$

where $\mathbf{h}_{i,j}(s,t)$ is an estimate of $h(x_0, y_0, s, t)$, (x_0, y_0) being the center of the current window $\mathbf{z}_{i,j}$, and $\|.\|$ is the L_2 norm. Regularization helps reduce the noise arising from the imprecise model.

The kernel estimation procedure (10) can fail. Such kernels must be identified, removed and replaced by the average of adjacent (valid) kernels. There are basically two reasons why kernel estimation fails. Therefore we need two different measures to decide which kernel is wrong. To identify textureless regions we compute entropy of the kernels and take those with entropy above some threshold. The other, more serious case of failure is pixel saturation, that is pixel values above the sensor range. This situation can be identified by computing the sum of kernel values, which should be close to one for valid kernels. Therefore, we simply remove kernels whose sum is too different from unity, again above some threshold.

For the restoration step, we use an energy minimization approach with total variation as an image regularization term, which belongs to the category of constrained least squares estimators [14]. Notice that it has the same form as (6). Total variation behaves satisfactorily for most photographs since it removes noise efficiently while not oversmoothing edges. It also helps to some extent to suppress artifacts caused by pixel saturation.

The restoration phase of the proposed algorithm can be described as minimization of the functional

$$E(\mathbf{u}) = \frac{1}{2}\|\mathbf{u} *_v h - \mathbf{z}\|^2 + \lambda \int |\nabla\mathbf{u}| \tag{11}$$

with respect to the unknown sharp image $\mathbf{u}$, where the second term is the total variation of the image.

Its derivative can be written as

$$\partial E(\mathbf{u}) = (\mathbf{u} *_v h - \mathbf{z}) \circledast_v h - \lambda \operatorname{div}\left(\frac{\nabla \mathbf{u}}{|\nabla \mathbf{u}|}\right), \tag{12}$$

where $\circledast_v$ is the operator adjoint to space-variant convolution

$$\mathbf{u} \circledast_v \mathbf{h}\,[x,y] = \int \mathbf{u}(x-s, y-t)\mathbf{h}(x,y;-s,-t)\,dsdt. \tag{13}$$

To minimize functional (11) we used a half-quadratic iterative approach, reducing this problem to a sequence of linear subproblems [14].

Alternatively, to speed up the restoration step, we could use a variant of the Richardson-Lucy algorithm, similar to methods [15,16]. Figures 4–6 show an example of a real image restored by this method.

In our opinion, this is the best of the three deblurring approaches. It is quite fast and reliable. Because of its stability it can be used to estimate the space-variant PSF, which makes it more applicable for a much larger range of situations. Another plus is that it can be used to segment moving objects, which is hardly possible from one image.

Fig. 4 Image of a shopping center taken in an evening with shutter speed $1/2$ s (*left*), results of our algorithm with PSF adjustment (*right*). Close-ups are shown in Figure 6.
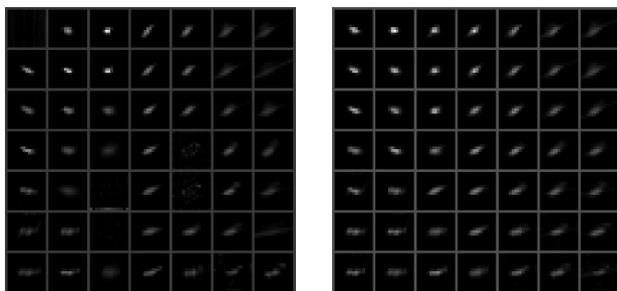

Fig. 5 Fourty-nine convolution kernels estimated in the shopping center image (*left*). Notice the wrong kernels at places of low-contrast texture (*upper left corner*) and pixel saturation (lights inside the building). Adjusted kernels on the right.
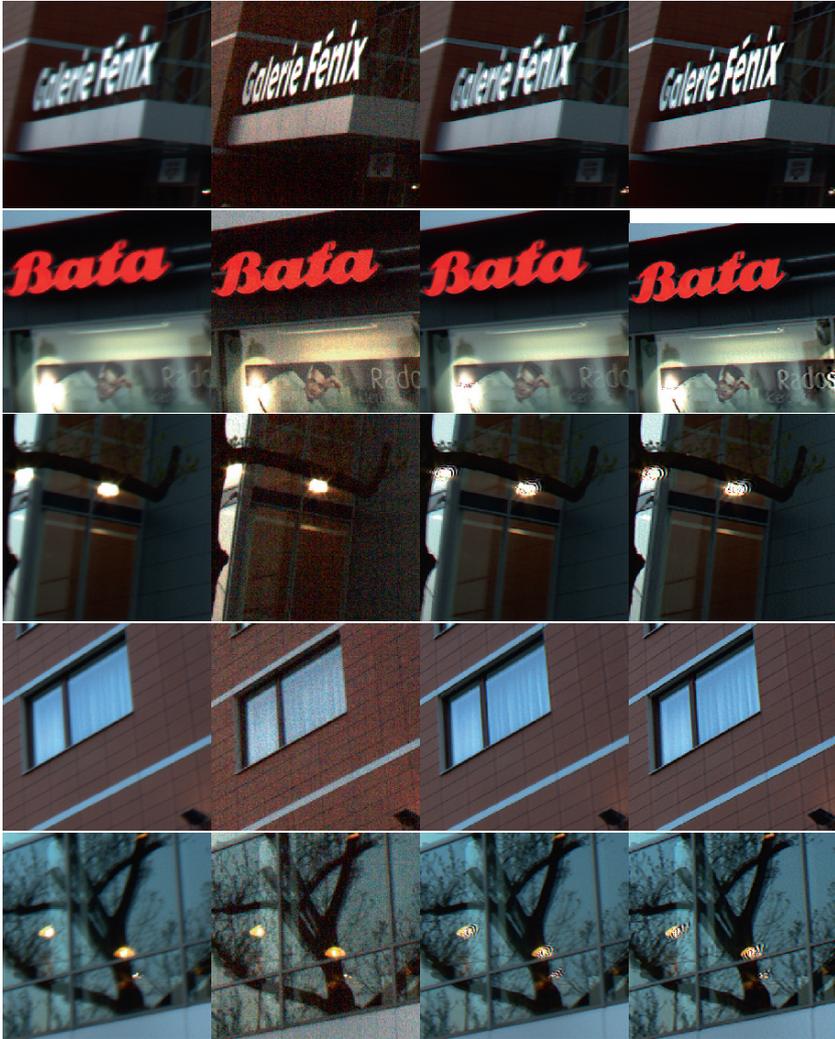
Fig. 6 Details of the shopping center image. From left to right – the blurred image, noisy image, result of deconvolution and our result.

8 Summary

In this chapter, we reviewed approaches to software image stabilization in the sense of removing blur caused by camera motion (Table 1).

The first possibility is to avoid blur from the beginning by taking a sequence of underexposed images. This idea is impractical because of the time needed for sensor read-out. We followed with the description of a deblurring algorithm from a single image. Although there are usable algorithms for this case, the main disadvan-

Table 1 Summary of approaches to image stabilization.

Approach	Speed	Quality	Main problem
Multiple underexposed images	High	High	Slow read-out, precise registration
Single-image deconvolution	Slow/medium	Medium	Homogenous blur only
Multi-image deconvolution	Slow	Medium/high	Slow computation
One blurred and one noisy image	Medium	Medium	More artifacts than multi-image deconvolution

tages are speed and difficulties with the segmentation of moving objects. The third approach was deconvolution from a sequence of blurred images. The main disadvantage of existing algorithms from this category is speed. They are even slower than single image deconvolution methods.

The last and, in our opinion, most advantageous approach is to use a pair of images, one blurred and one underexposed. Its main assets are relative speed, reliability, ability to deal with space-variant blur and the potential to segment moving objects.

References

1. M. R. Banham and A. K. Katsaggelos. Digital image restoration. *IEEE Signal Process. Mag.*, 14(2):24–41, March 1997.
2. R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graphics, SIGGRAPH 2006 Conf. Proc., Boston, MA*, 25:787–794, 2006.
3. M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.
4. R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University, Cambridge, 2nd edition, 2003.
5. S. H. Lim and D. A. Silverstein. Method for deblurring an image. US Patent Application, Pub. No. US2006/0187308 A1, August 24 2006.
6. J. Miskin and D. MacKay. Ensemble learning for blind image separation and deconvolution. In M. Girolami (Ed.), *Advances in independent component analysis*. Berlin: Springer-Verlag, (pp. 123–141), 2000.
7. Q. R. Razligh and N. Kehtarnavaz. Image blur reduction for cell-phone cameras via adaptive tonal correction. pages I: 113–116, 2007.
8. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

9. Qi Shan, J. Jia, and A. Agarwala. High-quality motion deblurring from a single image. *ACM Trans. Graphics (SIGGRAPH)*, 27(3) 2008.

10. F. Šroubek and J. Flusser. Multichannel blind deconvolution of spatially misaligned images. *IEEE Trans. Image Process.*, 14(7):874–883, July 2005.

11. M. Tico, M. Trimeche, and M. Vehvilainen. Motion blur identification based on differently exposed images. In *Proc. IEEE Int. Conf. Image Process.*, pp. 2021–2024, 2006.

12. A. Tikhonov and V. Arsenin. *Solution of Ill-Posed Problems*. Wiley, New York, 1977.

13. D. Tschumperlé and R. Deriche. Vector-valued image regularization with pdes: A common framework for different applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):506–517, 2005.

14. M. Šorel and J. Flusser. Space-variant restoration of images degraded by camera motion blur. *IEEE Trans. Image Process.*, 17(2):105–116, February 2008.

15. L. Yuan, J. Sun, L. Quan, and H.-Y. Shum. Image deblurring with blurred/noisy image pairs. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, p. 1, ACM, New York, 2007.

16. L. Yuan, J. Sun, L. Quan, and H.-Y. Shum. Progressive inter-scale and intra-scale non-blind image deconvolution. *ACM Trans. Graphics (SIGGRAPH)*, 2008.