Combined Wavelet Domain and Temporal Video Denoising

Aleksandra Pizurica, Vladimir Zlokolica and Wilfried Philips
Ghent University
Department for Telecommunications and Information Processing
Sint-Pietersnieuwstrat 41, 9000 Ghent, Belgium
Aleksandra.Pizurica@telin.ugent.be

Abstract

We develop a new filter which combines spatially adaptive noise filtering in the wavelet domain and temporal filtering in the signal domain. For spatial filtering, we propose a new wavelet shrinkage method, which estimates how probable it is that a wavelet coefficient represents a "signal of interest" given its value, given the locally averaged coefficient magnitude and given the global subband statistics. The temporal filter combines a motion detector and recursive time-averaging. The results show that this combination outperforms single resolution spatio-temporal filters in terms of quantitative performance measures as well as in terms of visual quality. Even though our current implementation of the new filter does not allow real-time processing, we believe that its optimized software implementation could be used for real- or near real-time filtering.

1 Introduction

Video denoising is important in numerous applications, including restoration of old movies, television broadcasting systems, teleconferencing and video surveillance. In surveillance applications video cameras often operate in bad-lighting conditions, which may yield extremely noisy video output. Not only the noise produces unpleasant perceptual effect, but in certain extreme cases it can even seriously hamper the interpretation of the video content. Development of advanced video denoising schemes is thus essential for a consumer-grade visual quality. Due to constrains of real-time implementation current video denoising schemes are mostly confined to relatively simple spatio-temporal filters [1–5]. While achieving the required computational speed, these filters usually cannot suppress noise sufficiently without blurring edges and other image discontinuities.

The *wavelet transform* [6–8] as a tool for a *multiresolution analysis* facilitates noise removal from signals and images without excessive smoothing of actual discontinuities. For denoising *still* images, numerous wavelet based filtering techniques were proposed, e.g., [9–23], and proved their advantages over standard single-resolution filters. In contrast to still images, wavelet based denoising for *video* is much less studied so far, due to real-time processing constraints. However, wavelet-based video compression and coding techniques are becoming increasingly popular [24] and multimedia trends indicate that the next generation of videocoders will be based on wavelets. Given these trends it is to be expected that wavelet-based video denoising will soon be implementable in real-time as well.

The primary goal of this study is to investigate how big an improvement can be expected from using wavelet based video denoising instead of standard spatio-temporal filters. Despite the importance of this issue, to our knowledge, such studies did not appear in literature yet. Here we first introduce a new spatially adaptive wavelet domain image denoising method. Like in [10, 16], we use a measurement computed from a local window around each coefficient in order to adapt the estimator to the spatial context in the image. We call this measurement *local spatial activity indicator* (LSAI). The novelty of our approach is that we combine the coefficient value, the value of LSAI and the global coefficient histogram to compute the *generalized likelihood ratio* and to estimate accordingly the probability that a given coefficient represents a *signal of interest*, where the notion of the signal of interest is related to the noise level. The proposed method is an extension of our work [19], in the sense that the empirical parameter estimation from [19] is now replaced by analytical expressions.

We analyze the use of the proposed wavelet domain denoising method within the following two scenarios: (1) denoising individual video frames in the wavelet domain and (2) combining wavelet based image denoising and temporal filtering. For comparison, as representatives of single resolution spatio-temporal filters we use a well established 3D rational filter of [1], and the recently proposed 3D KNN (K-Nearest Neighbors) filter of [5].

Our results demonstrate that even though wavelet based denoising of individual frames can suppress noise remarkably well in each frame, the visual quality of the resulting video is often not satisfactory. This is due to the fact that the residual noise as well as inevitable degradations and denoising artifacts differ from frame to frame, which causes an unpleasant visual effect. To overcome this problem, we combine wavelet domain denoising with simple temporal filtering. Our temporal filter is based on a simple pixel-based motion detector and on selective recursive time averaging of the spatially filtered frames. At the positions where motion is detected, the recursive filter resets, preventing in this way edge blurring. The new 3D filter yields a better visual quality and much higher signal-to-noise ratio than the single-resolution 3D filters from [1] and [5].

The paper is organized as follows. In Section 2, the wavelet transform and wavelet based denoising are reviewed. In Section 3, the proposed two dimensional (2D) wavelet denoising method is described and its performance on video is demonstrated. In Section 4, a new 3D filter, which combines wavelet denoising with temporal filtering is described and its results are presented and discussed. The concluding remarks are in Section 5.

2 Wavelet domain noise filtering

Here we review briefly the wavelet decomposition and its use in noise filtering. For a comprehensive treatment of wavelets see [6–8].

2.1 The Discrete Wavelet Transform (DWT)

The standard (bi-)orthogonal discrete wavelet transform can be seen as a *filter bank* algorithm iterated on the lowpass output [7]. A filter bank is a pair of lowpass and highpass filters followed by downsampling by two. The lowpass filtering produces an *approximation* of the signal, which is expressed by the *scaling* coefficients, while the highpass filtering reveals the *details* (i.e., the differences between two successive approximations) that are expressed by the *wavelet* coefficients. At the reconstruction, the scaling and the wavelet coefficients are first up-sampled (by introducing a zero between each two samples) and then filtered with a lowpass and a highpass filter, respectively, followed by summation of the filtered outputs. If the wavelet transform is *orthogonal*, the reconstruction highpass and the lowpass filter coefficients are simply the mirrored versions of their counterparts at the decomposition stage.

The conventional *separable* two-dimensional (2D) DWT follows from applying the filter bank algorithm successively to the rows and to the columns of an image.

The above described DWT is criticaly sampled (non-redundant). It is well known that the noise suppression performance of each denoising method improves when it is implemented in a *redundant* representation instead of using the critically sampled one. In this respect, two approaches are common: (1) *Cycle spining* [12]: apply the orthogonal DWT to several cyclically shifted image versions and average over unshifted denoising results and (2) Denoising in a *non-decimated* (*stationary*) wavelet representation, which is computed with the algorithm *à trous* [7].

2.2 Denoising by wavelet shrinkage

Assume the input image $\mathbf{f} = [\mathbf{f_1}, ..., \mathbf{f_n}]$ is contaminated with additive white Gaussian noise of zero mean and variance σ^2. Due to linearity of the wavelet transform, the noise remains additive in the transform domain as well

$$w_i = y_i + \epsilon_i, \quad i = 1, ..., n \tag{1}$$

where y_i are unknown noise-free wavelet coefficients and ϵ_i are noise contributions. If the wavelet transform is orthogonal, ϵ_i are independent identically distributed (i.i.d.) normal random variables $\epsilon_i \sim N(0, \sigma^2)$. In case where the input standard deviation σ is not known, one usually estimates it as the median absolute deviation of the highest-frequency subband coefficients divided by 0.6745 [14].

Regardless of the type of the employed discrete wavelet transform (e.g., critically sampled or non-decimated), noise reduction is commonly done by *wavelet shrinkage*: the magnitude of each coefficient is reduced by a given amount depending on the noise level and depending on how likely it is that a given coefficient represents an actual discontinuity. A common shrinkage approach is *thresholding* [14], which sets the wavelet coefficients with "small" magnitudes to zero while keeping ("hard-thresholding") or shrinking in magnitude ("soft-thresholding") the remaining ones.

As an alternative to somewhat *ad-hoc* thresholding rules, one can derive the wavelet shrinkage estimators adapted to the statistical properties of images using *Bayesian* approaches [9,11,21,23]. Moreover, advanced wavelet shrinkage methods make also use of image context [10], inter-scale dependencies [13, 15, 22] or intrascale (spatial) correlations [16–18] between image wavelet coefficients. We shall introduce a new spatially adaptive Bayesian wavelet shrinkage approach that is adapted to the global subband statistics and to the local spatial context.

For natural noise-free images, the wavelet coefficient histograms in each subband are typically long-tailed, sharply peaked at zero and are commonly modelled by the generalized Laplacian (also called generalized Gaussian) density [7, 10, 21]

$$p(y) = \frac{\nu}{2s\Gamma(\frac{1}{\nu})} \exp(-|y/s|^{\nu}), \quad s, \nu > 0, \tag{2}$$

where $\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt$ is the Gamma function. For natural images, the shape parameter ν is typically $\nu \in [0, 1]$.

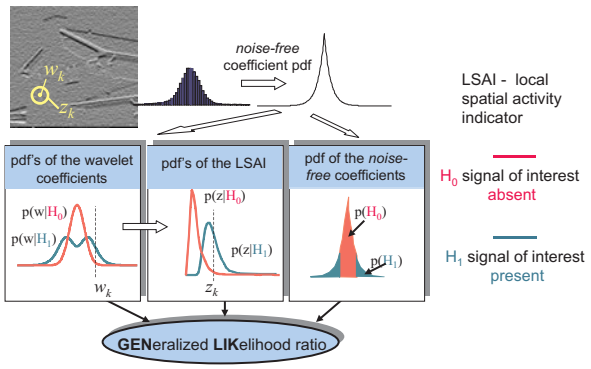

Figure 1. An illustration of the proposed denoising method, where pdf denotes the probability density function.

In case of additive white Gaussian noise, the model parameters ν and s are estimated from the noisy coefficient histogram using the following equations [21]

$$\sigma_w^2 = \sigma^2 + \frac{s^2\Gamma(\frac{3}{\nu})}{\Gamma(\frac{1}{\nu})}, \quad m_{4,w} = 3\sigma^4 + \frac{6\sigma^2 s^2 \Gamma(\frac{3}{\nu})}{\Gamma(\frac{1}{\nu})} + \frac{s^4\Gamma(\frac{5}{\nu})}{\Gamma(\frac{1}{\nu})}, \tag{3}$$

where σ_w^2 and $m_{4,w}$ denote the variance and the fourth moment of the noisy histogram, respectively. A special case in the family (2) with $\nu = 1$, called Laplacian density, is often used due to analytical tractability. The scale parameter is then estimated as

$$s = [0.5(\sigma_w^2 - \sigma^2)]^{1/2}. \tag{4}$$

In our experiments, this simplification usually does not produce a significant degradation in performance.

In the following, the term *density* is used for the probability density function. Also, as it is usual in the related literature, the density $p(y)$ is called "prior (model) for y".

3 The proposed 2D wavelet denoising

In this Section, we investigate two dimensional wavelet based denoising of video sequences.

3.1 The proposed denoising algorithm

Here proposed denoising method is related to our previous work on still image denoising [18, 19], which is described in more detail in [20], where the proposed approach is also shown to compare well with the state-of-the-art in the field. The proposed method estimates how probable it is that the coefficient represents a "signal of interest" given

- the wavelet coefficient value w_l,
- the locally averaged coefficient magnitude in a small window $\delta(l)$: $z_l = \sum_{k\in\delta(l)} |w_k|$ and
- the global statistical distribution of the coefficients in a given subband.

We use the locally averaged magnitude z_l as a *local spatial activity indicator* (LSAI). It is well known that large-magnitude coefficients of noise free images tend to occur near each other within subbands. Many existing methods, e.g., [10, 16], use some form of a local measurement (such as the locally averaged magnitude or local variance) in order to adapt the estimator to the spatial context in the image. Our approach differs in the sense that we also employ the conditional probability density functions of the LSAI given the presence and given the absence of a "signal of interest".

Starting from the following two hypotheses H_1 - "signal of interest is present" and H_0 - "signal of interest is absent", we proposed in [19] the following shrinkage estimator

$$\hat{y}_l = \frac{\rho\xi_l\eta_l}{1+\rho\xi_l\eta_l} w_l, \tag{5}$$

where

$$\rho = \frac{P(H_1)}{P(H_0)}, \quad \xi_l = \frac{p(w_l|H_1)}{p(w_l|H_0)} \quad \text{and} \quad \eta_l = \frac{p(z_l|H_1)}{p(z_l|H_0)}. \tag{6}$$

While in our related method [19] ρ, ξ_l and η_l were estimated empirically from the image using a preliminary coefficient classification, here we derive these expressions analytically.

Starting from the generalized Laplacian prior (2) and defining the hypotheses H_0 and H_1 explicitly as

$$H_0 : |y| < T \quad \text{and} \quad \mathrm{H}_1 : |\mathrm{y}| \geq \mathrm{T}, \tag{7}$$

where the threshold T defines the notion of the signal of interest, we have

$$p(y|H_0) = \begin{cases} \propto \exp(-|y/s|^\nu) & \text{if } \mathrm{y} < \mathrm{T}, \\ 0 & \text{if } \mathrm{y} \geq \mathrm{T}, \end{cases} \tag{8}$$

$$p(y|H_1) = \begin{cases} 0 & \text{if } \mathrm{y} < \mathrm{T}, \\ \propto \exp(-|y/s|^\nu) & \text{if } \mathrm{y} \geq \mathrm{T}. \end{cases} \tag{9}$$

Since the additive white noise model $w = y + \epsilon$, with $\epsilon \sim N(0, \sigma^2)$ is assumed, the densities of noisy coefficients $p(w|H_0)$ and $p(w|H_1)$ result from convolving the normal density $N(0, \sigma^2)$ with $p(y|H_0)$ and with $p(y|H_1)$, respectively. Further on, for the model (7), $P(H_1)$ is clearly

$$P(H_1) = \int_{-\infty}^{-T} p(\beta)d\beta + \int_{T}^{\infty} p(\beta)d\beta. \tag{10}$$

Based on this, for the prior (2) we derive

$$\rho = \frac{P(H_1)}{P(H_0)} = \frac{1 - \Gamma_{inc}\left((\lambda T)^\nu, \frac{1}{\nu}\right)}{\Gamma_{inc}\left((\lambda T)^\nu, \frac{1}{\nu}\right)}, \tag{11}$$

where $\Gamma_{inc}(x,a) = \frac{1}{\Gamma(a)}\int_0^x t^{a-1}e^{-t}dt$ is the *incomplete gamma function*. For the Laplacian prior, with $\nu = 1$, the previous expression reduces to

$$\rho = \frac{P(H_1)}{P(H_0)} = \frac{\exp(-\lambda T)}{1-\exp(-\lambda T)}. \quad (12)$$

Accurate modelling of the conditional densities $p(z_l|H_0)$ and $p(z_l|H_1)$ is difficult. The simplest solution is to assume that the coefficients are conditionally independent given H_0 and H_1 and that all the coefficients within the small window are equally distributed. With these simplifications, the densities of the locally averaged coefficient magnitudes z_l over the window $W\mathrm{x}\mathrm{W} = \mathrm{N}$ are modelled by N convolutions:

$$p(Nz_l|H_1) = \underbrace{p(m_l|H_1) \star p(m_l|H_1) \star ... \star p(m_l|H_1)}_{N},$$
$$p(Nz_l|H_0) = \underbrace{p(m_l|H_0) \star p(m_l|H_0) \star ... \star p(m_l|H_0)}_{N}, \quad (13)$$

where m_l denotes the coefficient magnitude $m_l = |w_l|$, with $p(m_l|H_{0,1}) = 2p(w_l|H_{0,1})$, $m_l > 0$. The complete method is illustrated in Fig. 1, where the product $\rho\xi_l\eta_l$ from (5) is denoted as *generalized likelihood ratio*, due to which we call this method *GenLik*. The figure illustrates that all the required expressions in the proposed method are estimated directly from the noisy coefficient histogram and that the method estimates how probable it is that a given coefficient presents useful information, based on its value, based on a measurement from the local surrounding and based on the global subband statistics.

A more extensive analysis of the proposed method is in [20]. We show there that the optimal value of the threshold in terms of the mean squared error is $T = \sigma$, and we also show that on standard test images the proposed method yields results that are among the best reported in the wavelet denoising literature. For most test images and noise levels, the optimal window size is 5x5.

3.2 Implementation details and complexity

The computation of the convolutions such as (13) can be time consuming. We avoid this problem in our practical application by quantizing appropriately all the involved densities. Since the densities of the wavelet coefficients are symmetrical around zero we have $\xi_l = p(\omega_l|H_1)/p(\omega_l|H_0)$, and in practice we quantize it to

$$\xi(\omega_l) = p(\lfloor\omega_l\rfloor|H_1)/p(\lfloor\omega_l\rfloor|H_0), \quad (14)$$

where $\lfloor\cdot\rfloor$ denotes the integer operation. We store for each subband a vector $\xi = [\xi(\mathbf{0}), ..., \xi(\mathbf{MAX})]$, where MAX is the maximum expected magnitude of the wavelet coefficient in a given subband. Similarly, we also calculate and store the vector $\eta = [\eta(\mathbf{0}), ..., \eta(\mathbf{MAX})]$. We also normalize the input image intensity such that the maximum magnitudes of the wavelet coefficients are relatively small, because then the convolutions in (13) are computed over short vectors.

Regarding the computation complexity, we need to calculate *per subband* the expressions (11) and (3) *or* (12) and (4) depending on the chosen prior and the vectors $\xi = [\xi(\mathbf{0}), ..., \xi(\mathbf{MAX})]$ and $\eta = [\eta(\mathbf{0}), ..., \eta(\mathbf{MAX})]$. The numerical operations *per coefficient* are three multiplications, one division and one addition and calculating the locally averaged magnitude in a 5x5 square window.

In practice, we apply the method in the non-decimated wavelet transform, with four decomposition levels. The employed wavelet was *symmlet* [6] with 8 vanishing moments. Here we used the generalized Laplacian prior (2). In our experiments, the simplified algorithm with $\nu = 1$ usually yielded 0.1dB - 0.3dB worse signal to ratio.

Our Matlab implementation takes several seconds for 256x256 images on the Pentium processor of 1 GHz.

3.3 Results of the 2D wavelet filter

We tested the proposed filter on three grey scale sequences: "salesman", "tennis" and "flowers". corrupted with artificial additive white Gaussian noise. As representatives of single-resolution filters, we used the 3D rational filter of [1] and the 3D KNN filter of [5].

The results from Fig. 2 correspond to the noise standard deviation $\sigma = 10$. These results demonstrate that the wavelet filter yields surprisingly high peak signal to noise ratio (PSNR) values with respect to other two filters even though it does not use filtering in time. For example, for the "flowers" sequence the advantage over the 3D KNN filter [5] is 2dB during the whole sequence. The 2D wavelet filter yields also much higher PSNR with respect to the 3D KNN filter for most of the frames in other two sequences. For higher noise levels the advantage of the wavelet filter is even more pronounced.

However, the visual quality of the denoised video sequences is often not as good as the high PSNR values suggest. We believe that the main reason for this is that the residual noise as well as inevitable degradations and denoising artifacts differ from frame to frame. This causes an unpleasant visual effect, which can be described as a kind of background "flickering". To overcome this problem the filtering in time is necessary.

4 Wavelet domain and temporal filtering

Here we extend the proposed filtering method to include the time dimension as well. Our focus was on solutions that do not increase the complexity significantly.

4.1 The proposed 3D denoising algorithm

Based on experiments, we found that it is effective to combine the 2D wavelet denoising with temporal filtering of the denoised frames as we explain next. Let $\mathbf{f^k}$ denote the k-th frame of a *noise-free* video sequence and $\mathbf{d^k} = \mathbf{f^k} + \mathbf{n^k}$ the corresponding noisy frame, where $\mathbf{n^k}$ is the noise field. Further on, let

$$\hat{\mathbf{f}}^{2D,k} = [\hat{f}_1^{2D,k}, ..., \hat{f}_L^{2D,k}] \tag{15}$$

denote the k-th 2D wavelet *denoised frame* of the sequence. Our intention is to improve $\hat{\mathbf{f}}^{2D,k}$ based on a motion information and on temporal filtering. Towards this end, we define the *motion field* $\mathbf{m^k} = [\mathbf{m_1^k}, ..., \mathbf{m_L^k}]$ of the k-th frame with respect to the previous frame as follows

- $m_l^k = 0$ if there is *no* (significant) motion from the frame $k-1$ to the frame k at the spatial position l, meaning that $f_l^k \approx f_l^{k-1}$.
- $m_l^k = 1$ if there *is* motion from the frame $k-1$ to the frame k at the spatial position l, meaning that f_l^k and f_l^{k-1} differ significantly.

We estimate this motion field from the denoised frames as

$$\hat{m}_l^k = \begin{cases} 0, & \text{if } |\hat{f}_l^{2D,k} - \hat{f}_l^{3D,k-1}| < T, \\ 1, & \text{otherwise,} \end{cases} \tag{16}$$

where T is a threshold that we leave as a free (to be optimized) parameter. At the spatial positions where no motion was detected, we apply recursive time averaging, yielding the final 3D filtered pixel intensity:

$$\hat{f}_l^{3D,k} = \begin{cases} \alpha \hat{f}_l^{2D,k} + (1-\alpha)\hat{f}_l^{3D,k-1}, & m_l^k = 0, \\ \hat{f}_l^{2D,k}, & \text{otherwise,} \end{cases} \tag{17}$$

where $0 \leq \alpha \leq 1$. Note that this recursive filter accumulates and averages the pixel intensities at a given position from *all the previous* frames if the motion was not present at that position. The detection of a motion resets the filter.

4.2 Practical implementation and complexity

The increase of complexity with respect to the 2D filtering scheme is not significant. Additional operations are: computing the difference between the pixel intensities at the same position in two consecutive frames and computing a weighted average of the two corresponding intensities if their absolute difference is above a fixed threshold.

In our experiments, we set the weighting parameter α to a constant value $\alpha = 0.5$ and investigated the influence of the threshold T on the resulting PSNR value and on the visual quality. The PSNR values for the three test sequences and for different values of the threshold T are shown in Fig. 3. In case of the salesman sequence, the increase of the value T up to 20 yielded an increasing improvement in the PSNR. In this sequence the background is static and the advantage of using the temporal filtering is obvious. However, in the other two tested sequences, our extension to temporal filtering does not always produce a better PSNR as compared to the 2D wavelet filter. One can see from Fig. 3 that in (the parts of) the sequences with rapid movements the resulting PSNR is worse than in case of the 2D wavelet filtering only. Nevertheless, the 3D filtering yielded an improved *visual quality* of all the tested sequences.

4.3 Results of the 3D filter and discussion

Fig. 2 shows PSNR for the three tested sequences, for $\sigma = 10$, filtered by the new 3D filter, the 2D wavelet filter, the 3D rational filter [1], and the 3D KNN filter [5]. The new 3D filter yields an improvement of 1 to 2dB, with respect to the 3D KNN filter [5], and the improvement of up to 8dB with respect to the 3D rational filter [1]. The new 3D filter outperforms all the other tested ones in terms of visual quality as well (see Fig. 4). The advantage of the new filter is even greater when viewing the video sequence instead of comparing the individual frames. It should be also noted that the 3D rational filter is the fastest and the new filter is the most complex among the tested ones.

Here the results were presented only for $\sigma = 10$. Our initial experiments, indicate that for higher noise levels the gain of the new filter with respect to the single-resolution ones is even larger. However, further research is needed to automatize the adjustment of the parameters of the temporal filtering part for different noise levels.

5 Conclusion

In this paper, we studied an application of wavelet based image denosinig to video sequences. First, we investigated denoising of the individual frames by means of a spatially adaptive 2D wavelet domain filtering. The results show that 2D filtering in the wavelet domain outperforms most of the existing 3D (spatio-temporal) single resolution filters in PSNR. However, our results also demonstrated that the 2D wavelet denoised sequences show an unpleasant "flickering" artifact due to the lack of the filtering in time.

To improve the result of the 2D wavelet filtering, we combined it with a temporal filter, which combines a pixel-based motion detector and a recursive time-averaging. The results show that this combination of the 2D wavelet domain and temporal filtering outperforms recent 3D single resolution methods for video denosing in terms of quantitative performance measures and in terms of visual quality.

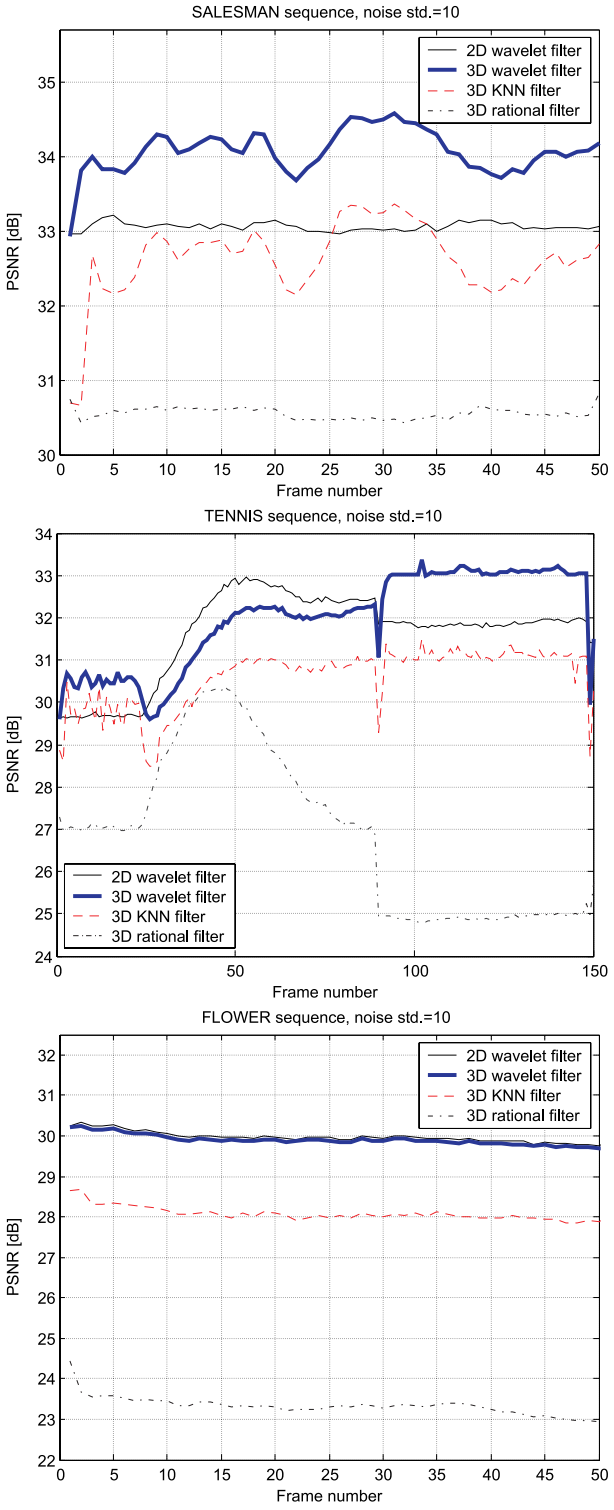

Figure 2. The quantitative performance of the proposed 2D and 3D filters in comparison with the 3D KNN filter [5] and the 3D rational filter [1].

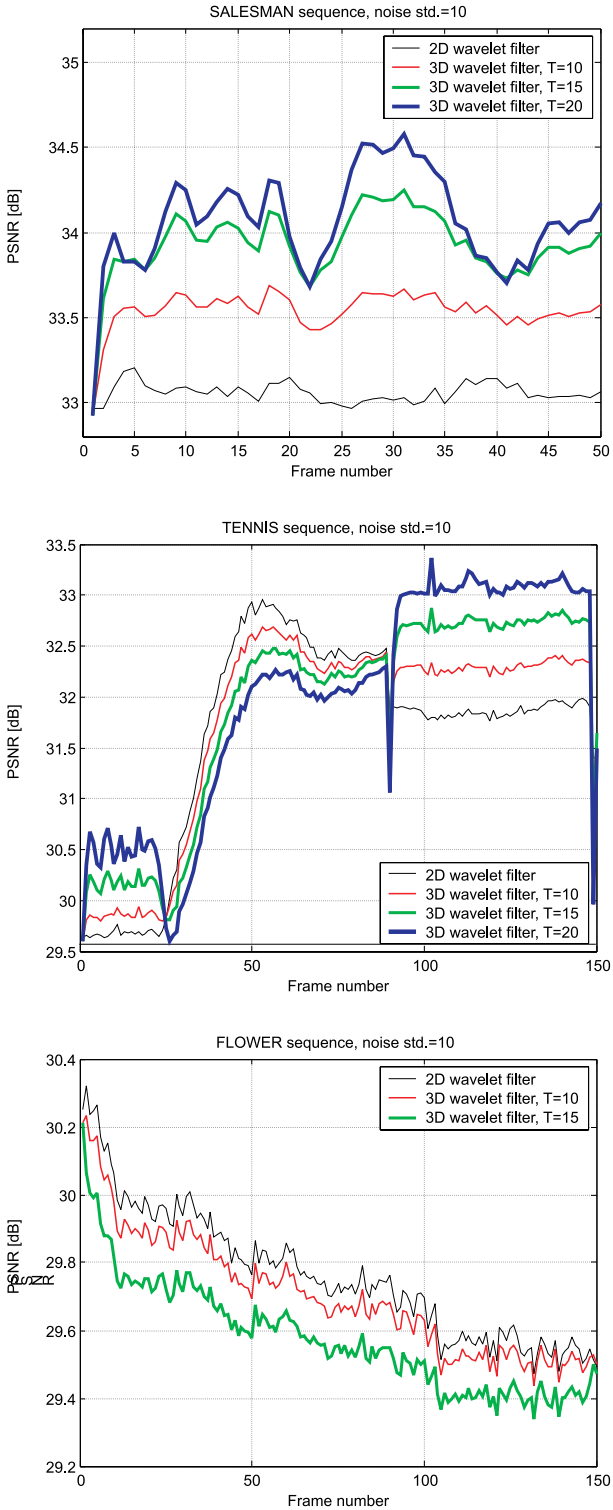

Figure 3. The influence of the parameter T on the performance of the proposed 3D filter, where $\alpha = 0.5$.

(a) Original

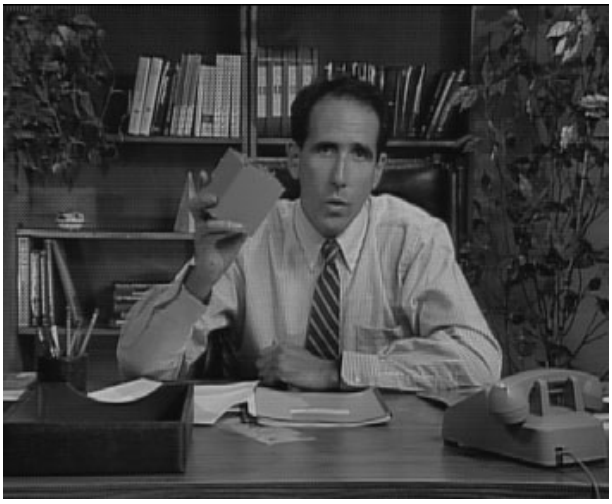

(b) Corrupted, PSNR=28.2dB

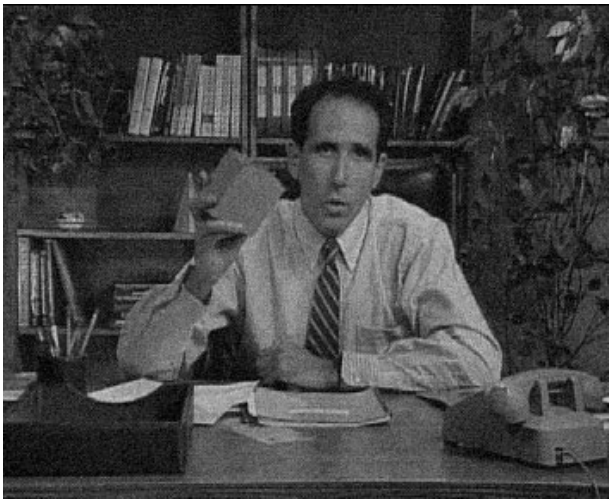

(c) Rational filter [1], PSNR=30.4dB

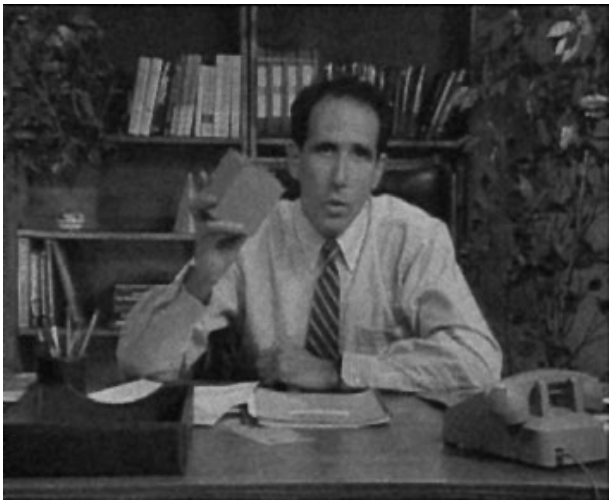

(d) 3D KNN filter [5], PSNR=33.3dB

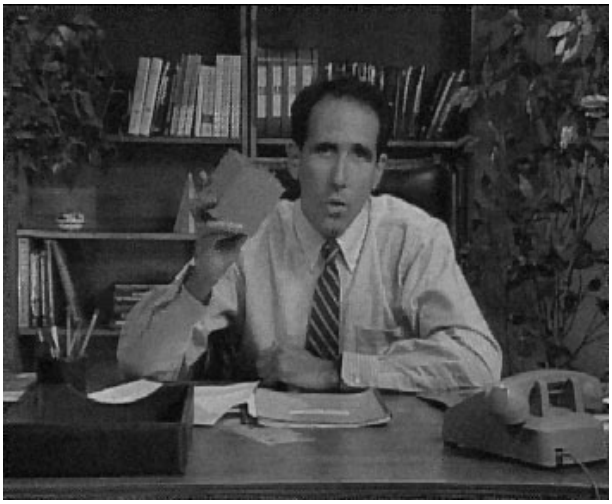

(e) 2D Wavelet filter, PSNR=33.0dB

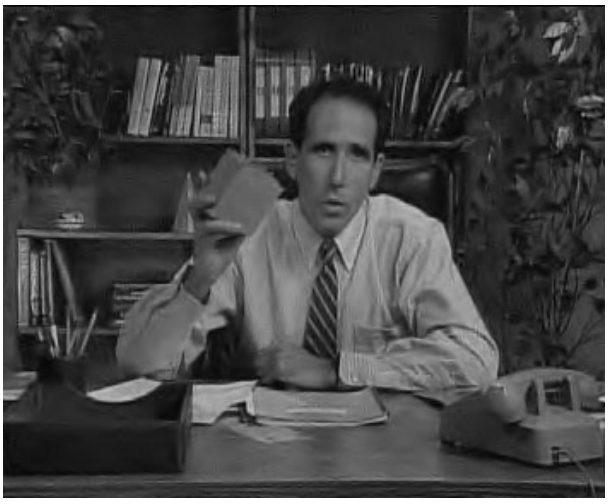

(f) Combined 2D wavelet and temporal filter, PSNR=34.6dB

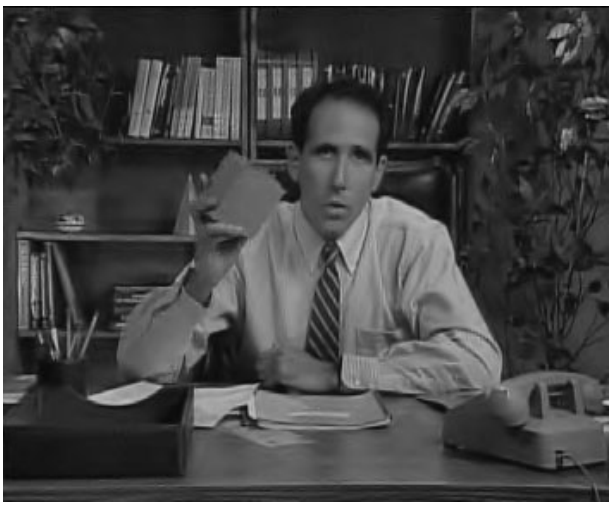

Figure 4. Visual results for the Salesman sequence with noise standard deviation 10.

References

[1] F. Cocchia, S. Carrato, and G. Ramponi, "Design and real-time implementation of a 3-D rational filter for edge preserving smoothing," *IEEE Trans. Consumer Electron.*, vol. 43, no. 4, pp. 1291–1300, Nov. 1997.

[2] G. De Haan, G.T. Kwaaitaal-Spassova, M. Larragy, and A.O. Olukayode, "Memory integrated noise reduction ic for television," *IEEE Trans. Consumer Electron.*, vol. 42, no. 2, pp. 175–180, 1996.

[3] K. Jostschulte, A. Amer, M. Schu, and H. Schroder, "Perception adaptive temporal tv-noise reduction using contour preserving prefilter techniques," *IEEE Trans. Consumer Electron.*, vol. 44, no. 3, pp. 1091–1096, 1998.

[4] A.O. Olukayode and G.T. Kwaaitaal-Spassova, "An algorithm for integrated noise reduction and sharpness enhacement," *IEEE Trans. Consumer Electron.*, vol. 46, no. 3, pp. 474–480, 2000.

[5] V. Zlokolica, W. Philips, and D. Van De Ville, "A new non-linear filter for video processing," in *IEEE Benelux Signal Processing Symposium*, Mar. 2002, pp. 221–224.

[6] I. Daubechies, *Ten Lectures on Wavelets*, Philadelphia: SIAM, 1992.

[7] S. Mallat, *A wavelet tour of signal processing*, Academic Press, London, 1998.

[8] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*, 1995.

[9] F. Abramovich, T. Sapatinas, and B.W. Silverman, "Wavelet thresholding via a bayesian approach," *J. of the Royal Statist. Society B*, vol. 60, pp. 725–749, 1998.

[10] S.G. Chang, B. Yu, and M. Vetterli, "Spatially adaptive wavelet thresholding with context modeling for image denoising," *IEEE Trans. Image Proc.*, vol. 9, no. 9, pp. 1522–1531, Sept. 2000.

[11] H.A. Chipman, E.D. Kolaczyk, and R.E. McCulloch, "Adaptive bayesian wavelet shrinkage," *J. of the Amer. Statist. Assoc.*, , no. 92, pp. 1413–1421, 1997.

[12] R.R. Coifman and D.L. Donoho, "Translation-invariant denoising," *Wavelets and Statistics*, pp. 125–150, Sept. 1995.

[13] M.S. Crouse, R.D. Nowak, and R.G. Baranuik, "Wavelet-based statistical signal processing using hidden markov models," *IEEE Trans. Signal Proc.*, vol. 46, pp. 886–902, 1998.

[14] D.L. Donoho and I.M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 8, pp. 425–455, 1994.

[15] S Mallat and W.L. Hwang, "Singularity detection and processing with wavelets," *IEEE Trans. Information Theory*, vol. 38, no. 8, pp. 617–643, Mar. 1992.

[16] M.K. Mihcak, I. Kozintsev, K. Ramchandran, and P. Moulin, "Low-complexity image denoising based on statistical modeling of wavelet coefficients," *IEEE Signal Proc. Lett.*, vol. 6, no. 12, pp. 300–303, Dec. 1999.

[17] M. Jansen and A. Bultheel, "Empirical bayes approach to improve wavelet thresholding for image noise reduction," *J. of the Amer. Statist. Assoc.*, vol. 96, pp. 223–242, 2001.

[18] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy, "A joint inter- and intrascale statistical model for wavelet based bayesian image denoising," *IEEE Trans. Image Proc*, vol. 11, no. 5, pp. 545–557, May 2002.

[19] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy, "A versatile wavelet domain noise filtration technique for medical imaging," *IEEE Trans. Medical Imaging, in press*, 2003.

[20] A. Pizurica, *Image Denoising Using Wavelets and Spatial Context Modeling*, Ph.D. thesis, Ghent University, Belgium, 2002.

[21] E.P. Simoncelli and E.H. Adelson, "Noise removal via bayesian wavelet coring," in *Proc. IEEE Internat. Conf. Image Proc. ICIP, Lausanne, Switzerland*, 1996.

[22] J.K. Romberg, H. Choi, and R.G. Baraniuk, "Bayesian tree structured image modeling using wavelet-domain hidden markov models," *IEEE Trans. Image Proc.*, vol. 10, no. 7, pp. 1056–1068, July 2001.

[23] B. Vidakovic, "Nonlinear wavelet shrinkage with bayes rules and bayes factors," *J. of the American Statistical Association*, vol. 93, pp. 173–179, 1998.

[24] K. Shen and E.J. Delp, "Wavelet based rate scalable video compression," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 9, no. 1, pp. 109–122, 1999.