



HAL
open science

Space-time adaptation for patch-based image sequence restoration

Jérôme Boulanger, Charles Kervrann, Patrick Bouthemy

► **To cite this version:**

Jérôme Boulanger, Charles Kervrann, Patrick Bouthemy. Space-time adaptation for patch-based image sequence restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2007, 29 (6), pp.1096-1102. 10.1109/TPAMI.2007.1064 . hal-02661591

HAL Id: hal-02661591

<https://hal.inrae.fr/hal-02661591>

Submitted on 30 May 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Space-Time Adaptation for Patch-Based Image Sequence Restoration

Jérôme Boulanger, Charles Kervrann, and Patrick Bouthemy

Abstract—We present a novel space-time patch-based method for image sequence restoration. We propose an adaptive statistical estimation framework based on the local analysis of the bias-variance trade-off. At each pixel, the space-time neighborhood is adapted to improve the performance of the proposed patch-based estimator. The proposed method is unsupervised and requires no motion estimation. Nevertheless, it can also be combined with motion estimation to cope with very large displacements due to camera motion. Experiments show that this method is able to drastically improve the quality of highly corrupted image sequences. Quantitative evaluations on standard artificially noise-corrupted image sequences demonstrate that our method outperforms other recent competitive methods. We also report convincing results on real noisy image sequences.

Index Terms—Image sequence restoration, denoising, nonparametric estimation, nonlinear filtering, bias-variance trade-off.

1 INTRODUCTION

IMAGE sequence restoration takes a crucial place in several important application domains. Infrared imaging, confocal microscopy, ultra-sound, and X-ray imaging are known to be noisy modalities. Restoring old films or videos is also of key importance for cultural heritage preservation. Image sequence restoration is then widely studied and still remains an active field of research. The main difficulties arise from nonstationarities observed in the spatio-temporal signals. Denoising or restoration methods must then preserve space-time discontinuities while minimizing the error between the unknown original noise-free image sequence and the denoised sequence.

A review of image sequence restoration methods can be found in [1]. These methods, especially designed for image sequences, take into account temporal correlation between images, and some of them involve a motion compensation/detection stage [2], [3], [4]. Other image sequence restoration methods can be exported from the still image denoising domain (see [5] for a recent review): Wavelet shrinkage [6], [7], Wiener filtering [8], or partial differential equations (PDE)-based methods [9] have been extended to process image sequences. However, other methods like “Total Variation” minimization [10] cannot be easily extended to the space-time domain. Recently, an extension of the recent *nonlocal means* filter [5], also related to the universal denoising—DUDE—algorithm [11], has been proposed to process image sequences and relies on the principle that the image sequence contains repeated patterns [12]. Such an approach is already popular in texture synthesis [13], inpainting [14], video completion [15], and has also been explored

- J. Boulanger and C. Kervrann are with INRA, UR341 Mathématiques et informatique appliquées, F-78352 Jouy-en-Josas and with IRISA-INRIA Rennes, Campus Universitaire de Beaulieu, F-35042 Rennes Cedex, France. E-mail: {jerome.boulanger, charles.kervrann}@irisa.fr.
- P. Bouthemy is with IRISA-INRIA Rennes, Campus Universitaire de Beaulieu, F-35042 Rennes Cedex, France. E-mail: patrick.bouthemy@irisa.fr.

Manuscript received 19 Apr. 2006; revised 28 Sept. 2006; accepted 11 Oct. 2006; published online 18 Jan. 2007.

Recommended for acceptance by S. Chaudhuri.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0305-0406. Digital Object Identifier no. 10.1109/TPAMI.2007.1064.

for image restoration [16]. Nevertheless, searching similar local patterns in the whole image sequence is infeasible in practice, even for 2D still images. Accordingly, a variant of the *nonlocal means* filter has been recently proposed in [17] and uses a preclassification of the pixels in the sequence in order to speed up the denoising procedure.

The original method we propose is a space-time patch-based adaptive statistical method for image sequence restoration. A preliminary version has been described in [18]. It is related to the statistical framework described for still images [19], [21], [22], [23], [24], and image sequences [20], [25]. Unlike robust anisotropic diffusion [26] and nonlinear Gaussian filtering [27], our approach supplies scale selection by estimating the appropriate space-time filtering window at each pixel. Moreover, the proposed method differentiates the space and time dimensions unlike other methods that consider the image sequence as an isotropic 3D volume [6], [9]. As a matter of fact, naive approaches can introduce motion blur and artifacts if the time dimension is merged with spatial dimension. In contrast to [21], [25], our approach is not based on a geometrical partition of the neighborhood in sectors. It uses a fixed neighborhood geometry but involves an appropriate weighted sum of data points in an adaptive neighborhood, which is far more flexible and efficient. The weights are defined by computing a distance between a local patch centered at the considered pixel \mathbf{x}_i and patches taken in an adapted space-time neighborhood. Additionally, a confidence level, that is the point-wise estimator variance, attached to each restored pixel is provided.

The remainder of the paper is organized as follows: Section 2 describes our general framework for image sequence restoration. In Section 2.2, the adaptive choice of the local space-time neighborhood is introduced. Section 2.3 deals with the similarity measure involved in the selection of patches in the space-time neighborhood. Section 3 gives details of the algorithm implementation. In Section 4, we report an important set of experimental results on artificially noise-corrupted video sequences as well as on a real noisy infrared image sequence. Intensive comparison with other recent methods has been carried out, demonstrating that our method outperforms most of them. We also present how our denoising method can be combined with a motion estimation method if required. We demonstrate from experiments that motion estimation and transforms into the frequential domain are not necessary for denoising image sequences with good results. In the case of camera motion represented by a quadratic model and reliably estimated using a robust statistical technique, the denoising results are slightly improved but additional computations are required. Finally, Section 5 contains concluding remarks.

2 PATCH-BASED SPACE-TIME APPROACH

2.1 Model Description

We consider the following statistical image model:

$$Y_i = u(\mathbf{x}_i) + \epsilon_i, \quad (1)$$

where $\mathbf{x}_i \in \Omega$ denotes the pixel location in the space-time volume $\Omega \subset \mathbb{R}^3$. The function $u_i = u(\mathbf{x}_i)$ is the ideal image to be recovered from observations Y_i . The errors ϵ_i are assumed to be independent zero-mean Gaussian variables with unknown variance τ^2 .

We need minimal assumptions on the structure of the image for recovering u . In what follows, we assume that there exists repeated small image patches of the patch centered at \mathbf{x}_i in the space-time neighborhood of pixel \mathbf{x}_i . However, the size and shape of this neighborhood will vary in the image sequence because of




Fig. 1. Patch-based space-time framework. To each point of the image sequence is associated an estimated space-time neighborhood $W_{i,n}$. To each point x_j of this neighborhood is associated a 3×3 patch (7×7 image patches are used in practice). Every weight ω_{ij} is defined as a function of the distance between the patch centered in point x_i and the patch centered in point $x_j \in W_{i,n}$.

nonstationarities and the presence of spatial and temporal discontinuities. If we can determine the adequate neighborhood for each pixel, then the regression function u can be estimated by optimizing a local maximum likelihood (ML) criterion. The proposed method addresses these two issues as described below.

As in [23], [24], we design a sequence of increasing nested space-time neighborhoods $\{W_{i,n}\}_{n \in [0:N]}$, centered at each point x_i , i.e., $W_{i,n} \subset W_{i,n+1}$, with N denoting the index of the largest neighborhood. Additional details about the neighborhood design are given in Section 2.2. As for initialization, we choose the smallest neighborhood (the eight nearest neighbors in the 2D space domain) as the *pilot* (starting) neighborhood $W_{i,0}$ at x_i . Then, we compute an initial estimate $\hat{u}_{i,0}$ of $u(x_i)$ and its associated variance $\text{var}(\hat{u}_{i,0})$ as

$$\hat{u}_{i,0} = \sum_{x_j \in W_{i,0}} \omega_{ij} Y_j \quad \text{and} \quad \text{var}(\hat{u}_{i,0}) = \hat{\tau}^2 \sum_{x_j \in W_{i,0}} \omega_{ij}^2, \quad (2)$$

where $\hat{\tau}^2$ is an empirical estimate of the noise variance τ^2 as described in Section 3. At the initialization step, the weights ω_{ij} are defined as a function of the distance between two spatial $p \times p$ image patches or space-time $p \times p \times q$ image patches (of fixed size) centered at x_i and x_j , respectively. In what follows, the $p \times p$ (or $p \times p \times q$) patch size is assumed to be fixed for all the iterations of the restoration procedure. In addition, there is no real difference for the proposed method between the spatial only and space-time patch since the spatial only patch can be considered as a space-time patch with the temporal dimension q equal to one.

At the first iteration, we consider a larger space-time neighborhood $W_{i,1}$ such that $W_{i,0} \subset W_{i,1}$ and calculate new estimates $\hat{u}_{i,1}$ and $\text{var}(\hat{u}_{i,1})$ over $W_{i,1}$. We continue this way, and at iteration n , we define the estimator as

$$\hat{u}_{i,n} = \sum_{x_j \in W_{i,n}} \omega_{ij} Y_j \quad \text{and} \quad \text{var}(\hat{u}_{i,n}) = \hat{\tau}^2 \sum_{x_j \in W_{i,n}} \omega_{ij}^2, \quad (3)$$

where the estimator $\hat{u}_{i,n}$ corresponds to a weighted average of the intensities Y_j located in the space-time neighborhood $W_{i,n}$. We propose to define weights ω_{ij} as a function of the distance between two image patches of fixed size $\hat{u}_{i,n-1}$ and $\hat{u}_{j,n-1}$ estimated at iteration $n-1$ and centered in x_i and x_j , respectively, as shown in Fig. 1. In Fig. 1, 3×3 patches are used for illustration purposes but 7×7 patches are typically used in our experiments to produce satisfying restoration results (see Section 4). In the two next

sections, we specify the sequence of neighborhoods $\{W_{i,n}\}_{n \in [0:N]}$ and we propose a suitable distance to compare image patches.

2.2 Space-Time Neighborhood Adaptation

2.2.1 Space-Time Neighborhood Geometry

One important contribution of this work is the online construction of the space-time neighborhood sequence $\{W_{i,n}\}_{n \in [0:N]}$. First, we consider a simple hyper-cube space-time volume as the neighborhood shape. Also, we separate the space dimension from the time dimension. Consequently, space-time neighborhoods are parametrized by their extent in the space domain and their extent in the time domain. The use of two distinct extents (one is for the space dimensions and the other one is for the time dimension) allows us to respect space-time discontinuities and the image sequence is not considered as an isotropic 3D volume. In Fig. 2a, the increasing neighborhood sequence is illustrated: At each iteration, the spatial extent and the temporal extent are alternatively increased until a stopping rule is satisfied. Then, the growth of the neighborhood is stopped for this direction (e.g., time) and continues in the other direction (e.g., space) until the stopping rule is again satisfied. The next paragraph will explicitly explain the considered stopping rule.

2.2.2 Space-Time Neighborhood Selection

A point-wise rule is used to guide the space-time neighborhood selection process. This rule aims at selecting the optimal neighborhood at x_i and is based on the measure of the closeness of the estimator $\hat{u}_{i,n}$ obtained at iteration n to the unknown function u_i given by the local L_2 risk. This local measure of performance, used to take into account the nonstationarities in the image sequence, can be decomposed in two terms, that is, the squared bias and the variance of the estimator as

$$\mathbb{E}[(\hat{u}_{i,n} - u_i)^2] = [\text{bias}(\hat{u}_{i,n})]^2 + \text{var}(\hat{u}_{i,n}), \quad (4)$$

where $\mathbb{E}[\cdot]$ denotes the mathematical expectation. In the sequel, we can reasonably assume that the squared bias is an increasing function of the neighborhood size and the variance is a decreasing function of the neighborhood size [19], [21], [22], [23], [24], [28]. Fig. 3 illustrates the fundamental bias-variance trade-off principle. Theoretically, it is not guaranteed that the pointwise bias and variance, respectively, increases and decreases with the window size mainly because the weights are computed from the restored image obtained




Fig. 2. (a) Spatio-temporal neighborhood: Colors correspond to iterations plotted in (b). (b) Confidence intervals: Circles represent estimates $\hat{u}_{i,n}$ obtained at each iteration n . The gray rectangles represent the intersection between the current confidence interval and the previous one. As long as the estimate belongs to this intersection, the estimation process is updated.

at the previous iterations of the procedure. However, we have proven that two theoretical upper bounds for each term correspond to strictly monotonous increasing and decreasing functions in [23], [24]. In that case, it is implicitly assumed that the image is continuous Lipschitz: $|u(\mathbf{x}_j) - u(\mathbf{x}_i)| \leq C_1 |\mathbf{x}_j - \mathbf{x}_i|$, $C_1 < \infty$. This means that, in some critical situations, this assumption could be violated. Finally, optimal neighborhood $W_{i,n}^*$ will then be one achieving an optimal compromise between the two upper bounds for the bias and variance terms. Furthermore, it is admitted that this compromise is reached when the squared bias and the variance are nearly the same [28]: $\mathbb{E}[(u_{i,n}^* - u_i)^2] \approx \text{var}(u_i^*)$.

A practical rule corresponding to the optimal compromise and based on a pairwise comparison of successive estimates can be derived to select the optimal neighborhood [22]. It amounts to defining the largest neighborhood satisfying the point-wise statistical rule [19], [22], [23], [24], [28]:

$$|\hat{u}_{i,n} - \hat{u}_{i,n'}| < \eta \text{std}(\hat{u}_{i,n'}), \forall n' < n \quad (5)$$

as the optimal neighborhood where $\text{std}(\hat{u}_{i,n})$ is the standard deviation of $\hat{u}_{i,n}$. The related demonstration can be found in [23], [24] and this rule can be interpreted as follows: While the successive estimates $\hat{u}_{i,n}$ are sufficiently close to each other, we continue the estimation process. More specifically, the estimation process is continued, while new estimates belong to the intersection of previously estimated confidence intervals $[\hat{u}_{i,n} - \eta \text{std}(\hat{u}_{i,n}), \hat{u}_{i,n} + \eta \text{std}(\hat{u}_{i,n})]$ (see Fig. 2). Besides, let us point out that we do not need to store all the previous estimates $\{\hat{u}_{i,n'}\}_{n' < n}$, but only the current intersection of confidence intervals, the last estimate, and its variance for each pixel. Finally, the factor η can be easily chosen in the range [2, 4] as justified with statistical arguments in [19], [21], [23], [24].

2.3 Similarity Measure for Patch Selection

In contrast to geometry-based approaches [21], [25], we use weights that allow us to select the correct data points in the neighborhood for averaging. This selection is based on the similarity between a given spatial $p \times p$ or a space-time $p \times p \times q$ image patch at \mathbf{x}_i and $p \times p$ (or a $p \times p \times q$) image patches at \mathbf{x}_j belonging to the space-time neighborhood $W_{i,n}$. Such patches are able to capture texels and local geometry in images. The L_2 distance is widely used for a similarity measure between image patches. However, in order to take into account the local variance of the estimator, we introduce the following symmetric distance between image patches $\hat{\mathbf{u}}_{i,n-1}$ and $\hat{\mathbf{u}}_{j,n-1}$:

$$\Delta_{ij} = \frac{1}{2} \left[(\hat{\mathbf{u}}_{i,n-1} - \hat{\mathbf{u}}_{j,n-1})^\top \hat{\mathbf{V}}_{i,n-1}^{-1} (\hat{\mathbf{u}}_{i,n-1} - \hat{\mathbf{u}}_{j,n-1}) + (\hat{\mathbf{u}}_{i,n-1} - \hat{\mathbf{u}}_{j,n-1})^\top \hat{\mathbf{V}}_{j,n-1}^{-1} (\hat{\mathbf{u}}_{i,n-1} - \hat{\mathbf{u}}_{j,n-1}) \right], \quad (6)$$

where the two vectors $\hat{\mathbf{u}}_{i,n-1}$ and $\hat{\mathbf{u}}_{j,n-1}$ denote the spatial $p \times p$ (or space-time $p \times p \times q$) image patches respectively centered at \mathbf{x}_i and \mathbf{x}_j . The two matrices $\hat{\mathbf{V}}_{i,n-1}$ and $\hat{\mathbf{V}}_{j,n-1}$ are diagonal with the diagonal elements equal to $\text{var}(\hat{u}_{i,n})$ and $\text{var}(\hat{u}_{j,n})$, respectively. We decide that the two vectors $\hat{\mathbf{u}}_{i,n-1}$ and $\hat{\mathbf{u}}_{j,n-1}$ are similar with a probability of false alarm $1 - \alpha$, under the hypothesis that they are Gaussian distributed, using a classical χ^2 test with p^2 ($p^2 q$ for space-time) degrees of freedom. In other words, if $\Delta_{ij} < \lambda_\alpha$, with λ_α chosen as a quantile of a $\chi_{p^2, 1-\alpha}^2$ distribution ($\chi_{p^2 q, 1-\alpha}^2$ for space-time), we can decide that the two patches are similar. In our experiments, we use a confidence level of 99 percent and set α to 0.01.

The distance Δ_{ij} is transformed into a similarity measure using the exponential kernel. We compute the similarity measure for all the points of the neighborhood and normalize it to obtain the following expression for the weights:

$$\omega_{ij} = \frac{\exp\left(-\frac{\Delta_{ij}}{2\lambda_\alpha}\right)}{\sum_{\mathbf{x}_j \in W_{i,n}} \exp\left(-\frac{\Delta_{ij}}{2\lambda_\alpha}\right)}. \quad (7)$$

If the distance Δ_{ij} between two patches is large, then the weight ω_{ij} associated to pixel \mathbf{x}_j is small and the pixel will not participate in the estimation at point \mathbf{x}_i . Consequently, this weighting provides an efficient and flexible way to select the appropriate pixels




Fig. 3. Illustration of the bias-variance trade-off principle. When the size of the neighborhood increases, the bias of the estimator increases while taking into account more and more intensity samples, and the variance decreases. The parameters α and β are two unknown constants.

contributing to the estimation of u_i in the adaptive space-time neighborhood while effectively preserving space-time discontinuities. Let us note that the process is entirely data-driven and does not require any particular geometry adapted to image contents as proposed in [25].

2.4 Motion Compensation

The motion of the image sequence can be taken into account in order to apply the proposed method along the direction of the motion. However, dense motion estimation is known to be a difficult task in noisy contexts [12]. Then, we propose robustly estimating a global parametric motion model only, which is able to capture the dominant image motion due to the camera movement. This is achieved using the multiresolution robust method described in [29]. A similar exploitation of a parametric motion compensation was proposed in [8] and associated with a 3D Wiener filtering technique.

Once the dominant motion has been estimated, the filtering along the motion direction can be considered in three ways. A naive one would be to warp all the frames in the referential of the first frame. Because of the accumulation of errors and interpolations, it turns out that such a scheme does not improve the performance of the denoising process. The second way is to compensate the motion within the space-time neighborhood by warping the frames into the referential of the current frame, but it also involves interpolations. Then, we propose to avoid interpolation by adapting the shape of the space-time neighborhood according to the estimated dominant motion. This is achieved by translating the patch at point $\mathbf{x}_j = (x_j, y_j, t_j)$ with displacement given by the estimated parametric motion model at the center (x_i, y_i, t_i) of the neighborhood $W_{i,n}$. Moreover, when using $p \times p \times q$ space-time image patches, the motion also has to be compensated into the patch. Experiments show that this third way is able to improve the performance of the proposed method.

3 ALGORITHM IMPLEMENTATION

At the initialization, the noise variance τ^2 first has to be estimated. It can be robustly estimated by calculating pseudoresiduals ε_i as described in [30]. We consider four spatial neighbors and two temporal ones, and pseudoresiduals are compactly represented by $\varepsilon_i = (8Y_i - \Delta Y_i)/\sqrt{42}$, where ΔY_i is the discrete Laplacian of Y_i at \mathbf{x}_i (see (1)) and the constant $\sqrt{42}$ is introduced to ensure that $\mathbb{E}[\varepsilon_i^2] = \tau^2$. Given the residuals ε_i , the noise variance τ^2 can then be robustly estimated as

$$\tau = 1.4826 \operatorname{med}_i(|\varepsilon_i - \operatorname{med}_j|\varepsilon_j|). \quad (8)$$

Let us recall that parameter λ_α is set to the 0.99 quantile of the $\chi_{p^2,0.99}^2$ distribution with a size of patch p chosen within $\{3, 5, 7, 9\}$. The last parameter η is set to $2\sqrt{2}$ to ensure a good accuracy of the estimation [19], [22], [28]. During the estimation, spatial and temporal extents of the space-time neighborhoods are alternatively increased (see Section 2.2).

Furthermore, the algorithm can be easily parallelized. Indeed, the estimation steps use only local information and, thus, we have distributed the computation load over eight CPUs dividing the computation time by eight. Finally, another possibility to speed up the algorithm is to use a dyadic scheme when increasing the extent of the neighborhood. The proposed method is very simple to implement and does not require the fine adjustment of parameters λ_α and η which control the estimation process.




Fig. 4. Performance of our denoising method for several noise levels and for several patch sizes including a no-patch version (pixel-wise). The PSNR is used to measure the overall performance of the filtering and the test sequence is “Akyio” ($176 \times 144 \times 300$). This sequence mainly exhibits low motion. Experiments show that the introduction of patches improves the PSNR of at least 2 dB. We can also point out that the results for patches 7×7 and 9×9 are similar.

4 EXPERIMENTAL RESULTS

In this section, a large set of experiments are reported to validate our patch-based space-time adaptive estimation method. We first consider real image sequences with artificially added Gaussian white noise. Using this usual protocol, we compare our method to other recent methods for denoising image sequences. For an objective performance evaluation, we consider the global measure given by the Peak-Signal-to-Noise-Ratio defined as $\text{PSNR} = 20 \log_{10}(255/\sqrt{\text{MSE}})$ (MSE denotes the mean squared error between the original noise-free image sequence and the denoised image sequence). The visual quality of the image sequence is also taken into account. We then discuss the usefulness of motion compensation and we comment on the respective performance of 2D spatial patches and 3D space-time patches. Finally, the proposed method is applied to a real noisy infrared image sequence.

4.1 Performance Assessment

We first report experiments to evaluate the influence of the noise level and the spatial patch size on the overall method performance. The patch size is generally a free-parameter, assumed to be fixed for the whole restoration procedure. Fig. 4 plots PSNR values obtained for eight noise levels and five patch sizes. First, we can note that the patch-based method performance is smoothly affected with the increase of the noise level. The improvement gained by introducing patches (to be compared to the no-patch version) is clearly demonstrated. As it could be expected, it is useless to consider patches that are too large. When the size of the patch increases, the PSNR increases too; however, results for fixed sizes 7×7 and 9×9 are quite similar while the computation time is proportional to the number of pixels in the patch.

The well-known sequence “Flower garden” is shown in Fig. 5 to illustrate the visual quality of the denoised image sequence in very noisy conditions. In order to give insights into the spatio-temporal behavior of the denoising method, we have displayed XT slices of the image sequence. The reported results demonstrate that our method can cope with the presence of motion while preserving temporal discontinuities and reaches a PSNR value of 23.59 dB using fixed 7×7 patches and six iterations. By applying our own



Fig. 5. Sequence “Flower garden.” (a) One image of the original sequence. (b) The same image of the noisy sequence with an additive Gaussian white noise of standard deviation $\tau = 30$ (PSNR = 18.58 dB). (c) The corresponding image of the denoised sequence, PSNR=23.59 dB. (d), (e), and (f) represent corresponding XT slices in the space-time domain. The camera motion induces lines in the XT slices.

implementation of the *nonlocal means* algorithm [12] on the “Flower garden” sequence using fixed 7×7 patches in fixed $21 \times 21 \times 3$ neighborhoods and choosing a bandwidth equal to $h = 12\tau$, we get a PSNR value equal to 21.14 dB only.

4.2 Comparison with Other Recent Methods

We have compared our method with four other recent methods: a combination of a spatial Wiener filter with a motion-compensated temporal Kalman filter [2], a space-time nonlinear adaptive K-NN filter [4], a 3D wavelet-based method [6], a 3D point-wise adaptive estimate using different neighborhood geometries [25], and a local adaptive sliding window DCT (SDCT) denoising method [20]. For these experiments, the motion-compensation stage is not applied.

Eight test sequences corrupted with noise of different levels are used. For a fair evaluation, we have considered the results supplied by the authors themselves in the referenced papers. Therefore, we cannot provide the PSNR measures for all the test sequences and Table 1 contains all the available results. Our method clearly outperforms most of the other methods since it supplies the best PSNR results for nearly all the tested sequences, sometimes with quite a significant improvement (up to 4 dB). Let us also stress that the implementation of our method is straightforward and our method involves no parameter tuning for the tested sequences. The more competitive (SDCT) denoising method [20] involves a hard thresholding of 3D DCT coefficients and minimization techniques

for block selection in the temporal domain. This approach is then somewhat based on motion estimation. In our approach, neither motion estimation is required nor image transformation into the frequential domain. Finally, it appears conceptually simpler with similar results for many image sequences.

4.3 Motion Compensation

In this section, we aim at evaluating the impact of a motion compensation stage on the performance of our method. We have applied this version to the “Avenger” sequence which contains two moving cars tracked by the camera mounted in an helicopter. Figs. 6a and 6b, respectively, show one image extracted from the original image sequence and one of the noisy image sequence with an additive Gaussian white noise of standard deviation $\tau = 20$. Fig. 6c shows the result of the version of our method without motion-compensation, while Fig. 6e shows the improvement supplied by the motion-compensation stage. The PSNR difference between the two image sequences is about 1 dB and the visual quality is also improved. We can then conclude that, when the global motion the image sequence is well described by a 2D parametric model, the proposed motion compensation scheme improves the quality of the denoising process. In that case, a quadratic motion model has been considered. Let us add that the motion of the two cars is not handled by the dominant motion model. However, since our method involves a data-driven adaptation scheme, the

TABLE 1
PSNR Results for Eight Test Sequences and Five Denoising Methods

Sequence name	Sequence size	noisy sequence	(a)	(b)	(c)	(d)	(e)	(f)
<i>Akiyo</i>	$176 \times 144 \times 300$	22 dB	–	–	–	33.86 dB	–	34.31 dB
<i>Salesman</i>	$176 \times 144 \times 449$	28 dB	34.4 dB	32.5 dB	–	–	37.01 dB	35.13 dB
		24 dB	31.1 dB	–	–	34.83 dB	32.60 dB	
<i>Flower garden</i>	$352 \times 240 \times 115$	28 dB	–	28.2 dB	–	–	31.25 dB	31.33 dB
<i>Miss America</i>	$176 \times 144 \times 108$	28 dB	–	35.3dB	–	–	–	39.39 dB
<i>Miss America</i>	$128 \times 128 \times 128$	7 dB	–	–	26.36 dB	–	–	26.69 dB
		12 dB	–	–	28.16 dB	–	–	29.63 dB
		17 dB	–	–	30.46 dB	–	–	32.05 dB
		22 dB	–	–	32.66 dB	–	–	34.20 dB
<i>Suzie</i>	$176 \times 144 \times 150$	28 dB	34.8 dB	–	–	–	–	37.07 dB
		24 dB	32.0 dB	–	–	–	–	35.11 dB
<i>Trevor</i>	$176 \times 144 \times 90$	28 dB	33.9 dB	34.1 dB	–	–	–	36.68 dB
		24 dB	31.3 dB	–	–	–	–	34.79 dB
<i>Foreman</i>	$176 \times 144 \times 300$	28 dB	33.9 dB	–	–	–	–	34.94 dB
		24 dB	31.1 dB	–	–	–	–	32.90 dB

(a) Joint Kalman and Wiener denoising with motion compensation using a dense motion field [2]. (b) Adaptive K-NN space-time filter [4]. (c) Wavelet-based method for image sequence denoising: TIWP3D [6]. (d) Three-dimensional nonparametric regression approach [25]. (e) Local adaptive sliding window DCT (SDCT) denoising method [20]. (f) The proposed adaptive method with 7×7 patches and six iterations without motion compensation. Numerical results for the other methods are taken from the related publications. The symbol “–” means that results were not provided by the authors for the tested sequence.




Fig. 6. Sequence "Avenger." (a) $384 \times 288 \times 12$ original sequence. (b) Noisy sequence with an additive Gaussian white noise of standard deviation $\tau = 20$, denoised sequence with the proposed method using: (c) fixed 5×5 patches and no motion compensation stage (5 min), (d) fixed $3 \times 3 \times 3$ patches and no motion compensation stage (6 min), (e) fixed 5×5 patches with the motion compensation stage using a quadratic motion model (9 min), and (f) fixed $3 \times 3 \times 3$ patches with the motion compensation stage using a quadratic motion model (20 min). The computation time is indicated for a Linux PC with 8×3 GHz CPU.

neighborhoods for the pixels belonging to the two cars are essentially 2D spatial neighborhoods.

4.4 Space-Time Patches

Fig. 6d shows one image of the "Avenger" sequence denoised using the proposed method with space-time fixed $3 \times 3 \times 3$ patches. It can be compared to the result obtained with spatial fixed 5×5 patches and shown in Fig. 6c. In the two cases, the number of intensity values used to compute the similarity measure (see (7)) is approximately the same and the PSNR difference is negligible. Nevertheless, the use of space-time patches increases the temporal stability of the reconstructed structures along the image sequence, which is an important point for visual quality.

4.5 Experiment on a Real Image Sequence

Fig. 7 reports an experiment on a real infrared sequence which is naturally noisy. It is taken from a plane approaching a harbor with boats and a moving vehicle on the pier. The noise standard deviation is estimated to 7.3. The contrast of some structures on the ground is very low and the sequence is shaking due to the plane vibrations. We once again use a quadratic motion model to estimate the dominant motion. Let us recall that this model is exact in the case of a rigid motion and a planar scene. The proposed method (Fig. 7c) can be

favorably compared to our own implementation of the *nonlocal means* algorithm (Fig. 7b) [12]. The details of the images like the small vehicle are better restored while the noise has been well removed. Finally, temporal discontinuities of the sequence due to the vibrations are also preserved.

5 CONCLUSION

We have described a novel and unsupervised method able to drastically reduce additive noise in image sequences. The proposed method is based upon an adaptive estimation statistical framework. It can specify, in a simple data-driven way, the most appropriate space-time neighborhood and associate weights to select the data points involved in the intensity estimation process at each pixel. Moreover, it involves a patch-based approach extended to the space-time domain. All the parameters of the algorithm are well calibrated and our method does not require any fine parameter tuning. The fixed patch size is the only free parameter to be chosen by the user. Quite satisfactory results have been obtained on several image sequences. Furthermore, it was experimentally demonstrated that our method outperforms most of recent methods. The visual quality improvement of the denoised image sequences is noticeable since noise is well smoothed out while spatial and



Fig. 7. Infrared sequence. (a) One image of the original sequence. (b) The corresponding image of the denoised sequence with the *nonlocal means* filter applied with a patch of fixed size 5×5 pixels and the bandwidth equal to four times the noise standard deviation. (c) Denoised image with the proposed motion compensated filter, (b1) Cropped region of (b) and (c1) Cropped region of (c).

temporal discontinuities are well preserved. Actually, the computation time (for a 8×3 GHz CPU) is 60 sec/frame (384×228 images) for both 5×5 and $3 \times 3 \times 3$ image patches (with no motion compensation) but a fast implementation has been recently proposed to reduce the computation by a factor of 8. Finally, some improvements are proposed to incorporate a motion-compensation stage improving the performance of the proposed method.

REFERENCES

- [1] J. Brailean, R. Kleihorst, S. Efstratiadis, A. Katsaggelos, and R. Lagendijk, "Noise Reduction Filters for Dynamic Image Sequences: A Review," *Proc. IEEE*, vol. 83, no. 9, pp. 1272-1291, 1995.
- [2] R. Dugad and N. Ahuja, "Video Denoising by Combining Kalman and Wiener Estimates," *Proc. IEEE Int'l Conf. Image Processing*, vol. 4, pp. 152-156, Oct. 1999.
- [3] H.-Y. Cheong, A. Tourapis, J. Llach, and J. Boyce, "Adaptive Spatio-Temporal Filtering for Video-Denoising," *Proc. IEEE Int'l Conf. Image Processing*, vol. 2, pp. 965-968, Oct. 2004.
- [4] V. Zlolkica and W. Philips, "Motion and Detail Adaptive Denoising in Video," *Proc. SPIE—Image Processing: Algorithms and Systems III*, vol. 5298, pp. 403-412, 2004.
- [5] A. Buades, B. Coll, and J. Morel, "A Review of Image Denoising Algorithms, with a New One," *Multiscale Modeling and Simulation: A SIAM Interdisciplinary J.*, vol. 4, no. 2, pp. 490-530, 2005.
- [6] N. Rajpoot, Z. Yao, and R. Wilson, "Adaptive Wavelet Restoration of Noisy Video Sequences," *Proc. IEEE Int'l Conf. Image Processing (ICIP '04)*, pp. 957-960, 2004.
- [7] F. Shi and I.V. Selesnick, "Video Denoising Using Oriented Complex Wavelet Transforms," *Proc. IEEE Int'l Conf. Acoustics, Speech, and Signal Processing*, vol. 2, pp. 949-952, 2004.
- [8] F. Dekeyser, P. Bouthemy, and P. Pérez, "Spatio-Temporal Wiener Filtering of Image Sequences Using a Parametric Motion Model," *Proc. IEEE Int'l Conf. Image Processing*, pp. 208-211, Sept. 2000.
- [9] S.H. Lee and M.G. Kang, "Spatio-Temporal Video Filtering Algorithm Based on 3-D Anisotropic Diffusion Equation," *Proc. IEEE Int'l Conf. Image Processing*, vol. 3, no. 2, pp. 447-450, 1998.
- [10] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear Total Variation Based Noise Removal Algorithms," *Physica D*, vol. 60, nos. 1-4, pp. 259-268, 1992.
- [11] G. Motta, E. Ordentlich, I. Ramirez, G. Seroussi, and M. Weinberger, "The Dude Framework for Continuous Tone Image Denoising," *Proc. IEEE Int'l Conf. Image Processing*, vol. 3, pp. 345-348, 2005.
- [12] A. Buades, B. Coll, and J. Morel, "Denoising Image Sequences Does Not Require Motion Estimation," Preprint CMLA, 2005.
- [13] A. Efros and T. Leung, "Texture Synthesis by Nonparametric Sampling," *Proc. Seventh IEEE Int'l Conf. Computer Vision*, pp. 1033-1038, 1999.
- [14] A. Criminisi, P. Pérez, and K. Toyama, "Region Filling and Object Removal by Exemplar-Based Inpainting," *IEEE Trans. Image Processing*, vol. 13, no. 9, pp. 1200-1212, 2004.
- [15] Y. Wexler, E. Shechtman, and M. Irani, "Space-Time Video Completion," *IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 120-127, 2004.
- [16] D. Zang and Z. Wang, "Image Information Restoration Based on Long-Range Correlation," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 12, no. 5, pp. 331-341, May 2002.
- [17] M. Mahmoudi and G. Sapiro, "Fast Image and Video Denoising via Nonlocal Means of Similar Neighborhoods," *SPLetters*, vol. 12, no. 12, pp. 839-842, Dec. 2005.
- [18] J. Boulanger, C. Kervrann, and P. Bouthemy, "Adaptive Spatio-Temporal Restoration for 4D Fluorescence Microscopic Imaging," *Proc. Eighth Int'l Conf. Medical Image Computing and Computer Assisted Intervention*, Oct. 2005.
- [19] J. Polzehl and V. Spokoiny, "Adaptive Weights Smoothing with Applications to Image Restoration," *J. Royal Statistical Soc.: Series B (Statistical Methodology)*, vol. 62, no. 2, pp. 335-354, 2000.
- [20] D. Rusanovskyy and K. Egiazarian, "Video Denoising Algorithm in Sliding 3D DCT Domain," *Proc. Seventh Advanced Concepts for Intelligent Vision Systems*, Sept. 2005.
- [21] V. Katkovnik, K. Egiazarian, and J. Astola, "Adaptive Window Size Image Denoising Based on Intersection of Confidence Intervals (ICI) Rule," *J. Math. Imaging and Vision*, vol. 16, no. 3, pp. 223-235, 2002.
- [22] C. Kervrann, "An Adaptive Window Approach for Image Smoothing and Structures Preserving," *Proc. Eighth European Conf. Computer Vision*, pp. 132-144, 2004.
- [23] C. Kervrann and J. Boulanger, "Optimal Spatial Adaptation for Patch-Based Image Denoising," *IEEE Trans. Image Processing*, vol. 15, no. 10, pp. 2866-2878, 2006.
- [24] C. Kervrann and J. Boulanger, "Unsupervised Patch-Based Image Regularization and Representation," *Proc. Ninth European Conf. Computer Vision*, May 2006.
- [25] C. Ercole, A. Foi, V. Katkovnik, and K. Egiazarian, "Spatio-Temporal Pointwise Adaptive Denoising in Video: 3D Non Parametric Approach," *Proc. First Int'l Workshop Video Processing and Quality Metrics for Consumer Electronics*, 2005.
- [26] M. Black, G. Sapiro, D. Marimont, and D. Heeger, "Robust Anisotropic Diffusion," *IEEE Trans. Image Processing*, vol. 7, no. 3, pp. 421-432, 1998.
- [27] D. Barash, "A Fundamental Relationship between Bilateral Filtering, Adaptive Smoothing, and the Nonlinear Diffusion Equation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 6, pp. 844-847, June 2002.
- [28] O. Lepski, "Asymptotically Minimax Adaptive Estimation 1: Upper Bounds," *SIAM J. Theory of Probability and Application*, vol. 36, no. 4, pp. 654-659, 1991.
- [29] J. Odobez and P. Bouthemy, "Robust Multiresolution Estimation of Parametric Motion Models," *J. Visual Comm. and Image Representation*, vol. 6, no. 4, pp. 348-365, Dec. 1995.
- [30] T. Gasser, L. Sroka, and C.J. Steinmetz, "Residual Variance and Residual Pattern in Nonlinear Regression," *Biometrika*, vol. 73, pp. 625-633, 1986.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.