Multichannel Image Recovery

*Nikolas P. Galatsanos, **Miles N. Wernick, ***Aggelos K. Katsaggelos and ****Rafael Molina

(*)Department of Computer Science, University of Ioannina,
P.O. Box 1186, GR 45110, Ioannina, Greece
(**)Dept. of Electrical and Computer Engineering
Illinois Institute of Technology
Chicago, IL 60616
(***)Dept. of Electrical and Computer Engineering
Northwestern University
Evanston, IL 60208
(****) Dept. of Computer Science and Artificial Intelligence
University of Granada, 18071 Granada

1   Introduction

Color images, video images, medical images obtained by multiple scanners, and multispectral satellite images consist of multiple image frames or *channels* (Fig. 1). These image channels depict the same scene or object observed either by different sensors or at different times, and thus have substantial commonality among them. We use the term *multichannel image* to refer to any collection of image channels which are not identical, but exhibit strong between-channel correlations.

In this chapter we focus on the problem of image recovery as it applies specifically to multi-channel images. *Image recovery* refers to the computation of an image from observed data that alone do not uniquely define the desired image. Important examples are image denoising, image deblurring, decoding of compressed images, and medical image reconstruction.

In image recovery, ambiguities in inferring the desired image from the observations usually arise from uncertainty produced by noise. These ambiguities can only be reduced if, in addition to information provided by the observed data, one also has prior knowledge about the desired image. In many applications the most powerful piece of prior information is that the desired image is smooth (spatially correlated), whereas the noise is not. Multichannel images offer the possibility of exploiting correlations *between* the channels in addition to those within each channel. By utilizing this extra information, multichannel image recovery can yield tremendous benefits over separate recovery of the component channels.

In a broad category of image recovery techniques, the image is computed by optimizing an objective function that quantifies correspondence of the image to the observed data as well as prior knowledge about the true image. Two frameworks have been developed to describe the use of prior information as an aid in image recovery: the deterministic formulation with regularization [8] and the statistical formulation [4]. Conceptually, these frameworks are quite different, but in many applications they lead to identical algorithms.

In the deterministic formulation, the problem is posed in terms of inversion of the observation model, with regularization used to improve the stability of the solution (for more details see Chapter 3.11). In the statistical formulation, the statistics of the noise and the desired image are incorporated explicitly, and are used to describe the desired characteristics of the image. In principle, these formulations apply equally well to multichannel images and single-channel images but in practice two significant problems must be addressed.

First, an appropriate model must be developed to express the relationships among the image channels. Regularization and statistical methods reflecting within-channel relationships are much better developed than those describing between-channel relationships, and are often easier to work with. For example, while the power spectrum (the Fourier transform of the autocorrelation) is a useful statistical descriptor for one channel, the cross power spectrum (the Fourier transform of the cross-correlation) describing multiple channels is less tractable because it is complex. Second, a multichannel image has many more pixels than each of its channels, so approaches that minimize the computations are typically sought. Several such approaches are described in the following sections.

The goal of this chapter is to provide a concise summary of the theory of multichannel image recovery. We classify multichannel recovery methods into two broad approaches, each of which is illustrated through a practical application. In the first, which we term the *explicit* approach, all the channels of the multichannel image are processed collectively, and regularization operators or prior distributions are used to express the between- and within-channel relationships. In the second, which we term the *implicit* approach, the same effect is obtained indirectly by: 1) applying a Karhunen-Loève (KL) transform that decorrelates the channels, 2) recovering the channels separately in a single-channel fashion, and 3) inverting the KL transform. This approach has a substantial computational advantage over the explicit approach, as we explain later, but it can only be applied in certain situations.

The rest of this chapter is organized as follows. In Section 2 we present the multichannel observation model and basic image recovery approaches are reviewed in Section 3. In Section 4 we describe the explicit approach and illustrate it using the example of restoration of video image sequences. In Section 5, we explain the implicit approach and illustrate it using the example of reconstruction of time-varying medical images. We conclude with a summary in Section 6.

2 Imaging Model

Throughout this chapter we use symbols in boldface type to denote multichannel quantities. We assume the following discrete model for multichannel imaging:

$$\mathbf{g} = \mathbf{H}\mathbf{f} + \mathbf{n}, \tag{1}$$

where $\mathbf{g}$, $\mathbf{f}$ and $\mathbf{n}$ are vectors representing the observed (degraded) multichannel data, the true multichannel image, and random noise, respectively, and matrix $\mathbf{H}$ denotes the linear multi-

channel degradation operator. Lexicographic ordering is used to represent the images as vectors by stacking all their rows or columns in one long vector. Then, each multichannel image is a concatenation of its K component channels, i.e.,

$$\mathbf{g} = \begin{bmatrix} g_1 \\ g_2 \\ \vdots \\ g_K \end{bmatrix}, \quad \mathbf{f} = \begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_K \end{bmatrix}, \quad \mathbf{n} = \begin{bmatrix} n_1 \\ n_2 \\ \vdots \\ n_K \end{bmatrix}, \tag{2}$$

where g_k, f_k, and n_k denote individual channels of the observations, the true image, and noise, respectively. In its most general form, the linear multichannel degradation operator can be written as

$$\mathbf{H} = \begin{bmatrix} H_{11} & H_{12} & \cdots & H_{1K} \\ H_{21} & H_{22} & \cdots & H_{2K} \\ \cdots & \cdots & \cdots & \cdots \\ H_{K1} & H_{N2} & \cdots & H_{KK} \end{bmatrix}, \tag{3}$$

where the diagonal blocks represent within-channel degradations, and the off-diagonal blocks represent between-channel degradations. We will assume that each source channel f_i has $N \times M$ pixels and that each channel of the observations g_i has $L \times P$ pixels. Therefore, f_i is a $MN \times 1$ vector and g_i is a $LP \times 1$ vector; consequently, H_{ij} is a $LP \times MN$ matrix and $\mathbf{H}$ is a $KLP \times KMN$ matrix. If the degradation is shift-invariant the product $H_{ij} f_j$ should in principle represent an ordinary linear convolution operation. However, we would like to use the discrete Fourier transform (DFT) to compute the convolutions rapidly, and this can be done only for circular convolutions. Fortunately, circular convolution and linear convolution produce the same result if we first embed each image in a larger array of zeros (see [10], p. 145). A matrix H_{ij} that, when multiplied by an image f_j, produces the effect of circular convolution has what is known as a *block-circulant* structure. Block circulant matrices are diagonalized by the DFT [10], thus leading to simplified calculations.

For purposes of notational simplicity, we assume throughout that the multichannel source image $\mathbf{f}$ and noise $\mathbf{n}$ have zero mean. The source image usually does not obey this assumption in reality, but the equations can easily be modified to accommodate a non-zero mean by introducing appropriate corrections.

3 Multichannel Image Estimation Approaches

Image recovery is most often achieved by constructing an objective function to quantify the quality of an image estimate, then optimizing that function to obtain the desired result. Some important objective functions and solutions from estimation theory are reviewed in this section.

3.1 Linear minimum mean square error (LMMSE) estimation

The mean square error (MSE), defined as

$$MSE(\hat{\mathbf{f}}) = E\{||\mathbf{f} - \hat{\mathbf{f}}||^2\}, \tag{4}$$

is a measure of the quality of the image estimate $\hat{\mathbf{f}}$. Here, $E\{\cdot\}$ denotes the expectation operator. Among the images that are a linear function of the data (i.e., $\hat{\mathbf{f}} = \mathbf{Ag}$ where $\mathbf{A}$ a matrix), the image that minimizes the MSE is known as the linear minimum mean square error (LMMSE) estimate $\mathbf{f}_{LMMSE}$. Assuming $\mathbf{f}$ and $\mathbf{n}$ are uncorrelated [14], the LMMSE solution is found by finding the matrix $\mathbf{A}$ that minimizes the MSE. The resulting LMMSE solution is

$$\mathbf{f}_{LMMSE} = \mathbf{C}_f \mathbf{H}^T (\mathbf{H} \mathbf{C}_f \mathbf{H}^T + \mathbf{C}_n)^{-1} \mathbf{g}, \tag{5}$$

where $\mathbf{C}_n$ is the covariance matrix of the multichannel noise vector $\mathbf{n}$, and $\mathbf{C}_f$ is the $KNM \times KNM$ covariance matrix of the multichannel image vector $\mathbf{f}$, defined as

$$\mathbf{C}_f = E\{\mathbf{f}\,\mathbf{f}^T\} = \begin{bmatrix} C_{11} & C_{12} & \cdots & C_{1K} \\ C_{21} & C_{22} & \cdots & C_{2K} \\ \cdots & \cdots & \cdots & \cdots \\ C_{K1} & C_{K2} & \cdots & C_{KK} \end{bmatrix}, \tag{6}$$

where $C_{ij} = E\{f_j f_i^T\}$.

Using the matrix inversion lemma [14] it is easy to show that

$$\mathbf{C}_f \mathbf{H}^T (\mathbf{H} \mathbf{C}_f \mathbf{H} + \mathbf{C}_n)^{-1} = (\mathbf{H}^T \mathbf{C}_n^{-1} \mathbf{H} + \mathbf{C}_f^{-1})^{-1} \mathbf{H}^T \mathbf{C}_n^{-1}. \tag{7}$$

Thus, we can also write

$$\mathbf{f}_{LMMSE} = (\mathbf{H}^T \mathbf{C}_n^{-1} \mathbf{H} + \mathbf{C}_f^{-1})^{-1} \mathbf{H}^T \mathbf{C}_n^{-1} \mathbf{g}. \tag{8}$$

3.2 Regularized weighted least-squares (RWLS) estimation

The weighted least-squares (WLS) estimate of $\mathbf{f}$ is

$$\mathbf{f}_{WLS} = \arg\min_{\mathbf{f}}\{(\mathbf{Hf} - \mathbf{g})^T \mathbf{C}_n^{-1}(\mathbf{Hf} - \mathbf{g})\}, \tag{9}$$

where $\arg\min_{\mathbf{f}}\{J(\mathbf{f})\}$ denotes the vector $\mathbf{f}$ that minimizes $J(\mathbf{f})$. Here, $\mathbf{C}_n$ is the covariance matrix of the noise, which in this context is assumed to be diagonal. In the presence of noise, the WLS solution will usually be very noisy. This occurs because the matrix $\mathbf{H}$ is often ill-conditioned or singular, therefore some of its singular values are close to or equal to zero. In this case, the solution is unstable and highly sensitive to noise. Regularization is a well-known solution to this instability, in which the ill-posed problem is replaced by a well-posed problem for which the solution is an acceptable approximation (see [4] and Chapter 3.11).

In the WLS formulation, regularization can be achieved by adding to the WLS functional a term that takes on large values if the image is noisy. A term that achieves this is $\|\mathbf{Qf}\|^2$, where $\mathbf{Q}$ is a high-pass filter operator. Incorporating this term, the regularized WLS (RWLS) estimate is obtained as follows:

$$\mathbf{f}_{RWLS} = \arg\min_{\mathbf{f}}\{(\mathbf{Hf} - \mathbf{g})^T \mathbf{C}_n^{-1}(\mathbf{Hf} - \mathbf{g}) + \lambda\|\mathbf{Qf}\|^2\}, \tag{10}$$

in which λ, known as the regularization parameter, controls the trade-off between fidelity to the data (reflected by the first term of the objective function) and smoothness of the estimate (reflected by the second term) [8]. Solving for $\mathbf{f}_{RWLS}$ we obtain

$$\mathbf{f}_{RWLS} = (\mathbf{H}^T \mathbf{C}_n^{-1}\mathbf{H} + \lambda\mathbf{Q}^T\mathbf{Q})^{-1}\mathbf{H}^T\mathbf{C}_n^{-1}\mathbf{g}. \tag{11}$$

Note that the RWLS and the LMMSE estimates are equivalent when $\mathbf{C}_f^{-1} = \lambda\mathbf{Q}^T\mathbf{Q}$.

A special case of RWLS estimation occurs if the noise is white, i.e., if $\mathbf{C}_n = \sigma^2\mathbf{I}$. In this case the RWLS functional reduces to the regularized least-squares (RLS) functional:

$$\mathbf{f}_{RLS} = \arg\min_{\mathbf{f}}\{J(\mathbf{f})\} = \arg\min_{\mathbf{f}}\{(\mathbf{Hf} - \mathbf{g})^T(\mathbf{Hf} - \mathbf{g}) + \lambda\|\mathbf{Qf}\|^2\}. \tag{12}$$

3.3 Multichannel regularization

In multichannel image recovery, the operator $\mathbf{Q}$ enforces smoothness not only within each image channel, but also between the channels, thus achieving an additional measure of noise suppression,

and often producing dramatically better images. In ordinary single-channel recovery of two-dimensional (2D) images, one can define Q as a discrete two-dimensional Laplacian operator, which represents 2D convolution with the following mask

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{pmatrix}. \tag{13}$$

For multichannel recovery, one can use the three-dimensional (3D) Laplacian $\mathbf{Q}$, which implies correlations between channels. The 3D Laplacian [7] can be performed as a 3D convolution which can be written in equation form for the lth channel as:

$$
\begin{aligned}
[\mathbf{Q}\,\mathbf{f}]_l(i,j) \;=\; & -6f_l(i,j) + f_l(i-1,j) + f_l(i+1,j) + f_l(i,j-1) + f_l(i,j+1) \\
& f_{l-1}(i,j) + f_{l+1}(i,j)
\end{aligned} \tag{14}
$$

For the multichannel problem we could also use the 2-D mask in Eq. (13) or the first derivative mask which would produce

$$\|\mathbf{Qf}\|^2 = \sum_c \sum_{i,j} \left((f_c(i+1,j) - f_c(i,j))^2 + (f_c(i,j+1) - f_c(i,j))^2 \right), \tag{15}$$

however, this regularization, as well as the one defined by Eq. (13), does not take into account the correlations between channels and corresponds to an independent channel reconstruction.

In order to relate channels, not by grey levels but by the presence of similar regions in all channels, we first modify Eq. (15) by smoothing each channel only within homogeneous regions. This selective smoothing is carried out by the introduction of a line process $l_c[(i,j),(u,v)]$, where $(u,v) \in \mathcal{N}(i,j)$, $\mathcal{N}(i,j) = \{(i-1,j),(i+1,j),(i,j-1),(i,j+1)\}$, which takes the value zero if pixels (i,j) and (u,v) in channel c are not separated by an active line and one otherwise. Then, we replace $\|\mathbf{Qf}\|^2$ in Eq. (15) by a penalty function, $Q(f,l)$, that now depends on the multichannel image f and the line process l,

$$
\begin{aligned}
Q(f,l) \;=\; & \sum_c \sum_{i,j} \Big[(f_c(i+1,j) - f_c(i,j))^2 (1 - l_c[(i,j),(i+1,j)]) \\
& + \; (f_c(i,j+1) - f_c(i,j))^2 (1 - l_c[(i,j),(i,j+1)]) \\
& + \; \tau l_c[(i,j),(i+1,j)] + \tau l_c[(i,j),(i,j+1)] \Big].
\end{aligned}
$$

The introduction of an active line element in the position $[(i,j),(u,v)]$ with $(u,v) \in \mathcal{N}(i,j)$ is penalized by the term $\tau l_c[(i,j),(u,v)]$ since otherwise the above equation would obtain its

7

minimum value by setting all line elements equal to one. The intuitive interpretation of this line process is simple; it acts as an activator or inhibitor of the relation between two neighbor pixels depending on whether or not the pixels are separated by an edge, see [9] and [11] for details.

We want now, see [16] for details, to increase the probability of a new active line element in the position $[(i,j),(u,v)]$ in a given channel if the other channels have a line in the same position. To do so we include in the regularization term $\epsilon^{cc'} l_c[(i,j),(i+1,j)]l_{c'}[(i,j),(i+1,j)]$ and also $\epsilon^{cc'} l_c[(i,j),(i,j+1)]l_{c'}[(i,j),(i,j+1)]$ with $c,c' = 1,\ldots,L$, $c \neq c'$ and $\epsilon^{cc'} > 0$ and finally write

$$
\begin{aligned}
Q(f,l) \;=\; & \sum_c \sum_{i,j} \Big((f_c(i+1,j) - f_c(i,j))^2 (1 - l_c[(i,j),(i+1,j)]) \\
& + \; (f_c(i,j+1) - f_c(i,j))^2 (1 - l_c[(i,j),(i,j+1)]) \\
& + \; \tau l_c[(i,j),(i+1,j)] + \tau l_c[(i,j),(i,j+1)] \Big) \\
& - \; \sum_{\substack{c'=1 \\ c' \neq c}}^{L} \epsilon^{cc'} \sum_{i,j} l_c[(i,j),(i+1,j)]l_{c'}[(i,j),(i+1,j)] \\
& - \; \sum_{\substack{c'=1 \\ c' \neq c}}^{L} \epsilon^{cc'} \sum_{i,j} l_c[(i,j),(i,j+1)]l_{c'}[(i,j),(i,j+1)]
\end{aligned}
\tag{16}
$$

The introduction of these cross terms will help to recognize the same objects in all the channels even if they do not have similar intensities. Note that if $\epsilon^{cc'} = 0$, $\forall c,c'$ no edge information will be passed between channels. Furthermore, if $\epsilon^{cc'}$ is very large the restoration will tend to have the same edge map in all channels.

Though they would yield excellent results, the closed-form solutions given in the previous section usually cannot be computed directly because of the large number of dimensions of multichannel images. For example, to restore a three-channel color image having 512×512 pixels per channel, direct computation of $\mathbf{f}_{LMMSE}$ using Eq. (5) would require the inversion of matrices of dimension $786,432 \times 786,432$. To sidestep this dimensionality problem, computationally efficient algorithms must be designed. Next, we present two such approaches that lead to practical multichannel image recovery techniques. Unfortunately, the use of the regularization term in Eq. (16) prevents us from using computationally efficient algorithms and other alternative algorithms have to be used, see [16] for details.

4 Explicit Multichannel Recovery Approaches

4.1 Space-Invariant Multichannel Recovery

The difficulty of directly implementing the solutions reviewed in the previous section lies in the complexity of inverting large matrices. In this section we show that, if the multichannel imaging system is space-invariant and the noise and signal are stationary, then the required inversion is easily performed by using the fact that block-circulant matrices are diagonalized by the discrete Fourier transform (DFT) [10]. To simplify our discussion of space-invariant multichannel imaging we assume in this section that the observed and true image channels are square and have the same number of pixels, i.e., $M = N = L = P$.

Let us assume that the imaging system is space-invariant and that the channels of the source image $\mathbf{f}$ are jointly stationary, i.e.,

$$E\{f_i(x,y)f_j(x',y')\} = [\mathbf{C}_f]_{ij}(x-x', y-y'), \tag{17}$$

where $f_i(x,y)$ denotes pixel (x,y) of image channel f_i, and $[\mathbf{C}_f]_{ij}$ is the covariance matrix of f_i and f_j. Let us make the same assumption about the noise channels n_i. In this case the covariance matrices $[\mathbf{C}_f]_{ij}$ and $[\mathbf{C}_n]_{ij}$ can be approximated by $M^2 \times M^2$ block-circulant matrices. The matrices $\mathbf{H}$, $\mathbf{C}_f$, and $\mathbf{C}_n$ are composed of $M^2 \times M^2$ block-circulant blocks, but they are not themselves block-circulant.

Now let us define the multichannel DFT as $\mathbf{W} = diag\{W, W, \cdots, W\}$, where W is a $M^2 \times M^2$ matrix representing the two-dimensional DFT, and $\mathbf{W}$ has K blocks. Using $\mathbf{W}$ to transform the LMMSE solution in Eq. (5) we obtain

$$\mathbf{Wf}_{LMMSE} = \mathbf{WC}_f\mathbf{W}^{-1}\mathbf{WH}^T\mathbf{W}^{-1}(\mathbf{WHW}^{-1}\mathbf{WC}_f\mathbf{W}^{-1}\mathbf{WHW}^{-1} + \mathbf{WC}_n\mathbf{W}^{-1})^{-1}\mathbf{Wg} \tag{18}$$

or, equivalently,

$$\mathbf{F}_{LMMSE} = \mathbf{D}_{C_f}(\mathbf{D}_H)^h\mathbf{\Delta}^{-1}\mathbf{G} \tag{19}$$

where

$$\mathbf{\Delta} = \mathbf{D}_H\mathbf{D}_{C_f}(\mathbf{D}_H)^h + \mathbf{D}_{C_n} \tag{20}$$

and h denotes the Hermitian of a matrix. The $KM^2 \times 1$ vectors $\mathbf{F}_{LMMSE}$ and $\mathbf{G}$ are the DFTs of the multichannel vectors $\mathbf{f}_{LMMSE}$ and $\mathbf{g}$, respectively. The matrices $\mathbf{D}_{C_f}$, $\mathbf{D}_{C_n}$ and $\mathbf{D}_H$ are obtained by using $\mathbf{W}$ to transform $\mathbf{C}_f$, $\mathbf{C}_n$ and $\mathbf{H}$, repectively (e.g., $\mathbf{D}_{C_f} = \mathbf{W}^h\mathbf{C}_f\mathbf{W}$).

The matrix $\mathbf{\Delta}$ has a special form which allows the inversion to be readily performed, thus the difficulty of computing $\mathbf{f}_{LMMSE}$ has been eliminated.

Any matrix $\mathbf{C}$ having block-circulant blocks C_{ij} can be transformed using the multichannel DFT into a matrix $\mathbf{D}$ having diagonal blocks as follows:

$$\mathbf{D} = \mathbf{WCW}^{-1} = \begin{bmatrix} WC_{11}W^{-1} & WC_{12}W^{-1} & \cdots & WC_{1K}W^{-1} \\ WC_{21}W^{-1} & WC_{22}W^{-1} & \cdots & WC_{2K}W^{-1} \\ \cdots & \cdots & \cdots & \cdots \\ WC_{K1}W^{-1} & WC_{K2}W^{-1} & \cdots & WC_{KK}W^{-1} \end{bmatrix}. \tag{21}$$

Although the blocks of $\mathbf{D}$ are diagonal, $\mathbf{D}$ is not itself diagonal. Any matrix having this property is termed a *non-diagonal block-diagonal* (NDBD) matrix. Matrices $\mathbf{D}_{C_f}$, $\mathbf{D}_{C_n}$, and $\mathbf{D}_H$ are also NDBD [5, 6].

NDBD matrices have two useful properties that lead to a tractable method for inverting $\mathbf{\Delta}$ and thus obtaining the LMMSE solution in Eq. (19). First, the set of NDBD matrices is closed under addition, multiplication, and inversion [5, 6]. Therefore, because $\mathbf{D}_H$ and $\mathbf{D}_{C_f}$ are NDBD matrices, so is $\mathbf{\Delta}$. Second, a $KM \times KM$ NDBD matrix such as $\mathbf{\Delta}$ can be rearranged into a matrix having M nonzero $K \times K$ blocks along its diagonal by applying a row operation transformation $\mathbf{T}$ [6, 13] to obtain $\mathbf{T\Delta T}^T = diag\{R_1, R_2, \cdots, R_M\}$, where each R_j is a general $K \times K$ matrix.

Once transformed, the originally intractable problem of inverting the $KM \times KM$ matrix $(\mathbf{H}C_f\mathbf{H}^T + \mathbf{C}_n)$ in Eq. (5) is reduced to one of separately inverting the $K \times K$ blocks R_j, $j = 1, \ldots, M$, of $\mathbf{T\Delta T}^T$. Because the number of channels K is usually much smaller than the number pixels M in each channel, the inversion problem is greatly simplified.

4.2   Numerical Experiment

A numerical experiment with color images is shown to demonstrate the improvement which results from the application of multi-channel as compared to single-channel LMMSE restoration. For this experiment different distortions were applied to each of the red (R), green (G), and blue (B) channels of the original image and is shown at the upper-left of Fig. (2). The red channel was blurred by vertical motion blur over 7 pixels, the green channel by horizontal blur over 9 pixels, and the blue channel by a 7×7 pill-box blur. In all cases the blurs were symmetric around the origin.

The variance of the noise added to the blurred data is defined using the blurred signal-to-noise ratio ($BSNR$) metric. These metrics are given per channel i

$$BSNR \;=\; 10\log_{10}\frac{||H_i f_i||^2}{M\sigma^2}dB \tag{22}$$

where M the total number of pixel in f_i and σ^2 is the variance of the additive noise.

Noise was added to all three channels with corresponding $BSNRs$ of 20 dB, 30 dB, and 40 dB. The degraded image is shown in the upper-right of Fig. (2). The degraded image was restored using a single-channel LMMSE filter to restore each channel independently, and the multi-channel LMMSE filter. In both cases the required power spectra and cross-power spectra were evaluated using the original image, to establish the upper bound of perfomance, as well as, the avaible noisy-blurred image, as a more realistic scenario. They were computed in all cases using Daniell's Periodogram (the regular periodogram was spatially averaged using a 5×5 window).

The results of the single-channel LMMSE restoration are shown in the middle of Fig. (2) left (use of original image for power spectra estimation) and right (use of degraded image for power spectra estimation).

The results of the multichannel LMMSE restoration are shown in Fig. (2) bottom, left (use of original image for cross-power spectra estimation) and left (use of degraded image for cross-power spectra estimation). From these experiments it is clear that multichannel restoration produces visually more pleasing results than than single-channel restoration.

The regularization term in Eq. (16) was used to restore a multichannel astronomical image. Although the original image size is 512×512 pixels, only a small 192×192 region of interest is presented in Fig. 3. The observed image, depicted in Fig. 3a, corresponds to three different images of the same object taken at different wavelength. The image was scaled to the range $[0, 255]$ for printing since the range of the original bands was extremely low.

Since the images were taken at different times no cross-channel blurring is present. The blurring functions H_{cc} correspond to the mask defined by

$$h_{cc}(i, j) \propto (1 + [i^2 + j^2]/r^2)^{-\delta}, \quad c \in \{R, G, B\}.$$

We found $\delta \sim 3$ and $r \sim 3.4$ pixels in all the channels.

Fig. 3b shows the reconstruction using the regularization term in Eq. (16).

4.3 Space-Variant Multichannel Recovery Approaches

In many cases the degradation and/or the regularization/covariance matrix may not be space-invariant. In such cases the frequency-domain approach described in the previous section cannot be applied because the matrices involved are not NDBD and thus direction inversion of $\mathbf{\Delta}$ in Eq. (19) is not possible. Instead, the RLS solution $\hat{\mathbf{f}}$ that minimizes Eq. (12) must be computed iteratively. Taking the gradient of $J(\mathbf{f})$ in Eq. (12) yields:

$$(\mathbf{H}^T\mathbf{H} + \lambda\mathbf{Q}^T\mathbf{Q})\hat{\mathbf{f}} = \mathbf{H}^T\mathbf{g}. \tag{23}$$

This equation can be solved using the method of successive approximations [18] which yields the following iteration:

$$\begin{aligned}
\hat{\mathbf{f}}^{(0)} &= \mathbf{0} \\
\hat{\mathbf{f}}^{(k+1)} &= \hat{\mathbf{f}}^{(k)} + \alpha\left[\, \mathbf{H}^T\mathbf{g} - (\,\mathbf{H}^T\mathbf{H} + \lambda\mathbf{Q}^T\mathbf{Q}\,)\hat{\mathbf{f}}^{(k)} \,\right],
\end{aligned} \tag{24}$$

where $\hat{\mathbf{f}}^k$ is the image estimate at iteration k, and α, known as the *relaxation parameter*, is a scalar that controls the convergence properties of the iteration. It is easy to verify that a stationary point of this iteration satisfies Eq. (23).

4.4 Application to restoration of moving image seqences

With the recent explosion of multimedia applications, the restoration of image sequences is becoming an increasingly important problem. The purpose of image-sequence restoration is to recover information lost during image sensing, recording, transmission and storage. Usually, image sequences consist of image frames of the same object or scene taken at closely spaced time intervals; therefore, they often exhibit a high degree of between-frame correlation. In the context of multichannel image recovery, we refer to the image frames as *channels*.

The correlation structure in an image sequence is often much more complicated than in a still color image because there is motion between frames. The *displacement vector* (DV) represents the motion of an image patch from one frame to the next, and the *displacement vector field* (DVF), which describes the motion of various pixels, is indispensable for describing the between-frame correlations, for details on DVF estimation see . For example, if the image patch occupying pixel (i,j) in frame l has displacement vector (m,n), it appears in pixel $(i+m, j+n)$ in frame $l+1$. Thus, there will be strong between-frame correlation between $f_l(i,j)$ and $f_{l+1}(i+m, j+n)$. The

correlation structure described by the DVF is not space-invariant in most situations; therefore, in these cases, the frequency-domain approach described previously cannot be applied.

To accommodate motion in the RLS formulation, the regularization operator must be modified to reflect the fact that a pixel in frame l is not necessarily correlated with the same pixel in frames $l + 1$ and $l - 1$, but rather with pixels that are offset by the corresponding displacement vectors. To express this, we modify the 3D Laplacian operator $\mathbf{Q}$ defined in Eq. (14) to obtain the 3D motion compensated Laplacian (3DMCL) defined by:

$$
\begin{aligned}
[\mathbf{Q}_{3DMCL}\,\mathbf{f}]_l(i, j) \;=\; & -6 f_l(i, j)(i, j) + f_l(i - 1, j) + f_l(i + 1, j) + f_l(i, j - 1) + f_l(i, j + 1) \\
& f_{l-1}\left(i + m^{(i,j)}_{(l-1,l)}, j + n^{(i,j)}_{(l-1,l)}\right) + f_{l+1}\left(i + m^{(i,j)}_{(l+1,l)}, j + n^{(i,j)}_{(l+1,l)}\right)
\end{aligned}
\tag{25}
$$

where $(m^{(i,j)}_{(l+k,l)}, n^{(i,j)}_{(l+k,l)})$, $k = -1, 1$, represents the DV between frames l and $(l + k)$ for pixel (i, j). It is easy to generalize this operator to capture the temporal correlation between more than three channels.

The iterative algorithm in Eq. (24) is easily implemented using the regularization operator $\mathbf{Q}_{3DMCL}$ because it is assumed symmetric, i.e., $\mathbf{Q}_{3DMCL} = \mathbf{Q}^T_{3DMCL}$. Therefore, in Eq. (24) $\mathbf{Q}_{3DMCL}{}^T\mathbf{Q}_{3DMCL}\mathbf{f}$ is computed by applying $\mathbf{Q}_{3DMCL}$ twice (for more details see [1]). Integer DVs are used in the example shown in this section. The generalization of this approach to non-integer DVs for recovery of compressed video is presented in [19].

The application of multichannel image restoration to image sequences with motion is demonstrated by the following experimental example. Ten frames (each of size 256×256) from the "Trevor White" sequence were used as test images. The results obtained by multichannel restoration are compared with the results obtained by restoring each frame separately (henceforth referred to as Model 0) using an independent-channel version of Eq. (24) in which the regularization operator is $\mathbf{Q} = diag\{Q, Q, \cdots Q\}$, where Q represents the convolution with the 2-D Laplacian kernel defined in Eq. (13).

To apply the iteration in Eq. (24) the DVF must first be estimated. Four approaches were used for this task. We refer to these approaches, combined with the iteration in Eq. (24), as Models 1-4, which are defined as follows. In Model 1 the DVF is estimated directly from the degraded images. In Model 2 the DVF is estimated from the images restored by Model 0. In Model 3 the DVF is estimated from the images restored by Model 2. In Model 4 the original image sequence is used to obtain the DVFs. Model 4 is used to test the upper bound of performance of the proposed multichannel restoration algorithm.

In Models 1-4 the DVF is computed from either the degraded, the restored, or the source image. A block search algorithm (BSA) was used to estimate the between channel DVFs. The motion vector at pixel (i, j) between frames l and k, was found by matching a 5×5 window centered at pixel (i, j) of frame l to a 5×5 window in frame k. An exhaustive search over a 31×31 area centered at pixel (i, j) of frame k was used and the matching metric was the sum of the squared errors.

Two experiments are summarized here (more are described in [1]), in which all five models were tested and compared. The variance of the noise added to the blurred data is defined using the blurred $BSNR$ metric which was defined in Eq. (22). As an objective measure of performance of the restoration algorithms the improvement signal-to-noise ratio ($ISNR$) metric was used. This metric is given by

$$ISNR \;=\; 10 \log_{10} \frac{||f_i - g_i||^2}{||f_i - \hat{f}_i||^2} dB, \tag{26}$$

where the vectors f_i, g_i and $\hat{f}_i$ are the ith channel of the original image, the degraded image and the restored image, respectively.

In both experiments the relaxation parameter α was obtained numerically using a method, based on the Rayleigh quotient, described in [1]. The value of the regularization parameter λ was chosen to be equal to $\left(10^{\frac{BSNR}{10}} \right)^{-1}$ [6]. To restore all ten frames of the image sequence, six five-channel multichannel filters were used in which a 5-channel multichannel regularization operator similar to the one in Eq. (25) was used. Except for the first two and last two frames of the sequence a five-channel non-causal filter was used to restore each frame. This filter used both the two previous and the two following frames of the frame being restored.

Ten frames (frames 41-50) of the Trevor White sequence were blurred by an 11×11 uniform blur. The point spread function of this blur is given by

$$h(i, j) = \begin{cases} \frac{1}{121} & \text{if } -5 \leq i, j \leq 5 \\ 0 & \text{otherwise.} \end{cases} \tag{27}$$

Cases (i) and (ii) corresponding, respectively, to 10 and 30 dB $BSNR$ of additive white Gaussian noise were examined. Plots of the $ISNR$ are shown in Figs. 4 and 5. In Figs. 6, 7, and 8 the 8^{th} frame of this experiment is shown for cases (i) and (ii). The original and the degraded images are shown in Fig. 6. In Figs. 7 and 8 the restored images from this experiment are shown.

Both the visual and the $PSNR$ results of this experiment demonstrate that: 1) the multi-

channel regularization greatly improves the restored images, and 2) the accuracy of the between-channel knowledge that is incorporated is crucial to the quality of the results.

5 Implicit Approach to Multichannel Image Recovery

The purpose of multichannel image recovery is to make use of the correlations between channels of the source image for purposes of noise suppression. Unfortunately, the between-channel smoothing required to exploit this information can greatly increase the computational cost. In the implicit approach, the computational burden is dramatically reduced by applying a Karhunen-Loève (KL) transformation to the observed data prior to processing. For a review of the KL transform see [10], p.163.

The computational savings result from two important functions of the KL transform: decorrelation and compression. Because the KL transform decorrelates the source channels, it eliminates the need for cross-channel smoothing. In addition, because the KL transform compresses the significant signal information, it effectively reduces the number of channels that must be processed. For example, a 50-channel source image with highly correlated channels might be described almost perfectly by only five KL channels, with the remaining 45 channels dominated by noise. In such an example, only five channels of data would need to be processed instead of the original 50.

The basic steps of the implicit approach to multichannel image recovery are as follows: 1) apply a KL transformation to the data; 2) discard the channels of the KL-transformed data that are dominated by noise; 2) recover an image channel from each of the remaining KL-domain data channels separately; and 3) apply an inverse KL transform to the recovered channels to convert them back to the original domain.

Having outlined the steps of the basic algorithm, let us now explain and justify it. The implicit approach is based on the assumption that the multichannel covariance matrix $\mathbf{C}_f$ in Eq. (6) is separable into a spatial part $C_f^{(s)}$ of dimension $K \times K$ and a temporal part $C_f^{(t)}$ of dimension $N \times N$ as follows:

$$\mathbf{C}_f = C_f^{(t)} \otimes C_f^{(s)} \tag{28}$$

where $\otimes$ denotes the Kronecker product. In the recovery algorithm this calls for separate temporal and spatial regularization operations. This separability assumption is best suited for imaging

of motion-free objects or scenes; however, it has been shown in [17] to work extremely well in reconstructing image sequences of the beating heart as we will show later.

Decorrelation of the channels of the source image is achieved by application of a KL transformation Φ, which is the transpose of the eigenvector matrix of $C_f^{(t)}$, i.e.,

$$C_f^{(t)} \Phi^t = \Phi^t D \qquad (29)$$

where $D = diag\{d_1, \ldots, d_K\}$ and d_l is the lth eigenvalue of $C_f^{(t)}$.

The limitation of the implicit approach is that it involves, in addition to the separability condition in Eq. (28), the following assumptions which may not hold in some applications.

1) *The system matrix must be of the form* $\mathbf{H} = diag\{H, H, \ldots, H\} = I \otimes H$. Each block H in the multichannel system matrix $\mathbf{H}$ denotes the system matrix describing the degradation of one image channel. This form for $\mathbf{H}$ represents the situation in which every image channel is degraded in the same way, and the channels are degraded independently of one another.

2) *The multichannel noise covariance must be of the form* $\mathbf{C}_n = diag\{C_n, C_n, \ldots, C_n\} = I \otimes C_n$. This means that every channel of the multichannel observations must be have the same noise covariance matrix and that the noise channels must be uncorrelated.

5.1   KL Transformation of the Multichannel Imaging Model

The values of pixel i in the multichannel image and pixel m in the multichannel observations form, respectively, the $K \times 1$ vectors

$$\mathbf{f}(j) = [f_1(i), f_2(i), \ldots, f_K(i)]^T \text{ and } \mathbf{g}(i) = [g_1(m), g_2(m), \ldots, g_K(m)]^T. \qquad (30)$$

In terms of these vectors, the form of $\mathbf{C}_f$ in Eq. (28) can be written as:

$$E\{\mathbf{f}(i)\mathbf{f}^T(j)\} = \left[C_f^{(s)}\right]_{ij} C_f^{(t)} \text{ for } i, j = 1, \ldots, N. \qquad (31)$$

If we define the KL-domain quantity $\tilde{\mathbf{f}}(i) = \Phi\mathbf{f}(i)$, then

$$E\{\tilde{\mathbf{f}}(i)\tilde{\mathbf{f}}^T(j)\} = \left[C_f^{(s)}\right]_{ij} D. \qquad (32)$$

Equation (32) indicates that the transformed vector $\tilde{\mathbf{f}}(i)$ does not exhibit any between-channel correlations. Thus, if recovery is performed in the KL domain, the need for between-channel

smoothing is eliminated. As applied to the multichannel vector $\mathbf{f}$, the KL transform is represented by a multichannel transformation matrix $\mathbf{A}_M$ defined as

$$\mathbf{A}_M = \Phi \otimes I_M, \tag{33}$$

where I_M denotes the $M \times M$ identity matrix. Applying $\mathbf{A}_M$ to both sides of the multichannel imaging model in Eq.(1), we obtain:

$$\mathbf{A}_M E\{\mathbf{g}\} = \mathbf{A}_M \mathbf{H} \mathbf{f}, \tag{34}$$

where $E\{\cdot\}$ denotes expectation with respect to the noise $\mathbf{n}$. Using properties of the Kronecker product, we rewrite $\mathbf{A}_M \mathbf{H}$ as follows:

$$\begin{aligned}
\mathbf{A}_M \mathbf{H} &= (\Phi \otimes I_M)(I_K \otimes H) = (\Phi I_K) \otimes (I_M H) \\
&= (I_K \Phi) \otimes (H I_N) = (I_K \otimes H)(\Phi \otimes I_N) = \mathbf{H} \mathbf{A}_N.
\end{aligned} \tag{35}$$

Interchanging the transformation matrix and the expectation operator in Eq. (34) , using the result in Eq. (35), and defining the transformed quantities $\tilde{\mathbf{g}} = \mathbf{A}_M \mathbf{g}$ and $\tilde{\mathbf{f}} = \mathbf{A}_N \mathbf{f}$, yields the following transformed imaging model:

$$E\{\tilde{\mathbf{g}}\} = \mathbf{H} \tilde{\mathbf{f}}. \tag{36}$$

Note that Eq. (36) has precisely the same form as the original linear imaging model, thus solution in the KL domain can be accomplished by existing recovery approaches.

5.2 KL Transformation of the RWLS Cost Functional

We define a more general version of the multichannel RWLS functional introduced in Eq. (12) as follows:

$$J(\mathbf{f}) = (\mathbf{g} - \mathbf{H}\mathbf{f})^T \mathbf{C}_n^{-1} (\mathbf{g} - \mathbf{H}\mathbf{f}) + \lambda \mathbf{f}^T \mathbf{C}_f^{-1} \mathbf{f}, \tag{37}$$

In this section we show that this RWLS functional is simplified greatly by the KL transformation under the conditions described previously. In a statistical interpretation of this functional, $\mathbf{C}_f$ should be chosen to be the covariance matrix of the multichannel image $\mathbf{f}$; therefore, $C_f^{(t)}$ and $C_f^{(s)}$ should be chosen to be the covariance matrices expressing the between-channel and within-channel covariances, respectively. As described earlier, we choose $C_f^{(s)} = (\lambda Q^T Q)^{-1}$, where Q is a matrix representing a discrete approximation of the two-dimensional (2D) Laplacian operator.

To transform the multichannel RWLS cost functional, we begin by writing the quantities of interest in terms of their KL-domain counterparts (identified with a tilde) as follows:

$$\mathbf{g} = \mathbf{A}_M^T \tilde{\mathbf{g}} \text{ and } \mathbf{f} = \mathbf{A}_N^T \tilde{\mathbf{f}}. \tag{38}$$

Note that $\mathbf{A}_N$ and $\mathbf{A}_M$ are orthogonal matrices, so $\mathbf{A}_N^T = \mathbf{A}_N^{-1}$ and $\mathbf{A}_M^T = \mathbf{A}_M^{-1}$. Substituting for these quantities in Eq. (37), we obtain

$$\begin{aligned} J(\tilde{\mathbf{f}}) &= \left(\mathbf{A}_M^T \tilde{\mathbf{g}} - \mathbf{H}\mathbf{A}_N^T \tilde{\mathbf{f}}\right)^T \mathbf{C}_n^{-1} \left(\mathbf{A}_M^T \tilde{\mathbf{g}} - \mathbf{H}\mathbf{A}_N^T \tilde{\mathbf{f}}\right) \\ &\quad + \lambda\, \mathbf{A}_N^T \tilde{\mathbf{f}}^T \left[C_f^{(t)} \otimes C_f^{(s)}\right]^{-1} \mathbf{A}_N^T \tilde{\mathbf{f}}. \end{aligned} \tag{39}$$

In a manner similar to that used to derive Eq. (35), it can be shown that

$$\mathbf{H}\mathbf{A}_N^T \tilde{\mathbf{f}} = \mathbf{A}_M^T \mathbf{H}\tilde{\mathbf{f}}. \tag{40}$$

Using Eq. (40), distributing the transpose operations, and factoring out the transformation matrices, we obtain

$$\begin{aligned} J(\tilde{\mathbf{f}}) &= \left(\tilde{\mathbf{g}} - \mathbf{H}\tilde{\mathbf{f}}\right)^T \mathbf{A}_M \mathbf{C}_n^{-1} \mathbf{A}_M^T \left(\tilde{\mathbf{g}} - \mathbf{H}\tilde{\mathbf{f}}\right) \\ &\quad + \lambda\, \tilde{\mathbf{f}}^T \left\{\mathbf{A}_N \left[C_f^{(t)} \otimes C_f^{(s)}\right]^{-1} \mathbf{A}_N^T\right\} \tilde{\mathbf{f}}. \end{aligned} \tag{41}$$

Using Eq. (33), we rewrite the term in curly braces as follows:

$$\mathbf{A}_N \left[C_f^{(t)} \otimes C_f^{(s)}\right]^{-1} \mathbf{A}_N^T = D^{-1} \otimes \left[C_f^{(s)}\right]^{-1}, \tag{42}$$

where D is the eigenvalue matrix defined in Eq. (29). It is easy to show that

$$\mathbf{A}_M \mathbf{C}_n^{-1} \mathbf{A}_M^T = \mathbf{C}_{\tilde{n}}^{-1}, \tag{43}$$

where $\mathbf{C}_{\tilde{n}}$ is the covariance matrix of the observations in the KL domain. Using Eqs. (42) and (43), Eq. (41) becomes

$$J(\tilde{\mathbf{f}}) = \left(\tilde{\mathbf{g}} - \mathbf{H}\tilde{\mathbf{f}}\right)^T \mathbf{C}_{\tilde{n}}^{-1} \left(\tilde{\mathbf{g}} - \mathbf{H}\tilde{\mathbf{f}}\right) + \lambda\, \tilde{\mathbf{f}}^T \left\{D^{-1} \otimes \left[C_f^{(s)}\right]^{-1}\right\} \tilde{\mathbf{f}}. \tag{44}$$

Using the assumption $\mathbf{C}_{\tilde{n}}^{-1} = I \otimes C_n^{(s)^{-1}}$, since $\mathbf{H}$ and $\mathbf{C}_{\tilde{n}}^{-1}$ are block diagonal, and D and I are diagonal, Eq. (44) reduces to

$$J(\tilde{\mathbf{f}}) = \sum_{l=1}^{K} J_l(\tilde{f}_l), \tag{45}$$

in which

$$J_l(\tilde{f}_l) = \left(\tilde{g}_l - H\tilde{f}_l\right)^T C_{\tilde{n}_l}^{-1} \left(\tilde{g}_l - H\tilde{f}_l\right) + \left(\frac{\lambda}{d_l}\right) \tilde{f}_l^T \left[C_f^{(s)}\right]^{-1} \tilde{f}_l. \tag{46}$$

where $\tilde{f}_l$, and $\tilde{g}_l$ are the l^{th} KL components of $\mathbf{f}$, and $\mathbf{g}$, respectively and d_l is the eigenvalue associated with the l^{th} KL basis vector.

5.3 Space-invariant image restoration by the implicit approach

Image restoration (deblurring) problems can be solved especially easily when the degradation operator H and the covariance matrices $C_f^{(s)}$ and C_{n_l} are circulant. As in the general case, application of the implicit approach begins with computation of the covariance matrix $C_f^{(t)}$ and its eigenvectors in Φ, which are used to transform the observed multichannel image to the KL domain. Then the conventional Wiener filter [10] can be implemented in the DFT domain in closed form, and applied one-by-one to the significant KL-domain channels to restore them. Finally, the multichannel image is obtained by inverting the KL transform.

5.4 Space-variant image recovery by the implicit approach

When the degradation is not shift-invariant and/or the statistics are not stationary, the recovered KL-domain channels must be computed iteratively. The RWLS functional $J(\mathbf{f})$ can be minimized by minimizing $J_l(\tilde{f}_l)$ separately in Eq. (46). Since $J_l(\tilde{f}_l)$ is quadratic with respect to $\tilde{f}_l$ a number of iterative minimization methods, including the conjugate gradient algorithm [3], can be used to find $\tilde{f}_l$. Theoretically, the conjugate gradient method is guaranteed to converge in N steps (the dimension of each image channel), but a much smaller number of iterations is sufficient for good results in practice.

5.5 Multichannel Reconstruction of Medical Image Sequences

In this section we describe an application of the implicit multichannel recovery approach to an important problem in medical imaging, namely the reconstruction of time-sequences of images. We focus specifically on two emission tomography methods: positron emission tomography (PET) and single-photon emission computed tomography (SPECT) [2].

For purposes of computation, the imaging model for PET and SPECT sequences can be approximated by the set of matrix equations

$$E\{g_k\} = Hf_k, \ k = 1, 2, \cdots K. \tag{47}$$

This corresponds exactly to the previously discussed multichannel linear imaging model for the special case in which $\mathbf{H} = diag\{H, H, \ldots, H\}$. In this application, H represents a tomographic projection operator which is not shift-invariant. In an idealized model, the projection operator is

the discrete Radon transform, but more-realistic models include blur caused by various physical factors in the imaging process.

In dynamic PET, one obtains a time sequence of data g_k, $k = 1, \ldots, K$, from which an image sequence f_k, $k = 1, \ldots, K$, is to be reconstructed. Usually the reconstruction is performed image by image, but recent research [20, 21, 17, 15, 12] has shown that it is preferable to reconstruct all of the images collectively as a multichannel image $\mathbf{f}$ from all of the data in the multichannel observation vector $\mathbf{g}$. The following example from PET brain imaging illustrates this principle. The images shown depict slices of the brain of a monkey; the bright areas indicate tissues rich in dopamine receptors, the part of the brain's chemical communication system

In the implicit approach, one begins by applying a KL transformation along the time axis of the data (across the channels of $\mathbf{g}$). Fig. 9 shows example frames from a time sequence of 44 frames of tomographic projection data; Fig. 10 shows the first six frames following the KL transformation. The KL transform eliminates redundancy in the observations and compresses the useful information into the first three frames. The remaining frames, which all look similar, are dominated by noise and can be discarded. The importance of the first three frames is depicted quantitatively by the eigenvalue spectrum shown in Fig. 11. Fig. 12 shows the result of reconstructing images from the first three KL-domain observations. The inverse KL transform is applied to these three KL-domain images to obtain a sequence of images. Examples of these results are shown in Fig. 13 where they are compared with the results obtained by more-conventional approaches. Note that, not only are the images obtained by the implicit multichannel approach superior to the others, but they were obtained in less time because only three KL frames required reconstruction instead of 44 time-domain image frames. Quantitative performance evaluations of the implicit approach for image reconstruction can be found in [12, 21].

Figure 14 shows another application of the implicit approach to cardiac SPECT imaging. Two example frames are shown, each reconstructed by both a single-channel approach and by the implicit approach. Image features that are normally obscured by noise are clearly visible when reconstructed by the implicit approach. Because of the separability assumption in Eq. (28) one might expect the implicit approach to perform poorly when there is motion; however, these images of the beating heart show that the KL decomposition can capture motion information in some cases.

6   Acknowledgments

References

[1] Mun Gi Choi, N. P. Galatsanos and A. K. Katsaggelos, "Multichannel regularized iterative restoration of motion compensated image sequences," *Journal of Visual Communication and Image Representation,* Vol. 7, No. 3, pp. 244-258, Sept. 1996.

[2] Z.H. Cho, J.P. Jones, and M. Singh, *Foundations of Medical Imaging.* New York: Wiley, 1993.

[3] E. Z. P. Chong and S. H. Zak, *An Introduction to Optimization.* New York: Wiley, 1996.

[4] G. Demoment, "Image reconstruction and restoration: Overview of common estimation structures and problems," *IEEE Trans. Acoustics, Speech and Signal Processing*, Vol. 37, No. 12, pp. 2024-2036, March 1989.

[5] N. P. Galatsanos and R. T. Chin, "Digital restoration of multichannel images," *IEEE Trans. Acoustics, Speech and Signal Processing*, Vol. 37, No. 3, March 1989.

[6] N. P. Galatsanos, A. K. Katsaggelos, R. T. Chin and A. D. Hillery, "Least squares restoration of multichannel images", *IEEE Trans. Signal Processing* , Vol. 39, No. 10, pp. 2222-2236, Oct. 1991.

[7] N. P. Galatsanos, A. K. Katsaggelos, R. T. Chin and A. D. Hillery, "Least Squares Restoration of Multichannel Images", *IEEE Trans. Signal Processing* , Vol. 39, No. 10, Oct. 1991, pp. 2222-2236.

[8] N. P. Galatsanos and A. K. Katsaggelos, "Methods for choosing the regularization parameter and estimating the noise bariance in image restoration and their relation," *IEEE Trans. on Image Processing,* Vol. 1, No. 3, pp. 322-336, July 1992.

[9] S. Geman and D. Geman, "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images",*IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, no. 6, pp. 721-742, 1984.

[10] A. K. Jain, *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.

[11] F. C. Jeng and J. W. Woods, "Simulated Annealing in Compound Gaussian Random Fields", *IEEE Transactions on Information Theory*, vol. 36, no. 1, pp. 94-107, 1990.

[12] C.-M. Kao, J. T. Yap, J. Mukherjee, and M. N. Wernick,"Image reconstruction for dynamic PET based on low-order approximation and restoration of the sinogram," *IEEE Trans. Med. Imaging*, vol. 16, pp. 738-749, 1997.

[13] A. K. Katsaggelos, K. T. Lay, and N. P. Galatsanos, "A general framework for frequency domain multichannel signal processing," *IEEE Trans. Image Proc.*, vol. 2, no. 3, pp. 417-420, July 1993.

[14] S. M. Kay *Fundamentals of Statistical Signal Processing*. Prentice Hall, 1993.

[15] D. S. Lalush and B. M. W. Tsui, "A priori motion models for four-dimensional reconstruction in gated cardiac SPECT," *1996 Conf. Record of the Nucl. Sci. Symp. Med. Imag. Conf.*, 1996.

[16] R. Molina, J. Mateos, A.K. Katsaggelos and M. Vega, "Bayesian Multichannel Image Restoration Using Compound Gauss-Markov Random Fields", *IEEE Trans. Image Proc.*, vol. 12, no. 12, pp. 1642-1654, 2003.

[17] V. M. Narayanan, M. A. King, E. Soares, C. Byrne, H. Pretorius, and M. N. Wernick, "Application of the Karhunen-Loeve transform to 4D reconstruction of gated cardiac SPECT images," *1997 Conf. Record of the Nucl. Sci. Symp. Med. Imag. Conf.*, 1997.

[18] R. W. Schafer, R. M. Mersereau and M. A. Richards, "Constrained Iterative Restoration Algorithm," *Proc. of IEEE*, vol. 69, pp. 432-450, April 1981.

[19] Y. Yang M. C. Choi, and N. P. Galatsanos, "Image Recovery from Compressed Video Using Multichannel Regularization," in *Signal Recovery Techniques for Image and Video Compression and Transmission*, A. K. Katsaggelos and N. P. Galatsanos, eds., pp. 69-98. Kluwer, 1998.

[20] M. N. Wernick, G. Wang, C.-M. Kao, J. T. Yap, J. Mukherjee, M. Cooper, and C.-T. Chen, "An image reconstruction method for dynamic PET," *1995 Conf. Record of the Nucl. Sci. Symp. Med. Imag. Conf.*, 1995.

[21] M. N. Wernick, E. J. Infusino, M. Milošević, "Fast spatio-temporal image reconstruction for dynamic PET," *IEEE Trans. Med. Imaging*, 1999, to appear.
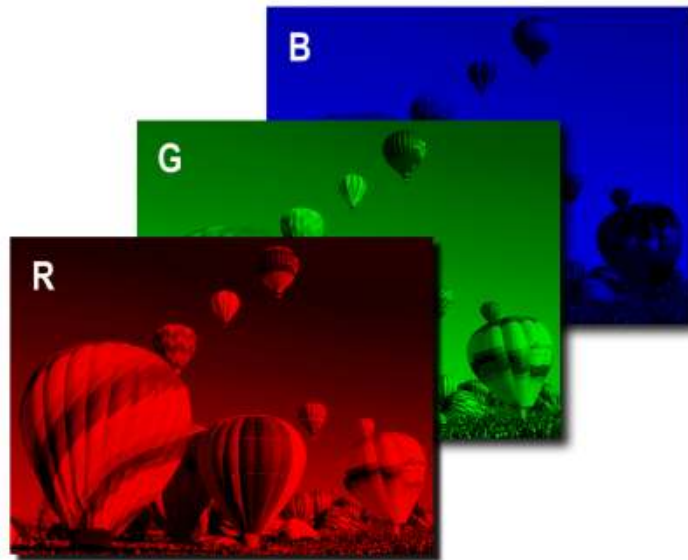
Figure 1: Example of a multichannel image. A color image consists of three color components (channels) which are highly correlated with one another. Similarly, a video image sequence consists of a collection of closely related images.

Figure 2: Example of a multichannel LMMSE restoration, original (upper left), degraded (upper right), restored single-channel power-spectrum from original (middle left), restored single-channel power-spectrum from degraded original (middle right), restored multi-channel power-spectrum from original (lower left), restored multi-channel power-spectrum from degraded (lower right).

a   b
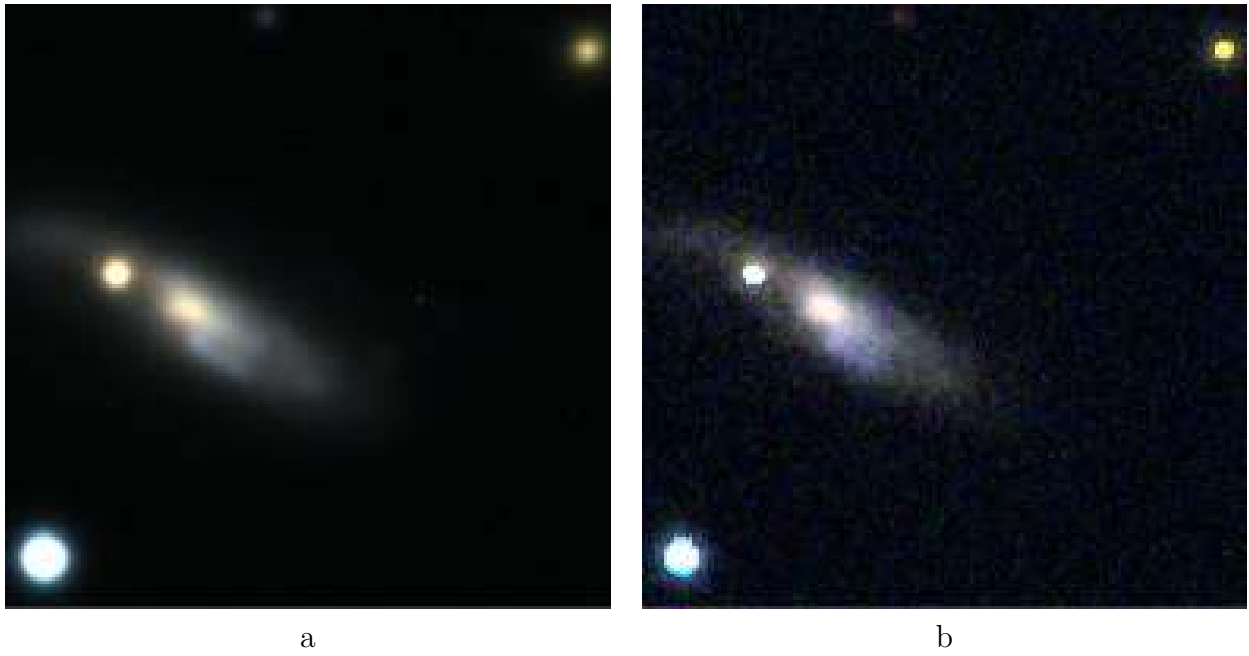
Figure 3: (a) Observed astronomical multichannel image. (b) Restoration using the regularization in Eq. (16).
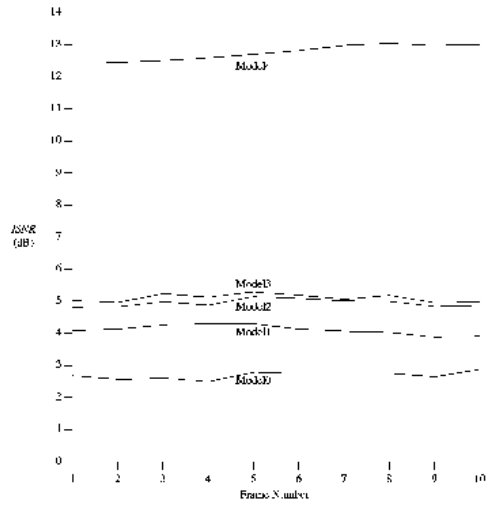
Figure 4: $ISNR$ plots for Experiment I, case (i): $BSNR = 10 \ dB$, 11×11 blur, $\alpha = 0.1$ and $\lambda = 0.1$.
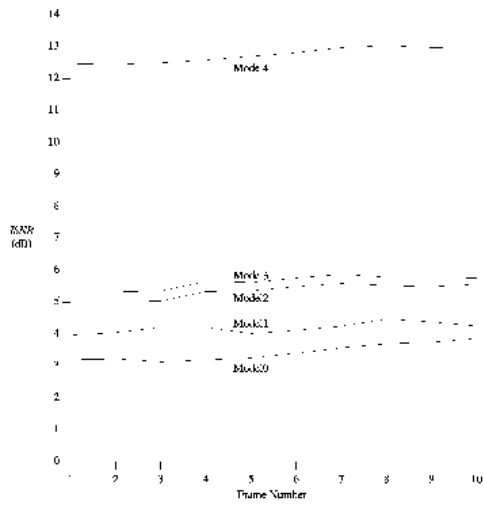
Figure 5: $ISNR$ plots for Experiment I, case (ii): $BSNR = 30$ dB, 11×11 blur, $\alpha = 2.0$ and $\lambda = 0.001$.

Figure 6: Original Image twy048 (top). Experiment I case (i) (bottom left): Degraded image, with 11×11 blur and 10 dB of $BSNR$ additive noise. Experiment I case (ii): Degraded image, with 11×11 blur and 30 dB of $BSNR$ additive noise. (bottom right)
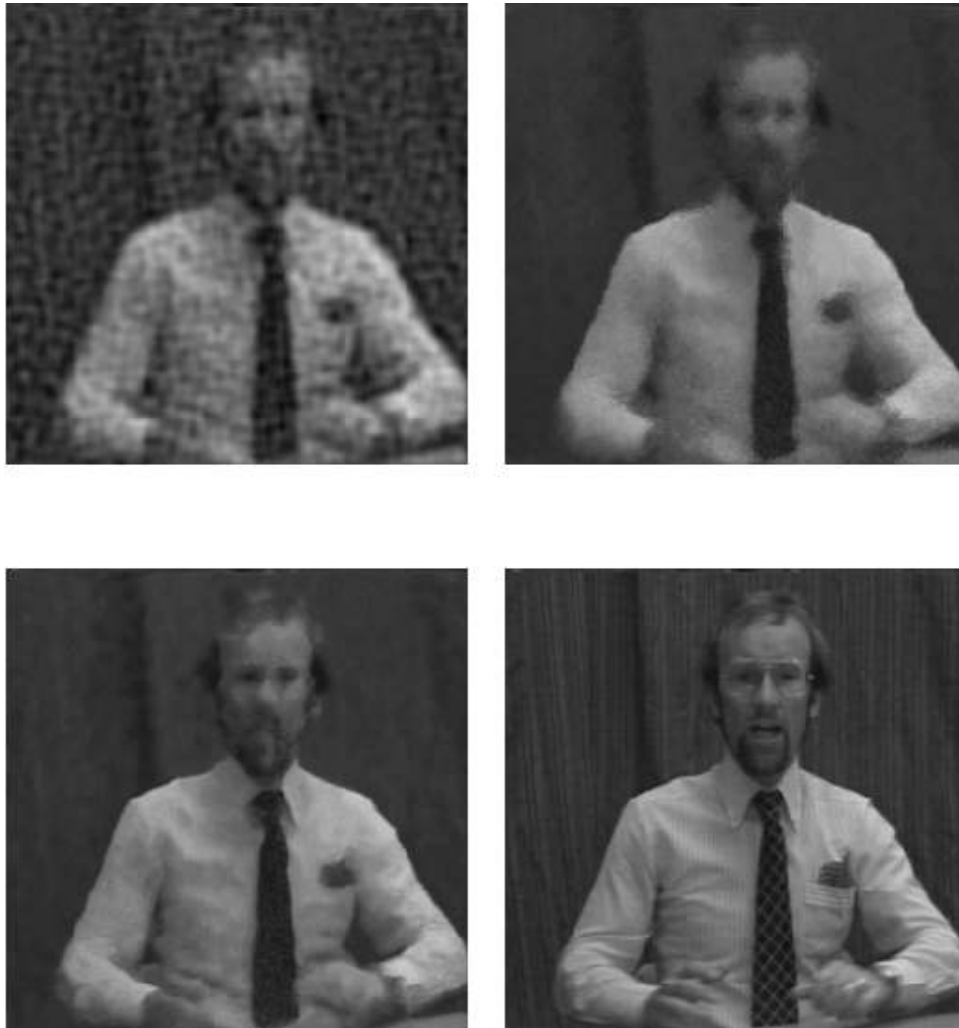
Figure 7: Experiment I case (i) restored images using Model 0 (upper left), Model 1 (upper right), Model 3 (lower left), and Model 4 (lower right).

Figure 8: Experiment *I* case (*ii*) restored images using Model 0 (upper left), Model 1 (upper right), Model 3 (lower left), and Model 4 (lower right).

Figure 9: Example frames (numbers 1, 2, 3, 10, 20, 40) from a sequence of 44 frames of dynamic positron emission tomography (PET) data. In dynamic PET the object typically is stationary, but is changing in time. Data courtesy of Jogeshwar Mukherjee.

Figure 10: First six Karhunen-Loève components of the dynamic PET data in Fig. 9. The remaining 38 components look similar to the sixth and are dominated by noise. Only the first three contain significant signal information.
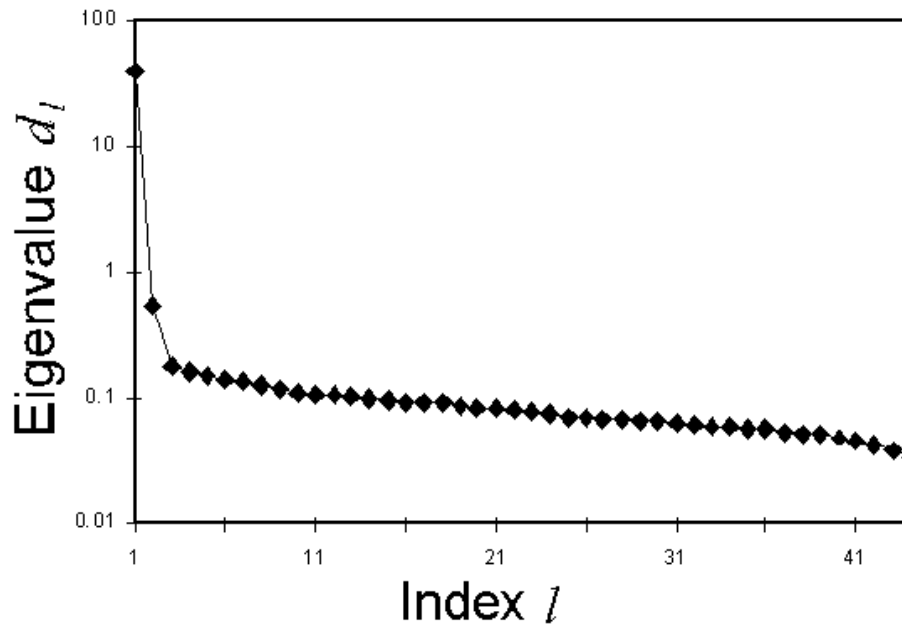

Figure 11: Spectrum of eigenvalues showing the dominance of the first two KL components. Usually the next component contains significant signal content as well.

Figure 12: Images reconstructed from the first six KL components of the data (see Fig. 10). All but the first three are dominated by noise and can be omitted from the computations.
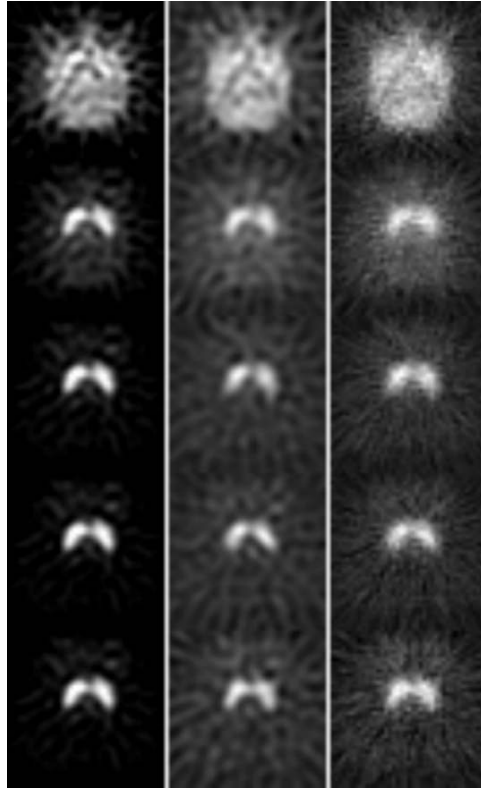
Figure 13: Example frames from the sequence of dynamic PET images reconstructed by the implicit approach (left column), by separate single-channel reconstruction by PWLS (center column) and filtered backprojection (FBP, right column). Because of the high noise level, single-channel PWLS fails to produce significantly better results than FBP, but multichannel regularization, provided by the implicit approach yields more-accurate images.
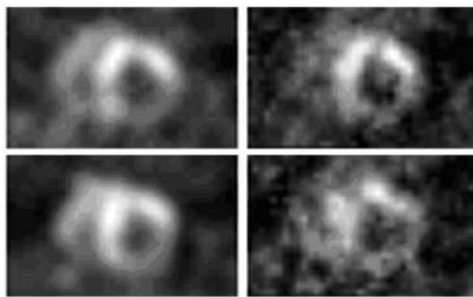

Figure 14: Example frames from a sequence of gated SPECT images of the heart. Images reconstructed by the implicit multichannel approach (left column) are less noisy than those obtained by single-channel reconstruction (right column). These images were obtained without accounting for blur in the system matrix $\mathbf{H}$ so no deblurring effect is apparent. Image results courtesy of V. Manoj Naryanan and Michael A. King.