Multiple-image deblurring with spatially-variant point spread functions[★]

R. Vio^{1,2}, J. Nagy³, and W. Wamsteker⁴

¹ Chip Computers Consulting s.r.l., Viale Don L. Sturzo 82, S.Liberale di Marcon, 30020 Venice, Italy
e-mail: robertovio@tin.it

² ESA-VILSPA, Apartado 50727, 28080 Madrid, Spain

³ Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA
e-mail: nagy@mathcs.emory.edu

⁴ ESA-VILSPA, Apartado 50727, 28080 Madrid, Spain
e-mail: willem.wamsteker@esa.int

Received 27 November 2003 / Accepted 3 January 2005

Abstract. We generalize a reliable and efficient algorithm, to deal with the case of spatially-variant PSFs. The algorithm was developed in the context of a least-square (LS) approach, to estimate the image corresponding to a given object when a set of observed images are available with different and spatially-invariant PSFs. Noise is assumed additive and Gaussian. The proposed algorithm allows the use of the classical LS single-image deblurring techniques for the simultaneous deblurring of the observed images, with obvious advantages both for computational cost and ease of implementation. Its performance and limitations, also in the case of Poissonian noise, are analyzed through numerical simulations. In an appendix we also present a novel, computationally efficient deblurring algorithm that is based on a Singular Value Decomposition (SVD) approximation of the variant PSF, and which is usable with any standard space-invariant direct deblurring algorithm. The corresponding Matlab code is made available.

Key words. methods: data analysis – methods: statistical – techniques: image processing

1. Introduction

Vio et al. (2004) presented a computationally efficient solution for the problem of simultaneous deblurring of a set of observed images (of a fixed object), where each observed image is degraded by a different spatially-invariant point spread function (PSF) and the noise is additive and Gaussian. The most serious limitation of the other techniques available in the literature is the use of algorithms that work cyclically on one image at a time (see Bertero & Boccacci 2000). For this reason they are computationally expensive and demanding of computer memory. The main idea of Vio et al. (2004) is to combine the observed images in a single one containing all the relevant information. In this way, apart from the obvious benefit concerning the memory requirements, it is possible to improve the convergence rate of Least-Squares (LS) iterative algorithms. Moreover, there is the additional benefit that standard algorithms (and software) for image restoration can be used.

There are important fields where potentially this methodology can be exploited. An example is the multi-frequency observations of the Cosmic Microwave Background (CMB)

obtained from satellites such as *PLANCK*. Although some other physical components are present at the microwave frequencies, the images taken at high galactic latitudes are expected to be almost entirely dominated by the CMB contribution over a large range of observing frequencies (e.g., see Stolyarov et al. 2002). Since the PSFs corresponding to the images obtained at the various frequencies can be quite dissimilar, simultaneous deblurring can be useful to improve the extraction of CMB information from the data. Unfortunately, since the PSF of *PLANCK*'s instruments is expected to be non-circular and, because of the observational strategy, to rotate during the observations (Arnau et al. 2002; Burigana & Saez 2003), the computational advantages provided by the methodology of Vio et al. (2004) cannot be exploited.

A similar situation can also hold for the restoration of images that will be obtained by the Large Binocular Telescope (LBT). This instrument consists of two 8.4 m mirrors on a common mount with a spacing of 14.4 m between the centers (Angel et al. 1998). For a given orientation of the telescope, the diffraction-limited resolution along the center-to-center baseline is equivalent to that of a 22.8 m mirror, while the resolution along the perpendicular direction is that of a single 8 m mirror. A possible way to obtain an image with an

[★] Figures 1 to 12 and Appendix A are only available in electronic form at <http://www.edpsciences.org>

improved and uniform spatial resolution is to simultaneously deconvolve the images taken with different orientations of the telescope. However, because of the different corrections of the adaptive optics across the image domain, the PSF for LBT can be space-variant (Bertero & Boccacci 2000; Carbillet et al. 2002).

To our knowledge, multiple-image deblurring in the case of a spatially-variant PSF has not been considered in the current literature. Given that the single-image technique undoubtedly offers some important benefits, in the present paper we extend it to deal with the case of PSFs that change across the images. In particular, in Sect. 2 we formalize the problem and propose an efficient solution in Sect. 3. In Sect. 4 its performance is studied through numerical experiments. In Sect. 5 we close with final comments and conclusions.

2. Formalization of the problem

Suppose one has p observed images, $g_j(x, y)$, $j = 1, 2, \dots, p$, $x, y \in \mathbb{R}$, of a fixed two-dimensional object $f_0(x, y)$, each of which is degraded by a different blurring operator. In the case of a generic blurring operator, $g_j(x, y)$ can be modeled as

$$g_j(x, y) = \iint K_j(x, y, x', y') f_0(x', y') dx' dy', \quad (1)$$

where, for each j , $K_j(x, y)$ is a (possibly) spacially-variant PSF.

In practice, the experimental images are discrete $N \times N$ arrays of pixels (for simplicity, we assume square images) contaminated by noise. Therefore model (1) has to be recast in the form

$$g_j(m, n) = \sum_{m', n'=0}^{N-1} K_j(m, n, m', n') f_0(m', n') + w_j(m, n), \quad (2)$$

where $0 \leq m, n < N$, and $w_j(m, n)$ is additive white noise. For the moment, we assume constant standard deviations across the images: $\sigma_{w_j} = \sigma_w$. The image restoration problem of interest is to find an estimate $f(m, n)$ of $f_0(m, n)$ given the observed images $g_j(m, n)$ and the known PSFs $K_j(m, n, m', n')$.

In the case of a spatially-invariant PSF, with a support much smaller than the image, model (2) can be well approximated by a cyclic convolution

$$g_j(m, n) = \sum_{m', n'=0}^{N-1} K_j(m - m', n - n') f_0(m', n') + w_j(m, n), \quad (3)$$

or equivalently

$$\widehat{g}_j(m, n) = \widehat{K}_j(m, n) \widehat{f}_0(m, n) + \widehat{w}_j(m, n), \quad (4)$$

where the symbol “ $\widehat{\cdot}$ ” denotes the two-dimensional Discrete Fourier Transform (DFT). One approach to compute an estimate of $f_0(m, n)$ is to pose the image restoration problem as a least-squares (LS) problem, which is equivalent to a maximum likelihood approach in the case of white Gaussian noise. In order to outline the LS approach, it is helpful to use matrix-vector notation: Eq. (3) can be rewritten as

$$\mathbf{g}_j = \mathbf{A}_j \mathbf{f}_0 + \mathbf{w}_j, \quad (5)$$

where the $N^2 \times 1$ arrays $\mathbf{g}$, $\mathbf{f}_0$ and $\mathbf{w}_j$ are obtained by column stacking $g_j(m, n)$, $f_0(m, n)$, and $w_j(m, n)$, respectively. $\mathbf{A}_j$ is the $N^2 \times N^2$ block-circulant matrix defined by cyclic convolution with the j th PSF. The LS estimate $\mathbf{f}$ can then be expressed in the form:

$$\mathbf{f}_{\text{LS}} = \arg \min_{\mathbf{f}} \sum_{j=1}^p \|\mathbf{A}_j \mathbf{f} - \mathbf{g}_j\|^2. \quad (6)$$

It is not difficult to see that $\mathbf{f}_{\text{LS}}$ is the solution of the normal equations

$$\sum_{j=1}^p \mathbf{A}_j^T \mathbf{A}_j \mathbf{f} = \sum_{j=1}^p \mathbf{A}_j^T \mathbf{g}_j, \quad (7)$$

where $\mathbf{A}_j^T$ denotes the transpose of the block-circulant matrix $\mathbf{A}_j$. The DFT of both sides of Eq. (7) provides

$$\widehat{f}(m, n) \sum_{j=1}^p |\widehat{K}_j(m, n)|^2 = \sum_{j=1}^p \widehat{K}_j^*(m, n) \widehat{g}_j(m, n), \quad (8)$$

with the symbol “ $*$ ” denoting complex conjugation.

The classical unbiased LS estimate given by the inverse DFT of

$$\widehat{f}(m, n) = \frac{\sum_j \widehat{K}_j^*(m, n) \widehat{g}_j(m, n)}{\sum_j |\widehat{K}_j(m, n)|^2} \quad (9)$$

can be very unstable. In fact, for white noise, the mean square error (MSE) of Eq. (9) is

$$\mathbb{E}[\|\mathbf{f} - \mathbf{f}_0\|^2] = \sigma_w^2 \sum_{m, n} \frac{1}{\sum_j |\widehat{K}_j(m, n)|^2}. \quad (10)$$

The quantities $\widehat{K}_j(m, n)$ can be very close to zero at high frequencies because the PSFs are almost band-limited in practice. Consequently, the MSE can be very large, which indicates that the LS estimate is unstable. This implies the necessity of stabilizing the solution by limiting the values of such terms. A classic approach is to exploit the semiconvergence property of the iterative inversion algorithms. The use of such algorithms is indicated whenever it is necessary to incorporate constraints on the solution (e.g., the non-negativity). The main limitation of the use of the normal equations with the iterative algorithms is that the resulting deblurring problem is much more ill-conditioned than implied by each of the models (3) because the coefficients $\widehat{K}(m, n)$ in Eq. (8) are squared. The most important consequence is that the convergence of the iterative algorithms can be remarkably slowed with obvious consequence for the computational cost.

In order to avoid this problem, Vio et al. (2004) proposed to use the normal Eq. (8) in the form

$$\widehat{K}_{j_0}(m, n) \widehat{f}(m, n) = \frac{\widehat{g}_{j_0}(m, n) + \sum_{j \neq j_0} [\widehat{K}_j^*(m, n) / \widehat{K}_{j_0}^*(m, n)] \widehat{g}_j(m, n)}{1 + \sum_{j \neq j_0} [|\widehat{K}_j(m, n)|^2 / |\widehat{K}_{j_0}(m, n)|^2]} = \widehat{\zeta}(m, n), \quad (11)$$

where $\widehat{K}_{j_0}(m, n)$ is given by

$$\widehat{K}_{j_0}(m, n) = \max[\widehat{K}_1(m, n), \widehat{K}_2(m, n), \dots, \widehat{K}_p(m, n)]. \quad (12)$$

Here, for each coordinate (m, n) , the operator $\max[\]$ extracts the element in the array with the largest modulus, and j_0 is the value of the corresponding index j . This equation shows that the multi-image LS deblurring problem (6) is equivalent to classical one-image LS deblurring where the “observed” image is $\widehat{\zeta}(m, n)$ (called the *mean image*) and the PSF is $\widehat{K}_{j_0}(m, n)$ (called the *mean PSF*). Model (11) permits one to improve the performance of the iterative algorithms. In fact, $\widehat{\zeta}(m, n)$ represents the DFT of the “observed” image with the PSF whose energy distribution is maximally concentrated in the spatial domain. In other words, the conditioning of system (11) is much better than that of system (8) (this is because $\widehat{K}_{j_0}(m, n)$ is not obtained by squaring the coefficients $\widehat{K}_j(m, n)$). It is well known that this means a faster convergence rate for the iterative algorithms (Vio et al. 2004). Moreover, the composition of the images $g_j(m, n)$ in the single-image $\zeta(m, n)$ permits the use of standard software for the deblurring operation.

From Eq. (11) it is evident that the single-image $\zeta(m, n)$ is obtained in the Fourier domain. From the computational point of view, this is quite advantageous since it permits us to exploit fast implementations of the DFT. However, in the case of spatially-variant PSFs some problems arise because model (2) has no representation in the Fourier domain. Consequently, model (11) has to be modified in such a way to take into account the changes of the PSF across the images. In the next section, a possible solution is provided.

3. The construction of the single-image in the case of a spatially-variant PSF

If s is an image,

$$s = \sum_{i=1}^{N_D} D_i \cdot s, \quad (13)$$

where “ $\cdot$ ” denotes element-wise multiplication, and D_i are $N \times N$ matrices with entries in the range $[0, 1]$ such that $\sum_i D_i(m, n) = 1$ for each value of (m, n) . Now, since the DFT is a linear operator, if

$$\widehat{s} = \text{DFT}[s], \quad (14)$$

we have that

$$\widehat{s} = \text{DFT}\left[\sum_{i=1}^{N_D} D_i \cdot s\right] = \sum_{i=1}^{N_D} \text{DFT}[D_i \cdot s] = \sum_{i=1}^{N_D} \widehat{s}_i, \quad (15)$$

where $\widehat{s}_i = \text{DFT}[D_i \cdot s]$. By using this equation with $s = g_j$, Eq. (11) can be written as

$$\widehat{\zeta}(m, n) = \sum_{i=1}^{N_D} \frac{\widehat{g}_{ij_0}(m, n) + \sum_{j \neq j_0} [\widehat{K}_j^*(m, n) / \widehat{K}_{j_0}^*(m, n)] \widehat{g}_{ij}(m, n)}{1 + \sum_{j \neq j_0} [|\widehat{K}_j(m, n)|^2 / |\widehat{K}_{j_0}(m, n)|^2]}, \quad (16)$$

with $\widehat{g}_{ij} = \text{DFT}[D_i \cdot g_j]$. Of course, this version of Eq. (11) is of no practical interest. However, if the matrices $D_i(m, n)$ have non-zero entries only in correspondence to square or rectangular regions of the image domain, and if in correspondence to the i th region the j th image $g_j(m, n)$ has its own spatially-invariant

PSF $K_{ij}(m, n)$ (i.e., the blur is spatially-invariant in the subregions corresponding to the images g_{ij}), then Eq. (16) can be written in the form

$$\begin{aligned} \widehat{\zeta}(m, n) &= \sum_{i=1}^{N_D} \frac{\widehat{g}_{ij_0}(m, n) + \sum_{j \neq j_0} [\widehat{K}_{ij}^*(m, n) / \widehat{K}_{ij_0}^*(m, n)] \widehat{g}_{ij}(m, n)}{1 + \sum_{j \neq j_0} [|\widehat{K}_{ij}(m, n)|^2 / |\widehat{K}_{ij_0}(m, n)|^2]} \\ &= \sum_{i=1}^{N_D} \widehat{\zeta}_i(m, n). \end{aligned} \quad (17)$$

In practical applications the PSFs change smoothly across the image. In this case, the image given by Eq. (17) provides an approximation whose accuracy is regulated by N_D . At this point, the problem becomes a definition of the matrices $D_i(m, n)$ that permits us to get satisfactory results.

The simplest form for $D_i(m, n)$ is obtained through *piecewise constant interpolation*, when the permitted values are either 0 or 1. In this case, the non-zero entries constitute a partitioning of the image domain. Although in many situations such a choice provides satisfactory results, sometimes the resulting $\zeta(m, n)$ suffers from blocking artifacts that are problematic for the deblurring operation. A more sophisticated alternative, able to avoid this kind of problem, is based on a linear interpolation approach (see Nagy & O’Leary 1998). The basic idea is simpler to understand in the one-dimensional case.

If a one-dimensional domain is partitioned into N_D subregions with dimension N_s , a useful set of overlapping masks, with dimension $2 \times N_s$, is shown in Fig. 1. In fact, since they go smoothly to zero, the boundary artifacts that affect each of the $\widehat{s}_i(m, n)$ are strongly weakened. Moreover, the overlap works in such a way as to further reduce these effects. The only exception is represented by the regions on the border of the domain.

In the two-dimensional case the argument is similar although made slightly more difficult by the larger number of masks that overlap at a given point of the image domain, and by the fact that there are two different situations for the subregions on the border of the frame. However, it is not difficult to see that, independent of the partitioning of the image domain, the resulting masks must have fixed values at the points marked in Fig. 1. Computationally, they can be obtained by interpolating these points through a *bilinear* interpolation (Press et al. 1992).

3.1. Two additional problems

So far we have assumed that the noise level is the same in all of the images ($\sigma_{w_j} = \sigma_w$), a condition that allows the combination of different images to provide the most interesting results. Moreover, this requirement is often satisfied in practice. The extension of the method to images with different noise levels is straightforward (Vio et al. 2004). If the different σ_{w_i} are known, one may account for the different variabilities by substituting Eq. (17) with

$$\begin{aligned} \widehat{\zeta}(m, n) &= \sum_{i=1}^{N_D} \frac{\widehat{Q}_{ij_0}(m, n) + \sum_{j \neq j_0} [\widehat{\mathcal{K}}_{ij}^*(m, n) / \widehat{\mathcal{K}}_{ij_0}^*(m, n)] \widehat{Q}_{ij}(m, n)}{1 + \sum_{j \neq j_0} [|\widehat{\mathcal{K}}_{ij}(m, n)|^2 / |\widehat{\mathcal{K}}_{ij_0}(m, n)|^2]} \\ &= \sum_{i=1}^{N_D} \widehat{\zeta}_i(m, n). \end{aligned} \quad (18)$$

Here, $\widehat{\mathcal{K}}_{ij}(m, n) = \widehat{K}_{ij}(m, n)/\sigma_{w_j}$ and $\widehat{\varrho}_{ij}(m, n) = \widehat{g}_{ij}(m, n)/\sigma_{w_j}$. The rest of the procedure remains unaltered.

A last question regards the images that present features close to their boundaries. This causes the appearance of edge effects. In the present context, the situation is made more troublesome by the fact that the masks $D_i(m, n)$ can overlap, thus causing a propagation of artifacts to the central regions of $\widehat{\zeta}_i(m, n)$. For this reason, as explained in Vio et al. (2004) (see also Von der Lühe 1993), before the calculation of $\widehat{\zeta}_i(m, n)$, all the images $g_j(m, n)$ should be windowed through the matrix

$$h(m, n) = \begin{cases} 0.25 \times \alpha \times \beta, & 0 \leq m, n < N_w; \\ 0.25 \times \alpha \times \beta, & N - N_w \leq m, n < N; \\ 1 & \text{otherwise.} \end{cases} \quad (19)$$

The parameters $\alpha = [1 - \cos(m\pi/N_w)]$ and $\beta = [1 - \cos(n\pi/N_w)]$ are chosen so that the pixels in the central subimage are not modified and the image has continuous first derivatives at the edges. This window approaches the classical two-dimensional Hanning window as $N_w \rightarrow N/2$ and tends to the rectangular window as $N_w \rightarrow 0$. The parameter N_w thus determines the filtering characteristics, in particular the frequency pass-band, of the window. Its “optimal” value depends on many factors such as the noise level, the form of the PSF and the specific realization of the process. For a Gaussian PSF, simulation results (Vio et al. 2003b) show that values of three or four times the dispersion of the PSF provide reasonably good results. The only unavoidable consequence of the procedure is that a border of thickness N_w has to be removed from the final deblurred image.

4. Some numerical experiments

We present results of some numerical experiments to show the reliability and the good performance of the single-image approach for the multiple-image deblurring in case of spatially-variant PSF. As claimed in the previous sections, in principle the deblurring of the image $\zeta(m, n)$ can be carried out by using any available algorithm able to deal with PSFs that are not constant across the images. Here we use two particular techniques. The first is based on an iterative method, which is a modified residual norm steepest descent (MRNSD) method that minimizes a least squares fit to data, using an exponential function reparameterization to enforce a non-negativity constraint; see Nagy & Strakoš (2000) for details. The MRNSD method is equivalent to an approach studied by Kaufman (1993), but the numerical results reported here are obtained using the implementation described by Nagy & Strakoš (2000), which views the approach as a preconditioned steepest descent method. In many respects (such as computational cost per iteration) MRNSD is similar to the Richardson-Lucy algorithm (Vogel 2002). The difference between the two methods is that MRNSD is based on a least squares criteria, where as Richardson-Lucy is based on a maximum likelihood criteria. Furthermore, MRNSD typically converges much more quickly than the Richardson-Lucy method. In either case, the main cost of the algorithm is matrix vector multiplications at each iteration. Efficient matrix vector multiplications of the space variant PSF, using the sectioning ideas given above, is described

in Nagy & O’Leary (1998). The second approach uses standard Tikhonov regularization (Vogel 2002), efficiently implemented as a filtering-based method (Vio et al. 2003a), using a novel Singular Value Decomposition (SVD) approximation of the variant PSF (see Appendix A). In this last case, the use of the single-image $\zeta(m, n)$ has no effect on the efficiency of the algorithm, but it makes coding the scheme much easier. As is well known, direct methods are often much more efficient than iterative methods. However, given the difficulty concerning the incorporation of constraints on the solution (e.g., nonnegativity), they have to be expected to provide results of inferior quality.

For each experiment we use a set of eight images of 400×400 pixels. Figure 2 shows the PSF corresponding to a set of pixels uniformly distributed across the first image. All the displayed PSFs are bidimensional Gaussian with dispersion along the major axis set to twelve pixels, and to four pixels along the minor axis. For the remaining pixels, the PSF is given by a bilinear interpolation of the displayed pattern. The PSF of the other figures is obtained by rotating this PSF by $22.5^\circ, 45^\circ, 67.5^\circ, 90^\circ, 112.5^\circ, 135^\circ, 157.5^\circ$ degrees, respectively. The non-stationary pattern of the PSF has been deliberately chosen quite peculiar in order to test the capabilities of the algorithm in extreme situations. The images g_j have been obtained by the direct computation of the discrete model (2).

Figures 3 and 4 show the deblurring results for a set of 169 identical point-like objects uniformly distributed across the image. Noise is additive and Gaussian with a standard deviation of 2% of the maximum value of the blurred images. We have chosen these particular examples since, the sources being of non-smooth functions, deblurring is a difficult problem. The image $\widehat{\zeta}_i(m, n)$ has been calculated by partitioning the *observed* images in 25 square subregions (i.e., $N_D = 25$), and by using piecewise constant interpolation for $D_i(m, n)$. For comparison, the deblurring of the first image of the set (that with the PSF shown in Fig. 2) is also presented; an MRNSD algorithm for spatially-invariant PSF has been used. In each case the improvement due to the use of the eight images is evident. The superiority of the results provided by the MRNSD approach is due to the fact that the Tikhonov method does not enforce a non-negativity constraint on the solution.

Although the use of piecewise constant interpolation for $D_i(m, n)$ can provide reasonable results, very often the linearly interpolated version is able to remarkably improve them. Figures 5–7 show the results of an experiment, carried out under the same conditions as the previous one, that confirms the superiority of the interpolated method (here $N_s = 80$).

In various situations the choice of the linear interpolation mask is a necessity. When the constant interpolation mask is applied to the object shown in Fig. 8 (NB. here noise is additive and Gaussian with a standard deviation of 2%, the standard deviation of the dispersion of the blurred images), the results are not so good since $\widehat{\zeta}_i(m, n)$ contains blocking artifacts that later are magnified by the deblurring operation. Consequently, the use of the linearly interpolated mask becomes necessary. Although this approach produces better results, some problems still remain in the form of fringes that contaminate $\zeta_i(m, n)$. This is linked to edge effects because the original image

contains significant features close to its boundaries. For this reason, before the calculation of $\widehat{\zeta}_i(m, n)$, all the images $g_j(m, n)$ should be windowed as explained at the end of Sect. 3.1. Figure 9 shows the results obtained after this operation with $N_s = 80$ and $N_w = 40$. A border of 40 pixels was removed by the windowing operation (see Vio et al. 2004). Again, the improvement due to the use of the eight images is evident. In this specific case, since the non-negativity of the solution is not as critical as in the experiments presented above, Tikhonov regularization is able to provide results as good as those provided by MRNSD, but at a much lower computational cost (and thus much more quickly).

Figures 10 and 11 show the results that can be obtained in presence of Poissonian noise. Here the simulated images are the same as those shown in Figs. 6 and 9. The only difference is that the Gaussian noise is not added to the blurred images, but to non-stationary Poissonian processes with the local mean given by the values of the pixels in the blurred images themselves (peak signal to noise ratio = 40 dB). In this way it is possible to get an idea of the degradation in the quality of deblurring operation due to the presence of non-Gaussian noise. The visual inspection of these figures, as well as the rms of the true residuals (i.e., $\|f - f_0\|/\sqrt{N}$), clearly indicate that, if present, such degradation is very small. This conclusion is not unexpected; various works have already shown that LS methods are able to provide results very similar (and often almost identical) to those obtainable with Maximum Likelihood techniques that take into account the Poissonian nature of the noise (e.g., see Bertero & Boccacci 2000). However, the possible worsening of the quality of the results is not due to the use of the mean image $\zeta(m, n)$, but is intrinsic to the adoption of a LS model with non-Gaussian noise.

Figure 12 shows the results obtainable when the proposed deblurring method is applied to simulated observations of the *PLANCK* Low Frequency Instrument (LFI). The images have been computed as explained in Vio et al. (2003a). The region we consider is a square patch of 370×370 pixels with a side of about 20° centered at $l = 90^\circ$, $b = 45^\circ$ (Galactic coordinates). The latitude is high enough that the Cosmic Microwave Background (CMB) dominates over the foreground. The *PLANCK*-LFI instrument works at frequencies 30, 44, 70, and 100 GHz, but we consider only the last three since at these frequencies the contamination by the foreground is negligible. We assume nominal noise and angular resolution¹. The maps are blurred through Gaussian PSFs with elliptical symmetry (axial ratio 1:1.3) and appropriate FWHMs (i.e., for the major axis, $\approx 23'$ at 44 GHz, $\approx 14'$ at 70 GHz, $\approx 10'$ at 100 GHz). Although the PSFs are not expected to change their shape during the observations, it is not improbable that they may suffer a rotation. Since the importance of such a phenomenon is not yet well determined, for all the frequencies we assume a spatially-variant PSF similar to that shown in Fig. 2, which certainly represents a worst-possible-case scenario. Finally, each map is summed together with simulated white noise with a rms level as expected for the considered

channels. Since we choose to work with a pixel size of about 3.5 arcmin, the noise rms are .049, .042 and .043 mK in antenna temperature at 44, 70, 100 GHz, respectively. The assumption of white noise is common in this kind of experiment (e.g., see Hobson et al. 1998; Stolyarov et al. 2002; Vio et al. 2003a). The CMB maps can have negative values, and this prevents the use of methods exploiting the non-negativity constraint of the solution. For this reason, only the Tikhonov deblurring method has been used. Because of the edge effects, the image $\zeta_i(m, n)$ has been obtained through the windowing procedure indicated above with $N_s = 60$ and $N_w = 30$. Using the three maps instead of only the one with the highest resolution provided final results whose rms of the true residuals is improved by about 1%.

5. Conclusions

We have considered the problem of simultaneously deblurring a set of images of a fixed object when the corresponding PSFs are space-variant and different from each other. We have developed a method, based on a LS approach, to efficiently transform a multi-image deblurring problem into a single-image one. This approach provides substantial savings in computational requirements, and can be implemented using numerical LS algorithms for single-image problems. These conclusions are confirmed by our numerical experiments. We have also presented a novel, computationally efficient, direct deblurring algorithm, which is based on a Singular Value Decomposition (SVD) approximation of the space variant PSF coupled with Tikhonov regularization.

Specific implementations used for the experiments are available in the Matlab package *RestoreTools* (Nagy et al. 2004b)².

References

- Angel, J. R. P., Hill, J. M., Strittmatter, P. A., Salinari, P., & Weigelt, G. 1998, in *Astronomical Interferometry*, ed. R. D. Reasenberg, Proc. SPIE, 3352, 881
- Arnau, J. V., Aliaga, A. M., & Saez, D. 2002, *A&A*, 382, 1138
- Bertero, M., & Boccacci, P. 2000, *A&A*, 147, 323
- Boden, A. F., Redding, D. C., Hanisch, R. J., & Mo, J. 1995, *J. Opt. Soc. Am. A*, 13, 1537
- Burigana, C., & Saez, D. 2003, *A&A*, 409, 423
- Carbillet, M., Correia, S., Boccacci, P., & Bertero, M. 2002, *A&A*, 387, 744
- Faisal, M., Lanterman, A. D., Snyder, D. L., & White, R. L. 1995, *J. Opt. Soc. Am. A*, 12, 2593
- Hobson, M. P., Jones, A. W., Lasenby, A. N., & Bouchet, F. R. 1998, *MNRAS*, 300, 1
- Kamm, J., & Nagy, J. G. 1998, in *Advanced Signal Processing Algorithms, Architectures, and Implementations VIII*, ed. F. T. Luk, SPIE, 3461, 358
- Kamm, J., & Nagy, J. G. 2000, *SIAM J. Matrix Anal. Appl.*, 22, 155
- Kaufman, L. 1993, *IEEE Trans. Med. Imag.*, 12, 200
- Nagy, J. G., & O'Leary, in *Advanced Signal Processing Algorithms, Architectures, and Implementations VIII*, ed. F. T. Luk, SPIE, 3462, 388

¹ <http://astro.estec.esa.nl/SA-general/Projects/Planck>

² Available at <http://www.mathcs.emory.edu/~nagy/RestoreTools/index.html>

- Nagy, J. G., & O'Leary, D. P. 1998, *SIAM J. Sci. Comput.*, 19, 1063
- Nagy, J. G., & Strakoš, Z. 2000, in *Mathematical Modeling, Estimation, and Imaging*, ed. D. C. Wilson, et. al., SPIE, 3461, 182, <http://www.mathcs.emory.edu/~nagy/research/pubs.html>
- Nagy, J. G., Ng, M. K., & Perrone, L. 2004a, *SIAM J Matrix Anal. and Appl.*, in press, <http://www.mathcs.emory.edu/~nagy/research/pubs.html>
- Nagy, J. G., Lee, K. P., & Perrone, L. 2004b, *Numerical Algorithms*, in press, <http://www.mathcs.emory.edu/~nagy/research/pubs.html>
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes* (Cambridge: Cambridge University Press)
- Stolyarov, V., Hobson, M. P., Ashdown, M. A. J., & Lasenby, A. N. 2002, *MNRAS*, 336, 97
- Von der Lühe, O. 1993, *A&A*, 268, 374
- Vio, R., Nagy, J., Tenorio, L., et al. 2003a, *A&A*, 389, 404
- Vio, R., Nagy, J., Tenorio, L., et al. 2003b, *A&A*, 408, 835
- Vio, R., Nagy, J., Tenorio, L., & Wamsteker, W. 2004, *A&A*, 416, 403
- Vogel, C. R. 2002, *Computational Methods for Inverse Problems*, Ser. in Front. in Appl. Math., Vol. 24 (Philadelphia: SIAM)

Online Material

Fig. 1. One- and two-dimensional bases for the computation of the masks using the linear interpolation approach. In the one-dimensional case the dashed line indicates the limits of each subregion, whereas in the two-dimensional case they indicate the limits of the masks. Symbols $\times$ denote the center of a subregion, whereas the values of the masks at the marked points are given by $\bullet = 1$, $\blacklozenge = 0.5$, $\blacktriangle = 0.25$, $\blacksquare = 0$.

Fig. 2. Spatial variation of the PSF across the first image of the set used in the numerical experiments (see text). The pair of numbers in the header of each panel indicates the coordinates of the pixels to which the displayed PSF corresponds. In the intermediate pixels the PSFs are obtained through a linear interpolation of the displayed ones. Each PSF is given by a two-dimensional Gaussian function with a dispersion of 12 pixels along the major axis and 4 pixels along the minor axis.

Fig. 3. Original image, blurred version and blurred + noise version of the first image of the set (see text), and corresponding mean PSF. The images are 400×400 pixels.

Fig. 4. Mean image $\zeta(m,n)$ (with $D_i(m,n)$ obtained using piecewise constant interpolation), MRNSD and Tikhonov deblurred versions of $\zeta(m,n)$, and MRNSD deblurred version of the first image of the set (see text) for the image shown in Fig. 3. For the MRNSD method the images shown are those obtained after 300 iterations of the algorithm, whereas for Tikhonov it is the one with the smallest rms of the true residuals; for the mean image and first image of the set the rms of the true residuals are, respectively, 0.028 (MRNSD) – 0.032 (Tikhonov), and 0.032.

Fig. 5. Original image, blurred version and blurred + noise version of the first image of the set (see text), and corresponding mean PSF. The images are 400×400 pixels.

Fig. 6. Mean image $\zeta(m, n)$ (with $D_i(m, n)$ obtained using piecewise constant interpolation), MRNSD and Tikhonov deblurred versions of $\zeta(m, n)$, and MRNSD deblurred version of the first image of the set (see text) for the image shown in Fig. 5. The images shown are the ones with the smallest rms of the true residuals; for the mean image and first image of the set these are, respectively, 0.082 (MRNSD) – 0.088 (Tikhonov), and 0.109.

Fig. 7. Mean image $\zeta(m, n)$ (with $D_i(m, n)$ obtained using linear interpolation), MRNSD and Tikhonov deblurred versions of $\zeta(m, n)$, and MRNSD deblurred version of the first image of the set (see text) for the image shown in Fig. 5. The images shown are the ones with the smallest standard deviation of the true residuals; for the mean image and first image of the set these are, respectively, 0.076 (MRNSD) – 0.082 (Tikhonov) and 0.109.

Fig. 8. Original image, blurred version and blurred + noise version of the first image of the set (see text), and corresponding mean PSF. The images are 400×400 pixels, but only the central 320×320 pixels are shown (see text).

Fig. 9. Mean image $\zeta(m, n)$ (with $D_i(m, n)$ obtained using linear interpolation), MRNSD and Tikhonov deblurred versions of $\zeta(m, n)$, and MRNSD deblurred version of the first image of the set (see text) for the image shown in Fig. 8. The images shown are the ones with the smallest rms of the true residuals; for the mean image and first image of the set these are, respectively, 0.036 (MRNSD) – 0.036 (Tikhonov), and 0.043.

Fig. 10. Mean image $\zeta(m,n)$ (with $D_i(m,n)$ obtained using linear interpolation), MRNSD and Tikhonov deblurred versions of $\zeta(m,n)$, and MNRSD deblurred version of the first image of the set (see text) for the image shown in Fig. 5 and Poissonian noise (peak signal to noise ratio =40 dB). The images shown are the ones with the smallest rms of the true residuals; for the mean image and first image of the set these are, respectively, 0.076 (MRNSD) – 0.082 (Tikhonov) and 0.109.

Fig. 11. Mean image $\zeta(m,n)$ (with $D_i(m,n)$ obtained using linear interpolation), MRNSD and Tikhonov deblurred versions of $\zeta(m,n)$, and MRNSD deblurred version of the first image of the set (see text) for the image shown in Fig. 8 and Poissonian noise (peak signal to noise ratio =40 dB). The images shown are the ones with the smallest rms of the true residuals; for the mean image and first image of the set these are, respectively, 0.037 (MRNSD) – 0.037 (Tikhonov), and 0.044 (NB. here the values of the pixels have been converted into photon counts).

Fig. 12. Original image, observed image at 100 GHz, mean image $\zeta(m,n)$ (with $D_i(m,n)$ obtained using linear interpolation), and Tikhonov deblurred versions of $\zeta(m,n)$.

Appendix A: SVD approximation for spatially-variant blurs

In matrix notation, model (2) can be written in the form

$$\mathbf{g} = \mathcal{K}\mathbf{f} + \mathbf{w}, \quad (\text{A.1})$$

where $\mathcal{K}$ is a large $N^2 \times N^2$ ill-conditioned matrix which models the blurring phenomena, $\mathbf{f}$ is a vector of length N^2 representing the true image, and $\mathbf{g}$ is a vector representing the *observed* image. Because the problem is ill-posed, regularization methods are needed to compute an approximation of $\mathbf{f}$. A useful tool in constructing and analyzing regularization methods is the singular value decomposition (SVD). The problem is that $\mathcal{K}$ is a very large matrix, and it may not be computationally feasible to explicitly compute its SVD.

In Vio et al. (2003a,b) it has been shown that for spatially-invariant blurs, spectral decompositions using fast transforms (such as DFT and the Discrete Cosine Transform) can sometimes be used in place of the SVD, but this depends on the PSF and the boundary condition. If it is not appropriate to use fast transforms, then another approach we considered was to compute an approximation of the SVD from a Kronecker product factorization of $\mathcal{K}$.

For spatially-variant blurs, the situation is a bit more difficult since the fast transform-based decomposition methods are no longer a viable approach. However, here we show that the SVD approximation based on a Kronecker product factorization of $\mathcal{K}$ can be extended to the case of space variant blurs. This work builds on ideas initially proposed by Kamm & Nagy (1998).

A.1. Space Invariant Blurs

Suppose the blur is spatially-invariant, and let $\mathbf{P}$ be an image representing the PSF. Then it can be shown that if $\mathbf{P}$ is decomposed as:

$$\mathbf{P} = \sum_{k=1}^r \mathbf{a}_k \mathbf{b}_k^T, \quad (\text{A.2})$$

then the blurring matrix, $\mathcal{K}$, can be decomposed as

$$\mathcal{K} = \sum_{k=1}^r \mathbf{A}_k \otimes \mathbf{B}_k. \quad (\text{A.3})$$

Here the notation $\otimes$ denotes Kronecker product, which is defined as

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \cdots & a_{1n}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \cdots & a_{2n}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}\mathbf{B} & a_{n2}\mathbf{B} & \cdots & a_{nn}\mathbf{B} \end{bmatrix}. \quad (\text{A.4})$$

The matrices $\mathbf{A}_k$ and $\mathbf{B}_k$ are defined by the vectors $\mathbf{a}_k$ and $\mathbf{b}_k$, respectively. In particular, if the center of the PSF is located at pixel (i, j) , then

- Zero boundary conditions imply $\mathbf{A}_k$ and $\mathbf{B}_k$ are Toeplitz matrices:

$$\mathbf{A}_k = \begin{pmatrix} a_i^{(k)} & \cdots & a_1^{(k)} & & \\ \vdots & \ddots & & \ddots & \\ a_n^{(k)} & & & & a_1^{(k)} \\ & \ddots & & \ddots & \vdots \\ & & a_n^{(k)} & \cdots & a_j^{(k)} \end{pmatrix}$$

$$\text{and } \mathbf{B}_k = \begin{pmatrix} b_j^{(k)} & \cdots & b_1^{(k)} & & \\ \vdots & \ddots & & \ddots & \\ b_n^{(k)} & & & & b_1^{(k)} \\ & \ddots & & \ddots & \vdots \\ & & b_n^{(k)} & \cdots & b_j^{(k)} \end{pmatrix}. \quad (\text{A.5})$$

- Similarly, for periodic boundary conditions, $\mathbf{A}_k$ and $\mathbf{B}_k$ are circulant matrices.
- In the case of reflexive boundary conditions, $\mathbf{A}_k$ and $\mathbf{B}_k$ are Toeplitz-plus-Hankel matrices:

$$\mathbf{A}_k = \begin{pmatrix} a_i^{(k)} & \cdots & a_1^{(k)} & & \\ \vdots & \ddots & & \ddots & \\ a_n^{(k)} & & & & a_1^{(k)} \\ & \ddots & & \ddots & \vdots \\ & & a_n^{(k)} & \cdots & a_i^{(k)} \end{pmatrix} + \begin{pmatrix} a_{i+1}^{(k)} & \cdots & a_n^{(k)} & & \\ \vdots & \ddots & & \ddots & \\ a_n^{(k)} & & & & a_1^{(k)} \\ & \ddots & & \ddots & \vdots \\ & & a_1^{(k)} & \cdots & a_{i-1}^{(k)} \end{pmatrix}, \quad (\text{A.6})$$

and

$$\mathbf{B}_k = \begin{pmatrix} b_j^{(k)} & \cdots & b_1^{(k)} & & \\ \vdots & \ddots & & \ddots & \\ b_m^{(k)} & & & & b_1^{(k)} \\ & \ddots & & \ddots & \vdots \\ & & b_m^{(k)} & \cdots & b_j^{(k)} \end{pmatrix} + \begin{pmatrix} b_{j+1}^{(k)} & \cdots & b_m^{(k)} & & \\ \vdots & \ddots & & \ddots & \\ b_m^{(k)} & & & & b_1^{(k)} \\ & \ddots & & \ddots & \vdots \\ & & b_1^{(k)} & \cdots & b_{j-1}^{(k)} \end{pmatrix}. \quad (\text{A.7})$$

The factorization of $\mathbf{P}$ can be computed using the SVD. In particular, it is possible to compute

$$\mathbf{P} = \sum_{k=1}^r \sigma_k \mathbf{u}_k \mathbf{v}_k^T, \quad (\text{A.8})$$

where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$, and thus $\mathbf{a}_k = \sqrt{\sigma_k} \mathbf{u}_k$ and $\mathbf{b}_k = \sqrt{\sigma_k} \mathbf{v}_k$. Using the SVD in this way, we obtain the Kronecker product factorization of $\mathcal{K}$ given in Eq. (A.3) with the property that $\mathbf{A}_1$ has more information than $\mathbf{A}_2$, etc., and similarly for $\mathbf{B}_k$. In particular, in some measure, $\mathbf{A}_1 \otimes \mathbf{B}_1$ is the best Kronecker product factorization of $\mathcal{K}$. This observation is important because we can exploit the Kronecker product structure of matrices in our computation. In particular, the SVD of a Kronecker product can be computed as:

$$\begin{aligned} \mathbf{A} \otimes \mathbf{B} &= (\mathbf{U}_a \Sigma_a \mathbf{V}_a^T) \otimes (\mathbf{U}_b \Sigma_b \mathbf{V}_b^T) \\ &= (\mathbf{U}_a \otimes \mathbf{U}_b) (\Sigma_a \otimes \Sigma_b) (\mathbf{V}_a \otimes \mathbf{V}_b)^T. \end{aligned} \quad (\text{A.9})$$

We do not need to explicitly form the big matrices $(\mathbf{U}_a \otimes \mathbf{U}_b)$, etc., because the properties of Kronecker products allow us to organize the computations so that we need only work with the smaller matrices $\mathbf{U}_a, \mathbf{U}_b$, etc.

The above remarks suggest that an SVD approximation can be obtained by computing the SVD of $\mathbf{A}_1 \otimes \mathbf{B}_1$. It is possible, though, to incorporate some information from $\mathbf{A}_2, \mathbf{A}_3, \dots$ and $\mathbf{B}_2, \mathbf{B}_3, \dots$. Following an approach outlined in Kamm & Nagy (2000); Nagy et al. (2004a), this can be done as follows:

- Let $\mathbf{A}_1 = \mathbf{U}_a \Sigma_a \mathbf{V}_a^T$ and $\mathbf{B}_1 = \mathbf{U}_b \Sigma_b \mathbf{V}_b^T$. We use $\mathbf{U}_a \otimes \mathbf{U}_b$ and $\mathbf{V}_a \otimes \mathbf{V}_b$ as the approximate singular vectors of $\mathcal{K}$. The approximate singular values are obtained by computing the diagonal matrix:

$$\begin{aligned} \Sigma &= \text{diag} \left[(\mathbf{U}_a \otimes \mathbf{U}_b)^T \mathcal{K} (\mathbf{V}_a \otimes \mathbf{V}_b) \right] \\ &= \text{diag} \left[(\mathbf{U}_a \otimes \mathbf{U}_b)^T \left(\sum_{k=1}^r (\mathbf{A}_k \otimes \mathbf{B}_k) \right) (\mathbf{V}_a \otimes \mathbf{V}_b) \right] \\ &= \text{diag} \left[\sum_{k=1}^r (\mathbf{U}_a \otimes \mathbf{U}_b)^T (\mathbf{A}_k \otimes \mathbf{B}_k) (\mathbf{V}_a \otimes \mathbf{V}_b) \right] \\ &= \text{diag} \left[\sum_{k=1}^r (\mathbf{U}_a^T \mathbf{A}_k \mathbf{V}_a) \otimes (\mathbf{U}_b^T \mathbf{B}_k \mathbf{V}_b) \right] \\ &= \sum_{k=1}^r \text{diag}(\mathbf{U}_a^T \mathbf{A}_k \mathbf{V}_a) \otimes \text{diag}(\mathbf{U}_b^T \mathbf{B}_k \mathbf{V}_b) \end{aligned} \quad (\text{A.10})$$

- That is, $\mathcal{K} \approx (\mathbf{U}_a \otimes \mathbf{U}_b) \Sigma (\mathbf{V}_a \otimes \mathbf{V}_b)^T$.

A.2. Space variant blurs

In the case of space variant blurs, a single PSF does not provide enough information to obtain a complete representation of the matrix, $\mathcal{K}$. One PSF provides only information about one column (or row) of $\mathcal{K}$. In order to get all N^2 columns, we would need N^2 PSFs, each one centered at a different pixel location. It would be extremely expensive to find all of these PSFs. So instead, we consider a simplifying assumption that the blur is locally spatially-invariant (this approach has been used in Boden et al. 1996; Faisal et al. 1995; Nagy & O’Leary 1998). Specifically, suppose that the image domain is partitioned into $p \times q$ rectangular regions, each with a PSF,

$$\begin{array}{|c|c|c|} \hline \mathbf{P}_{11} & \cdots & \mathbf{P}_{1q} \\ \hline \vdots & & \vdots \\ \hline \mathbf{P}_{p1} & \cdots & \mathbf{P}_{pq} \\ \hline \end{array}$$

Suppose the image has $N \times N$ pixels, and define the dimensions of the ij th region as $l_i \times m_i$, where $l_1 + \dots + l_p = m_1 + \dots + m_q = n$. Further, let $\mathbf{D}_{ij}$ be a $\{0, 1\}$ diagonal matrix representing the pixels in the ij th region. That is, $\mathbf{D}_{ij} = \text{diag}(d_1, \dots, d_N)$, where $d_k = d_{(s-1)n+t} = 1$ if the point (s, t) is in the ij th region; otherwise, $d_k = 0$. With this notation, we represent the matrix $\mathcal{K}$ as

$$\mathcal{K} = \sum_{i=1}^p \sum_{j=1}^q \mathbf{D}_{ij} \mathcal{K}_{ij}, \quad (\text{A.11})$$

where $\mathcal{K}_{ij}$ is a matrix defined by the spatially-invariant PSF $\mathbf{P}_{ij}$.

Now using the Kronecker product factorization for spatially-invariant blurs described in the previous section, we obtain (without loss of generality, we assume each of the PSFs has rank equal to r):

$$\mathcal{K} = \sum_{i=1}^p \sum_{j=1}^q \mathbf{D}_{ij} \left(\sum_{k=1}^r \mathbf{A}_k^{(i)} \otimes \mathbf{B}_k^{(j)} \right), \quad (\text{A.12})$$

where the matrices $\mathbf{A}_k^{(i)}$ and $\mathbf{B}_k^{(j)}$ are Toeplitz or Toeplitz-plus-Hankel matrices (depending on the boundary conditions), as described in the previous section. We can simplify this expression by observing that each of the diagonal matrices, $\mathbf{D}_{ij}$ is separable, and can be written as:

$$\mathbf{D}_{ij} = \widehat{\mathbf{D}}_i \otimes \tilde{\mathbf{D}}_j, \quad (\text{A.13})$$

where

$$\begin{aligned} \widehat{\mathbf{D}}_1 &= \text{diag}(\overbrace{1, \dots, 1}^{l_1}, 0, \dots, 0), \\ \widehat{\mathbf{D}}_2 &= \text{diag}(\overbrace{0, \dots, 0}^{l_1}, \overbrace{1, \dots, 1}^{l_2}, 0, \dots, 0), \\ &\vdots \\ \widehat{\mathbf{D}}_p &= \text{diag}(0, \dots, 0, \overbrace{1, \dots, 1}^{l_p}), \\ \tilde{\mathbf{D}}_1 &= \text{diag}(\overbrace{1, \dots, 1}^{m_1}, 0, \dots, 0) \\ \tilde{\mathbf{D}}_2 &= \text{diag}(\overbrace{0, \dots, 0}^{m_1}, \overbrace{1, \dots, 1}^{m_2}, 0, \dots, 0) \\ &\vdots \\ \tilde{\mathbf{D}}_q &= \text{diag}(0, \dots, 0, \overbrace{1, \dots, 1}^{m_q}). \end{aligned} \quad (\text{A.14})$$

Thus,

$$\begin{aligned} \mathcal{K} &= \sum_{i=1}^p \sum_{j=1}^q \mathbf{D}_{ij} \left(\sum_{k=1}^r \mathbf{A}_k^{(i)} \otimes \mathbf{B}_k^{(j)} \right) \\ &= \sum_{i=1}^p \sum_{j=1}^q (\widehat{\mathbf{D}}_i \otimes \tilde{\mathbf{D}}_j) \left(\sum_{k=1}^r \mathbf{A}_k^{(i)} \otimes \mathbf{B}_k^{(j)} \right) \\ &= \sum_{k=1}^r \sum_{i=1}^p \sum_{j=1}^q (\widehat{\mathbf{D}}_i \mathbf{A}_k^{(i)} \otimes \tilde{\mathbf{D}}_j \mathbf{B}_k^{(j)}) \\ &= \sum_{k=1}^r \left[\left(\sum_{i=1}^p \widehat{\mathbf{D}}_i \mathbf{A}_k^{(i)} \right) \otimes \left(\sum_{j=1}^q \tilde{\mathbf{D}}_j \mathbf{B}_k^{(j)} \right) \right] \\ &= \sum_{k=1}^r (\mathbf{A}_k \otimes \mathbf{B}_k), \end{aligned} \quad (\text{A.15})$$

where

$$\mathbf{A}_k = \sum_{i=1}^p \widehat{\mathbf{D}}_i \mathbf{A}_k^{(i)} \quad \text{and} \quad \mathbf{B}_k = \sum_{j=1}^q \tilde{\mathbf{D}}_j \mathbf{B}_k^{(j)}. \quad (\text{A.16})$$

Once $\mathbf{A}_k$ and $\mathbf{B}_k$ are constructed, the approximate SVD of $\mathcal{K}$ is computed just as it was in the spatially-invariant case described in Sect. A.1.