

Restoration of Noisy, Blurred, Undersampled Image Sequences Using a Parametric Motion Model

Fabien Dekeyser

Patrick Pérez

Patrick Boutheymy

IRISA/INRIA

Campus de Beaulieu

35042 Rennes Cedex, France

contact author : fabien.dekeyser@irisa.fr

Abstract

This paper introduces the use of a 2D parametric motion model to reduce noise in image sequences. We estimate with a robust method an affine motion model, accounting for the dominant image motion, and then cancel it before applying an adaptive spatio-temporal filter. We have compared the performance of several filtering schemes and evaluated the influence of the motion compensation step on this performance, including the comparison to the use of a dense motion field.

We then apply on the denoised image sequence a low cost scheme to reconstruct high resolution images. This scheme is derived from an iterative backprojection algorithm. We again exploit the parametric motion model to keep a good trade-off between accuracy and computation time.

Keywords: *image sequence restoration, 2D parametric motion model, super-resolution*

Résumé

Nous introduisons l'utilisation d'un modèle paramétrique de mouvement 2D pour la restauration de séquences d'images. A l'aide d'un estimateur robuste, nous estimons les paramètres d'un modèle de mouvement affine représentant le mouvement dominant dans l'image. Nous le compensons avant d'appliquer un filtre spatio-temporel adaptatif. Nous avons comparé les performances de plusieurs schémas de filtrage et évalué l'influence de la compensation de mouvement sur leurs performances. En particulier, nous avons comparé l'efficacité d'un modèle paramétrique avec celle d'un champ de déplacement dense.

Nous appliquons ensuite sur la séquence d'images débruitées un schéma de reconstruction d'images haute résolution à faible coût de calcul. Ce schéma est dérivé d'un algorithme de rétro-projection itérative. Le modèle paramétrique de mouvement est à nouveau exploité afin de garder un bon compromis entre la précision des résultats et le temps de calcul.

Mots-clé : *restauration de séquences d'images, modèle paramétrique de mouvement 2D, super-résolution*

1 Introduction

Image sequences are degraded by the resolution loss due to the downsampling (not meeting the Nyquist criterion) of the images and to the integration over sensor area. However, the knowledge of subpixel motion can allow one to reconstruct high resolution images from low resolution image sequences. An overview of proposed methods for reconstructing a high resolution image from a low resolution image sequence can be found in [6]. Some methods for reconstructing high resolution image sequences have appeared more recently [3, 7]. One of the computationally simplest algorithm was proposed by Irani and Peleg and is known as iterative backprojection [12]. However, this method does not explicitly account for the presence of noise in image sequences. Whatever recording means are, digital pictures remain noisy. They can be considered to be corrupted by a Gaussian white noise which is an appropriate model for quantization and electronic noises.

To reduce noise in image sequences, several temporal and spatio-temporal filtering schemes have been proposed. A survey can be found in [4]. More recently, methods appeared using wavelets [24] or PDEs [15]. Whichever scheme is used, it is always preferable to first compensate for apparent motion in image sequences, in order to avoid the blurring of moving objects by temporal filtering. It is also preferable to use adaptive techniques that will inhibit filtering where motion has not been compensated for accurately enough.

Most motion estimators used in these noise filtering schemes use either block matching or dense optic flow but not parametric motion models. Nevertheless, the latter are good approximations of apparent motion for a considerably lower

computation time [20]. Parametric motion models are thus attractive for spatio-temporal filtering of image sequences. The estimation of one single parametric motion model enables to account only for dominant motion, but residual motions can be handled efficiently by adaptive filtering.

Blurring is another type of degradation frequently encountered in image sequences. However, setting apart motion blur, it can be treated as a purely spatial problem.

The restoration scheme we propose in this paper consists first in reducing noise in an image sequence via parametric motion compensation and spatio-temporal Wiener filtering. We then use this denoised image sequence to construct a better resolved image still exploiting the estimated motion model. The deblurring is included in the super-resolution process.

In Section 2, we present the motion estimation process. Section 3 is devoted to the spatio-temporal Wiener filter, while Section 4 deals with the super-resolution algorithm. Section 5 contains some experimental results and concluding remarks.

2 Motion estimation using a 2D parametric model

The displacement between two images is modeled as a polynomial of the image point coordinates:

$$\vec{w}_\theta(s) = P(s)\theta$$

where parameter vector θ is made up with the coefficients of the concerned polynomial, $s = (x, y)^T$ denotes pixel location, $\vec{w}_\theta(s)$ stands for the displacement vector at point s supplied by the parametric motion model, and P is a matrix whose elements are monomials of x and y depending on the chosen motion model. The most commonly used model is the 6-parameter affine motion model. Its estimation is stated as the minimization of the following robust function as introduced and described in [20]:

$$H(\theta) \triangleq \sum_{s \in S} \rho [g_{n+1}(s + P(s)\theta) - g_n(s)] \quad (1)$$

where S is the pixel grid, $g_n(s)$ and $g_{n+1}(s)$ stand for the intensity at pixel location s in images n and $n + 1$ respectively, and ρ is a non-quadratic robust penalty function.

This minimization is achieved within an incremental Gauss-Newton-type multiresolution procedure. At each step, a linearization of the argument of ρ is performed around the current estimate $\hat{\theta}$, yielding the following function to be minimized with respect to motion parameter increment $\Delta\theta$:

$$H'(\Delta\theta; \hat{\theta}) = \sum_{s \in S} \rho \left[g_{n+1}(s + P(s)\hat{\theta}) - g_n(s) + \nabla g_{n+1}(s + P(s)\hat{\theta})^T P(s)\Delta\theta \right] \quad (2)$$

where ∇g denotes the spatial intensity gradient.

The use of a robust estimator allows us to correctly estimate the dominant motion which is often due to camera movement, even in presence of secondary motions (due to moving objects in the scene). The multiresolution scheme solves the problem of large motion magnitude. At each instant of the sequence, a low-pass pyramid of images is built from original images g_n . At each pyramid level, parameter estimate $\hat{\theta}$ is provided by the projection of the total parameter estimate obtained at previous level, and a new increment $\Delta\theta$ is iteratively obtained using an iterated reweighted least square technique. The final estimate obtained at finest resolution is denoted as $\hat{\theta}_n^{n+1}$.

From the estimated motion model, we can compute a warped image from time $n + 1$ to n denoted $\hat{g}_n^{n+1}(s) \triangleq g_{n+1}(s + P(s)\hat{\theta}_n^{n+1})$ using bilinear interpolation, since $s + P(s)\hat{\theta}_n^{n+1}$ is likely not to have integer coordinates. Inverting the model, we can also compute $\hat{g}_n^{n-1}(s) \triangleq g_{n-1}(s + P(s)\hat{\theta}_n^{n-1})$ where $\hat{\theta}_n^{n-1}$ stands for the ‘‘inverse’’ of $\hat{\theta}_n^n$.

To evaluate the performance of the use of this parametric motion model compared to the one of an estimated dense motion field, we also consider in the following the dense optic flow method described in [19].

3 Spatio-temporal filtering

Among the different noise reduction methods, spatio-temporal Wiener filtering provides a good trade-off between efficiency and computation time. A first attempt based on the use of local statistics was proposed in [23]. In [21] a Wiener filter explicitly computing the correlation between frames was presented. Kokaram proposed a simplification using a 3D Discrete Fourier Transform (DFT) [14]. Finally, in [2], Boo and Bose use an orthogonal transform and can then reduce the 3D problem in a set of 2D problems. In the following, we also evaluate a method based on an adaptive weighted average [22]. This method, contrary to the others we have considered is not based on Wiener filtering.

3.1 Problem modeling

The blurring is modeled by a shift-invariant linear operator and the noise is supposed to be additive white Gaussian. For each degraded image, we can write :

$$g_n(i, j) = (h_n * f_n)(i, j) + \eta_n(i, j) \quad (3)$$

where :

- $g_n(i, j)$ stands for the intensity of the pixel of spatial coordinates (i, j) of the n^{th} frame of the observed sequence;
- $f_n(i, j)$ stands for the intensity of the pixel of spatial coordinates (i, j) of the n^{th} frame of the original sequence;
- $\eta_n(i, j)$ stands for the additive noise for the n^{th} frame;
- $h_n(i, j)$ stands for the Point Spread Function (PSF) for the n^{th} frame;
- $*$ stands for convolution product.

This linear system can be rewritten using a matrix formulation:

$$g_n = H_n f_n + \eta_n \quad (4)$$

where g_n, f_n and η_n are lexicographically ordered vectors [1]. Let us note that in our model, the PSF could depend on frame. In practice, we assume it is constant in the whole sequence.

For a N -frame sequence, we can write:

$$g = \begin{bmatrix} g_1 \\ \vdots \\ g_N \end{bmatrix}, f = \begin{bmatrix} f_1 \\ \vdots \\ f_N \end{bmatrix}, \eta = \begin{bmatrix} \eta_1 \\ \vdots \\ \eta_N \end{bmatrix}, \text{ and } H = \begin{bmatrix} H_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & H_N \end{bmatrix} \quad (5)$$

The model can then be rewritten with only one equation :

$$g = H f + \eta \quad (6)$$

3.2 Wiener filtering

The Wiener filter is a linear estimator that minimizes the mean square error between the filtered and original sequence, assuming knowledge of the covariance matrices of the noise and the original image sequence. The estimate \hat{f} is given by :

$$\hat{f} = R_f H^T (H R_f H^T + R_\eta)^{-1} g \quad (7)$$

where R_f et R_η are respectively the image sequence covariance matrix and the noise covariance matrix. We work on signals whose mean has been subtracted. Covariance is then equivalent to correlation. For an image sequence, we compute the mean of each image and not of the group of frames. Filtering will thus be robust to global intensity variations.

If we denote $b = H f$ the vector representing a non-noisy blurred sequence and R_b its autocorrelation matrix, we have :

$$\hat{f} = H^{-1} R_b (R_b + R_\eta)^{-1} g \quad (8)$$

assuming H invertible.

This formulation emphasizes that the restoration process can be viewed as a noise reduction step, whose gain is $R_b (R_b + R_\eta)^{-1}$, followed by a deblurring step (H^{-1}). The inversion of H being usually impossible, we shall use for the deblurring step a regularization scheme such as truncated singular value decomposition or regularization via truncated iterations [16].

As already noted, the Wiener filter requires the knowledge of the covariance matrix of the original image or of the non-noisy blurred image sequence. This covariance matrix could be considered as known *a priori* and valid for a wide range of images [1]. But, in the case of adaptive filtering, we can use ergodic property to compute such an information from the sequence itself. As a matter of fact, we can formulate the expression of the Wiener filter in terms of the covariance matrix of the observed image sequence. Assuming noise and signal decorrelated, we can rewrite equation (8) as follows:

$$\hat{f} = H^{-1} (R_g - R_\eta) R_g^{-1} g \quad (9)$$

where R_g and R_η are covariance matrices of the observed sequence and noise respectively. In practice, matrix $R_g - R_\eta$ could have negative values. They are set to zero since image intensity is always positive.

Noise variance can be estimated by the intensity variance in uniform areas of an image [1]. Anyway, a coarse approximation is sufficient because it is not a very sensitive parameter for the quality of the restoration results.

We have still to estimate and invert covariance matrix R_g . This matrix has the following form:

$$R_g = E[gg^T] = \begin{bmatrix} g_1 g_1^T & g_1 g_2^T & \cdots & g_1 g_N^T \\ g_2 g_1^T & g_2 g_2^T & & \vdots \\ g_3 g_1^T & g_3 g_2^T & & \vdots \\ \vdots & & \ddots & \vdots \\ g_N g_1^T & \cdots & \cdots & g_N g_N^T \end{bmatrix} \quad (10)$$

For images of size $M \times M$, this matrix has size $NM^2 \times NM^2$. The estimation and the inversion of this matrix are computationally heavy.

Nevertheless, if the frames are assumed to be stationary, the covariance matrices of each frame $E[g_j g_j^T]$ are block-Toeplitz. They can be approximated by circulant matrices which can be diagonalized by 2D DFT [9]. In this case, the matrix inversion can be determined by the inversion of M^2 matrices of size $N \times N$ [21]. However, with this method, we have to calculate the spectral power density of each frame as well as the inter-frame power spectral density. The results are sensitive to the estimation of these quantities.

Moreover, the structure of matrix R_g shows that the use of a purely spatial filter on each frame is largely suboptimal. An independent restoration of each frame implicitly assumes that non-diagonal blocks are null; in other words, that the frames of the sequence are decorrelated. This hypothesis is obviously wrong.

Therefore, we have considered three other methods that we are now describing. All these methods are adaptive, involving local computation on limited spatial supports and short temporal windows. We thus account for spatial non-stationarities (edges) and temporal ones due to moving objects.

The first approach is a generalization of Lee filter to three dimensions [17]. It consists in using local statistics around each pixel. The mean and variance are computed using a cube of size $(2p+1) \times (2q+1) \times (2r+1)$. The local mean is defined as :

$$M_{i,j} = \frac{1}{(2p+1)(2q+1)(2r+1)} \sum_{k=i-p}^{i+p} \sum_{l=j-q}^{j+q} \sum_{m=n-r}^{n+r} g_m(k,l) \quad (11)$$

The variance is similarly defined as :

$$\nu_{i,j} = \frac{1}{(2p+1)(2q+1)(2r+1)} \sum_{k=i-p}^{i+p} \sum_{l=j-q}^{j+q} \sum_{m=n-r}^{n+r} (g_m(k,l) - M_{i,j})^2 \quad (12)$$

And for each pixel, the filter gain is given by :

$$w_{i,j} = \frac{\nu_{i,j} - \sigma_\eta^2}{\nu_{i,j}} \quad (13)$$

where σ_η^2 is the noise variance.

The second approach computes the filter in the Fourier domain [14]. If we assume R_g to be Toeplitz, that is if:

$$E[g_j g_k] = E[g_{j+1} g_{k+1}] \quad (14)$$

the matrix can be diagonalized by a 3D DFT technique. Then, the inversion simply consists in computing NM^2 scalar divisions in the frequency domain. The exploitation of Hermitian symmetries allows us to divide this number by two.

The variance of the image sequence is estimated from the power spectral density. The filter gain is given by :

$$W(\omega_1, \omega_2, \omega_3) = \frac{G^2(\omega_1, \omega_2, \omega_3) - \sigma_\eta^2}{G^2(\omega_1, \omega_2, \omega_3)} \quad (15)$$

where $G(\omega_1, \omega_2, \omega_3)$ is the 3D DFT of the group of observed frames.

The estimate is given by :

$$\hat{f}_n(i, j) = \mathcal{F}^{-1}(W(\omega_1, \omega_2, \omega_3)G(\omega_1, \omega_2, \omega_3)) \quad (16)$$

where \mathcal{F}^{-1} denotes the inverse 3D DFT.

The power spectral density is an appropriate estimate of the variance only if the signal is stationary. This is not true on the whole image. However, this can be verified on small blocks. Therefore, we process the signal on small overlapping blocks. The blocks overlap from half their horizontal and vertical dimensions. A Hamming window is then applied so that the final weight at each pixel is unitary. Furthermore, the Hamming window permits a better estimation of the power spectral density. This filter can be viewed as a 3D extension of the Lim filter [18].

Analyzing the formula of the gain for the two described methods, we can point out that they are robust to motion compensation errors. If motion is not accurately compensated, the variance of the image sequence becomes greater than the variance of the noise and the gain gets close to one.

The third approach was first used by Boo and Bose [2]. It assumes that spatial and temporal components are separable, which is a reasonable hypothesis [5]. From this hypothesis, Boo and Bose demonstrated the following theorem : *For a temporally stationary group of frames which is corrupted by interframe uncorrelated noise, there exists an orthogonal transform such that the 3D Wiener filter can be decorrelated into a set of 2D Wiener filters. This orthogonal transform is the Karhunen-Loeve transform.*

In practice, we use the Discrete Cosine Transform (DCT) as an approximation to the Karhunen-Loeve transform. We then apply a 2D Lee filter on each transformed image. This approach is a generalization to image sequences of the Hunt and Kübler method for multispectral images [11].

The 3D DCT supplies a transformed image which is the mean of the group of frames and other ones which are difference images. The latter are uniform except in some locations where the motion has not been compensated accurately enough. Temporal nonstationarities thus become spatial nonstationarities of the transformed images and can be handled efficiently by the adaptive 2D filter.

It appears that there exists a link between the last two approaches. As a matter of fact, the 3D DFT can be considered as an orthogonal transform approaching the Karhunen-Loeve transform.

3.3 Adaptive weighted averaging

For comparison purpose, we also used an Adaptive Weighted Averaging (AWA) method, which was first described in [22]. The intensity estimate at image location (i, j) for the n^{th} image is given by :

$$\hat{f}_n(i, j) = \sum_{(p, q, l) \in \mathcal{S}} w(p, q, l) g_{n-l}(i-p, j-q) \quad (17)$$

where \mathcal{S} is the filter support and where

$$w(p, q, l) = \frac{K}{1 + \alpha \max\left(\epsilon^2, (g_n(i, j) - g_{n-l}(i-p, j-q))^2\right)} \quad (18)$$

α and ϵ are parameters of the filter. K is a normalizing constant.

The contribution of a pixel in the average is reduced as a function of the mismatch between its value and the one of the considered pixel (i, j) . This avoids the blurring of edges and of moving areas. The α parameter tunes this decrease. The ϵ^2 parameter accounts for the noise: its value is typically set to twice the noise variance.

4 Super-resolution

Once the sequence has been filtered from noise, we can apply the super-resolution process.

We now assume that the sequence originates from a single high resolution image F . Each observed image g_n is viewed as the result of a geometrical warping W_n (related to the image motion) of the image F . The warped image is then convolved

by a shift-invariant Point Spread Function (PSF) h and finally downsampled. We denote D the downsampling operator. We can then write :

$$g_n = D \circ h \circ W_n(F) \quad (19)$$

The block diagram of the algorithm we have designed for super-resolution is outlined in Figure 1.

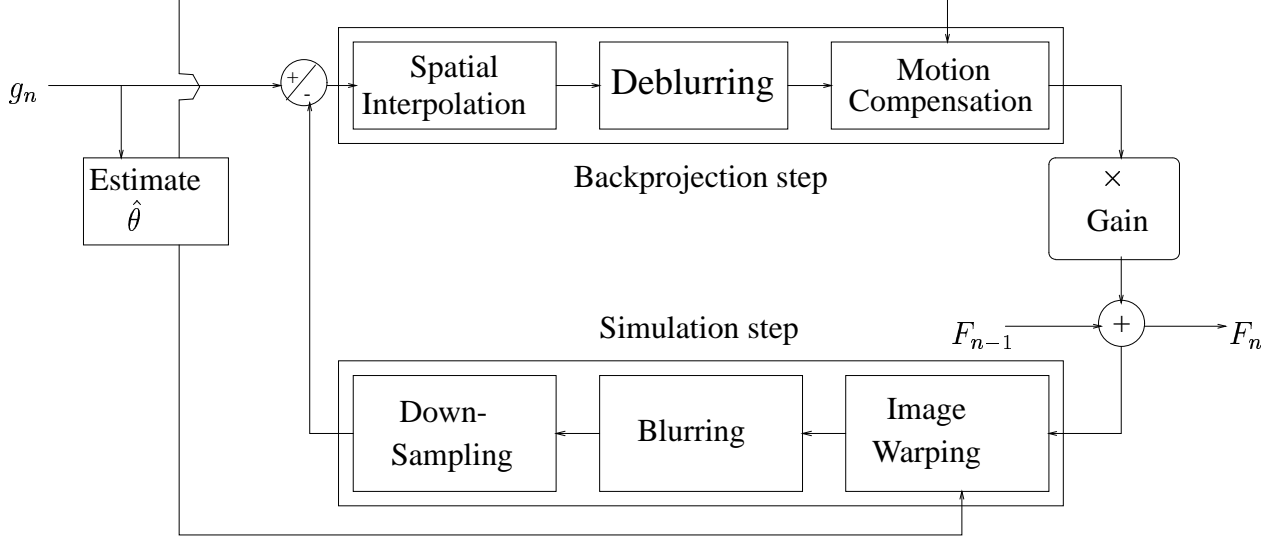


Figure 1 – Block diagram of the super-resolution algorithm

In the following, we denote \hat{F}_k the estimate of the super-resolved image at iteration k . A first estimate F_1 of the high resolution image is given by spatial interpolation of g_1 the first image of the sequence. We use cubic convolution interpolation at this early step [13]. Before starting the iterations, the forward (θ_n^{n+1}) and backward (θ_{n+1}^n) motion models between successive frames have been estimated for the N frames of the sequence ($1 \leq n \leq N - 1$). The PSF of the imaging system is supposed to be known.

In the “simulation step” (see Figure 1), we warp, blur and downsample the high-resolution estimate to get an estimate $\hat{f}_k = D \circ h \circ W_k(F_k)$ of the current image g_k of the low resolution sequence. The warping is performed according to the estimated motion model. For convenience, the super-resolved image is supposed to be reconstructed in the same frame as the first image of the sequence. Therefore we have to estimate the displacement $\vec{w}_{\theta_k^1}$ between the first and current frame of the sequence. This displacement is updated at each iteration $\vec{w}_{\theta_k^1}(s) = \vec{w}_{\theta_{k-1}^1}(s) + \vec{w}_{\theta_k^{k-1}}(s + \vec{w}_{\theta_{k-1}^1}(s))$. The warped image at each position s is given by $\hat{F}_k^c(s) = \hat{F}_k(s + \vec{w}_{\theta_k^1}(s))$ using bilinear interpolation. This warped image is blurred and downsampled to get \hat{f}_k .

The error $g_k - \hat{f}_k$ between this low-resolution estimate and the current image is then “backprojected” (see Figure 1) and added to the high-resolution estimate in order to improve it. We perform the same process for each image of the sequence. In the backprojection step, the displacement field used for motion compensation is $\vec{w}_{\theta_k^k}(s) = \vec{w}_{\theta_1^k}(s) + \vec{w}_{\theta_k^{k-1}}(s + \vec{w}_{\theta_1^k}(s))$.

The whole process can be iterated several times over the considered sequence.

This algorithm is similar to the iterative backprojection method proposed by Irani and Peleg[12]. However, in the process of backprojection, we backwarp a spatially interpolated error image instead of computing the so-called *receptive field* (in F) of the low resolution pixel. Another difference is that the high resolution estimate is updated for each frame and not after one pass of the sequence.

A high resolution image sequence can be obtained by repeating this process with successive reference images.

5 Experimental results

Our goal was to validate the use of 2D parametric motion model with simple noise reduction methods. Therefore, we first used it in combination with a temporal averaging filter [10].

We consider image triplets $\{\hat{g}_n^{n-1}, g_n, \hat{g}_n^{n+1}\}$ obtained by warping onto the central image n its two neighboring images $n - 1$ and $n + 1$. We then perform filtering over this three-frame temporal window.

We report on Figure 2 results obtained on an image sequence which involves moving objects in the foreground and a panning motion of the camera. We artificially added Gaussian noise of variance 200. Without motion compensation, the whole filtered image is blurred. After parametric motion compensation, we observe that the background is well filtered but moving objects in the foreground are still blurred. We can thus point out that motion compensation can reduce the blurring induced by the filtering, but that we have to introduce robustness to motion compensation error. This is done with the adaptive filters described in Section 3. We have reported the results for the spatio-temporal Wiener filter computed with a 3D DFT. In this case, the temporal part of the filtering is inhibited when the true motion is not accurately estimated. The noise reduction is still ensured by spatial components of the filter.

We have then determined the influence of the quality of motion estimation on the filtering performance. We have evaluated the quality of the filtered images by measuring the Signal-to-Noise Ratio improvement defined as follows :

$$R = 10 \log_{10} \frac{\sum_{i,j} (g(i,j) - f(i,j))^2}{\sum_{i,j} (\hat{f}(i,j) - f(i,j))^2} \quad (dB) \quad (20)$$

Results for two sequences are reported on Figure 3. We can point out that the motion compensation step largely improves the performance of the spatio-temporal filtering and that the use of a parametric motion model gives very similar results to those corresponding to a dense motion field while leading to a quite lower CPU load. In the particular case of the *calendar* sequence, the dominant motion being due to the global panning motion of the camera, a parametric motion model can even give better results. Furthermore, the results on the *interview* sequence show that the 3D DFT filter is less sensitive to motion estimation errors than the Boo and Bose filter, since the SNRI decreases rapidly for the non-compensated version of this latter filter as the motion goes more important.

Figure 4 shows the results of the different filters used in combination with a 6-parameter motion model estimator. We observe that Lee 3D filter has a rather poor behavior on edges. We have finally adopted the 3D DFT Wiener filter which supplies the best results in terms of restoration quality and computational cost. Let us stress that this 3D DFT filter preserves the fine details of the tree trunk in the background. These results are confirmed by results on another sequence reported on Figure 5, and by other experiments carried out on real noisy infrared aerial image sequences.

Results of the denoising and super-resolution scheme are reported on Figure 7. The sequence was degraded by an additive Gaussian noise of variance equal to 100. We considered a 6-parameter motion model. Results of the denoising (Figure 6) demonstrate that the noise is well reduced by Wiener filtering without introducing blur. Moreover, we can point out that the super-resolution algorithm produces better quality images (Figure 7.b, 7.d), i.e. more accurate, with more details and more contrasts, than spatial interpolation of the restored low resolution sequence (Figure 7.a, 7.c). Note for example that the path in the bottom right corner, which was indistinguishable on the noisy images becomes visible after processing. Furthermore, the computation time is very low. A 512×512 image can be obtained from a noisy sequence of ten noisy 256×256 images in about ten seconds on a Sun Ultra 60 at 360 Mhz.

6 Conclusion

We made a comparative study of several adaptive spatio-temporal filters for noise reduction in image sequences. These filters have been applied to image sequences which are compensated for the apparent motion beforehand. The use of parametric motion models proved to be efficient compared to a dense motion field estimator. It supplies quite satisfactory restoration quality for low computation time.

The filters that we compared were based on an adaptive weighted averaging or on the minimization of mean square error. These techniques are computationally efficient and easily tunable. Our preferred method is the spatio-temporal Wiener filter computed via 3D DFT because of its ability to preserve fine details.

The noise reduction filter was then considered as a preprocessing step for a super-resolution algorithm derived from the iterative backprojection method. In the super-resolution step, the parametric motion model is exploited again to keep a good trade-off between accuracy and computation time.

7 Acknowledgments

This work is partially supported by DGA/DSP/STTC (contract and student grant). It is carried out in cooperation with Thomson-CSF-Optrosys.

References


- [1] H. C. Andrews and B. R. Hunt. *Digital Image Restoration*. Prentice-Hall, 1977.
- [2] K. J. Boo and N. K. Bose. A motion-compensated spatio-temporal filter for image sequences with signal-dependent noise. *IEEE Transactions on Circuits and Systems for Video Technology*, 8(3):287–298, June 1998.
- [3] S. Borman and R. L. Stevenson. Simultaneous multi-frame map super-resolution video enhancement using spatio-temporal priors. In *Proceedings of the 6th IEEE International Conference on Image Processing*, Kobe, 1999.
- [4] J. C. Brailean, R. P. Kleihorst, S. N. Efstratiadis, A. K. Katsaggelos, and R. L. Lagendijk. Noise reduction filters for dynamic image sequences : A review. *Proceedings of the IEEE*, 83(9):1236–1251, September 1995.
- [5] M. P. Eckert, G. Buchsbaum, and A. B. Watson. Separability of spatiotemporal spectra of image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(12), December 1992.
- [6] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy and undersampled measured images. *IEEE Transactions on Image Processing*, 6(12):1646–1658, December 1997.
- [7] M. Elad and A. Feuer. Super-resolution reconstruction of image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(9):817–834, September 1999.
- [8] C. Fan and N. M. Namazi. Simultaneous motion estimation and filtering of image sequences. *IEEE Transactions on Image Processing*, 8(12):1788–1795, December 1999.
- [9] R. M. Gray. Toeplitz and circulant matrices: A review. Technical report, Stanford university, 1998. <http://www-isl.stanford.edu/~gray/toeplitz.pdf>.
- [10] T. S. Huang and Y. P. Hsu. Image sequence enhancement. In T. Huang, editor, *Image Sequence Analysis*, chapter 4, pages 289–309. Springer-Verlag, 1981.
- [11] B. R. Hunt and O. Kübler. Karhunen-Loeve multispectral image restoration, part I: Theory. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(3):592–599, June 1984.
- [12] M. Irani and S. Peleg. Motion analysis for image enhancement: Resolution, occlusion and transparency. *Journal of Visual Communications and Image Representation*, 4(4):324–335, December 1993.
- [13] R. G. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 29(6):1153–1160, December 1981.
- [14] A. C. Kokaram. *Motion Picture Restoration*. Springer, 1998.
- [15] P. Kornprobst, R. Deriche, and G. Aubert. Image sequence restoration : A PDE based coupled method for image restoration and motion segmentation. In *5th European Conf. on Computer Vision*, pages 548–562, Freiburg, June 1998.
- [16] R. L. Lagendijk and J. Biemond. *Iterative Identification and Restoration of Images*. Kluwer Academic Publishers, 1991.
- [17] J. S. Lee. Digital image enhancement and noise filtering by use of local statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2(2):165–168, March 1980.
- [18] J. S. Lim. Image restoration by short spectral subtraction. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 28(2):191–197, April 1980.
- [19] E. Mémin and P. Pérez. Optical flow estimation and object-based segmentation with robust techniques. *IEEE Transactions on Image Processing*, 7(5):703–719, May 1998.
- [20] J. Odobez and P. Bouthemy. Robust multiresolution estimation of parametric motion models. *Journal of Visual Communications and Image Representation*, 6(4):348–365, December 1995.
- [21] M. K. Özkan, A. T. Erdem, M. I. Sezan, and A. M. Tekalp. Efficient multiframe Wiener restoration of blurred and noisy image sequences. *IEEE Transactions on Image Processing*, 1(4):453–475, October 1992.
- [22] M. K. Özkan, M. I. Sezan, and A. M. Tekalp. Adaptive motion-compensated filtering of noisy image sequences. *IEEE Transactions on Circuits and Systems for Video Technology*, 3(4):277–288, August 1993.
- [23] R. Samy. An adaptative image sequence filtering scheme based on motion detection. In *Architectures and Algorithms for Digital Image Processing*, volume 596 of *SPIE Proceedings*, pages 135–144, 1985.
- [24] P. M. B. vanRoosmalen, S. J. P. Westen, R. Lagendijk, and J. Biemond. Noise reduction for image sequences using an oriented pyramid thresholding technique. In *Proceedings of 3rd IEEE International Conference on Image Processing*, pages 375–378, Lausanne, 1996.




a




b




c



d




e




f

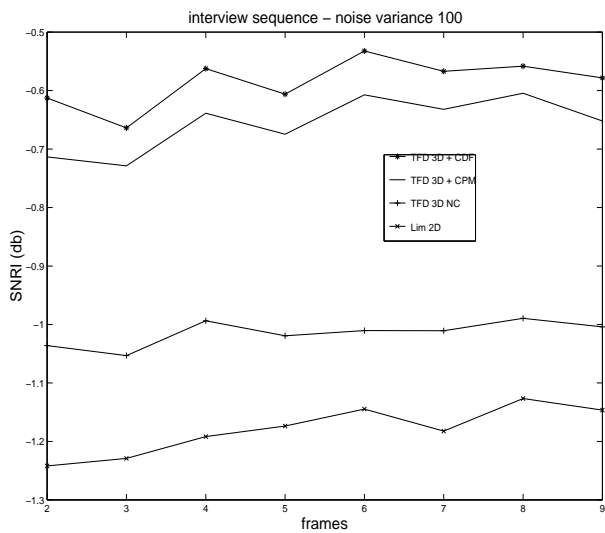
Figure 2 – (a) one image from the original image sequence, (b) one image from the noisy image sequence, (c) direct temporal averaging, (d) motion compensated temporal averaging (6-parameter model), (e) direct spatio-temporal Wiener filter computed with 3D DFT technique, (f) motion compensated spatio-temporal Wiener filter computed with 3D DFT technique



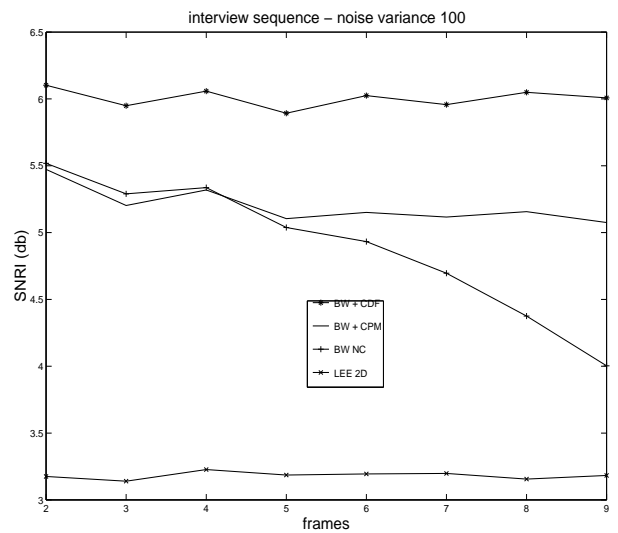
a



b



c




d


Figure 3 – Influence of motion compensation on the filtering results for (a and c) 3D DFT filter, (b and d) Boo and Bose filter. NC holds for “no motion compensation”, CDF for “compensation with a dense motion field”, and CPM is for “compensation with an affine parametric motion model”.




a




b



c



d



e



f

Figure 4 – (a) one image from the original image sequence, (b) one image from the noisy image sequence with noise variance equal to 200, filtering results using (c) 3D Lee filter, (d) AWA filter, (e) Boon and Bose filter, (f) the 3D DFT spatio-temporal filter.



a



b



c



d



e



f


Figure 5 – (a) one image from the original image sequence, (b) one image from the noisy image sequence with noise variance equal to 100, filtering results using (c) 3D Lee filter, (d) AWA filter, (e) Boo and Bose filter, (f) the 3D DFT spatio-temporal filter.



a



b



c



d

Figure 6 – (a) one image from the image sequence corrupted with Gaussian white noise of variance equal to 200, (b) image 6.a filtered with the 3D DFT Wiener filter, (c) one image from the image sequence corrupted with Gaussian white noise of variance equal to 200 and blurred with Gaussian mask of size 5×5 , (d) image 6.c filtered with the 3D DFT Wiener filter



e



f



g



h

Figure 7 – (a) image 6.b interpolated using cubic interpolation, (b) super-resolved image estimated using the denoised ten-frame sequence whose first image is given in 6.b, (c) image 6.d interpolated and deblurred, (d) super-resolved image estimated using the denoised ten-frame sequence of 6.d