SPATIO-TEMPORAL POINTWISE ADAPTIVE DENOISING OF VIDEO: 3D NON-PARAMETRIC REGRESSION APPROACH

Chiara Ercole[†]*, *Alessandro Foi*[‡]*, *Vladimir Katkovnik* and *Karen Egiazarian**

*Institute of Signal Processing, Tampere University of Technology, Tampere, Finland
[†]Dipartimento di Elettronica Applicata, Università degli Studi "Roma TRE", Roma, Italy
[‡]Dipartimento di Matematica, Politecnico di Milano, Milano, Italy

ABSTRACT

An adaptive spatio-temporal algorithm for video denoising is presented. The local polynomial approximation (*LPA*) is exploited in order to design $3D$ directional filtering kernels. For each specified direction in the $3D$ space-time domain, an "optimal" scale (size of the kernel's support) is selected using the intersection of confidence intervals (*ICI*) rule. In this way a pointwise adaptive spatio-temporal estimator is constructed. Experimental results show an advanced performance of the proposed method with a good noise attenuation and nearly perfect edges and change-point preservation.

1. INTRODUCTION

Noise is usually present in a video sequence because of transmission over noisy channels or acquisition with poor quality devices. Besides unpleasant visual effects, that sometimes can seriously compromise the perceiving and interpretation of the content, the main and most affecting problem is the degradation of the result of further processing such as video compression, segmentation, motion estimation. When addressing the problem of restoring a corrupted video sequence, the goal is to find a denoising scheme that can guarantee good performances of video processing algorithms and also a satisfactory visual quality. Since dealing with $3D$ data set, a good approach should take into account both spatial and temporal dimensions, so to exploit the spatial and temporal correlation in the video. Nevertheless, constraints of real-time implementation make all the efforts go in the direction of simple separable filters. Although they reach the required computational speed, these filters cannot suppress noise sufficiently well without introducing disturbing artifacts such as blurry edges or smoothing away salient characteristics like details and texture. The loss of these elements can heavily affect not only the further video processing, but also the subjective perception, since they encode a great amount of visual information contained in image sequences.

Recently, a novel accurate technique was developed [7]. It is based on a directional discrete approximation of the image intensity using the nonparametric regression technique called *LPA*. The neighborhood of this approximation is locally shape and size adaptive. The intersection of confidence intervals (*ICI*) rule [3, 4] is exploited for this adaptation, in order to find the largest local vicinity of the point of estimation where the estimate fits best to the data [5, 6, 7].

It is assumed in the algorithm proposed in [7] that the local neighborhood is a starshaped set, which can be approximated by some sectorial segmentation with K non-overlapping sectors. These estimators are equipped with univariate scale parameters defining the size of the supports in the sector. The *ICI* rule is applied K times, once for each sector, in order to find the optimal pointwise adaptive scales for each sector's estimates. These are then combined into the final one using an adaptive convex fusing.

In this paper we develop this idea and algorithm to $3D$, where time is the third dimension completing the $2D$ space of the image frame. In this way, we build an anisotropic $3D$ denoising filter for video.

Although this general approach has been extended also to a more general class of noise processes and non-linear estimators [8], all the results presented in this paper concern the Gaussian observation model

$$z(x) = y(x) + n(x), \quad x \in \mathbb{R}^3,$$

where y is a video intensity function and $n \sim \mathcal{N}(0, \sigma^2)$ is additive white Gaussian noise. Here $x = (x_1, x_2, x_3)$, with x_1, x_2 being the spatial coordinates and the third coordinate x_3 being interpreted as time or frame number.

2. PARTITION OF $3D$ NEIGHBORHOOD AND FILTERING

A partition of the neighborhood in the $3D$ space can be done in different ways. Here we discuss the partition based on spherical coordinate system. Figure 1 illustrates the meaning of the spherical angular coordinates θ and φ with respect

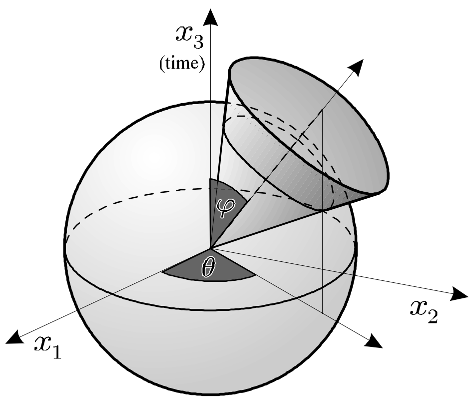

Fig. 1. Spherical coordinates: the angle θ lies in the spatial frame plane while φ is the angle measured from the temporal axis.

to the cardinal spatial and temporal coordinates: θ is the angular coordinate of a polar system in the frame plane, while φ is the temporal angular coordinate. Thus, purely temporal directions are obtained for $\varphi = 0 \bmod \pi$; purely spatial directions are obtained for $\varphi = \pi/2 \bmod \pi$, $n \in \mathbb{Z}$. So, referring to the frame plane, direction along axis x_1 is obtained for $\varphi = \pi/2$ and $\theta = 0$, while direction parallel to axis x_2 is obtained for $\varphi = \pi/2$ and $\theta = \pi/2$.

2.1. Spherical partition

Let $\partial B^3 = \{x \in \mathbb{R}^3 : \|x\| = 1\}$ be the unit sphere. Introduce a covering of ∂B^3 and a finite family $\{D_{\theta_i,\varphi_i}\}_{i=1,\dots,K}$ of non-overlapping contractible bodies (in the sphere topology) $D_{\theta_i,\varphi_i} \subset \partial B^3$ whose baricenters have spherical angular components (θ_i, φ_i). For any given $h \in \mathbb{R}^+$, $S^h_{\theta_i,\varphi_i} = \bigcup_{0\leq\alpha\leq h} \alpha D_{\theta_i,\varphi_i}$ are the corresponding $3D$ conical bodies the union of which, over all directions, fills the ball $B^3_h = \{x \in \mathbb{R}^3 : \|x\| \leq h\}$. These cones correspond to angular sectors having their vertex in the origin and oriented as defined by the angles (θ_i, φ_i), resulting from the union of the α-dilated D_{θ_i,φ_i} for every $0 \leq \alpha \leq h$.

Figure 2 illustrates such a partition. Practically, we do not need a covering of the whole sphere, and the considered partition can be restricted to a few narrow cones pointing at different directions and covering only a part of the $3D$ neighborhood for the considered point.

2.2. Multi-directional filtering

Let $g_{h,i} \triangleq g_{h,(\theta_i,\varphi_i)}$ be compactly supported kernels such that $\operatorname{supp} g_{h,i} = S^h_{\theta_i,\varphi_i}$ for all values of the scale parameter h. Then, the introduced anisotropic estimator has the

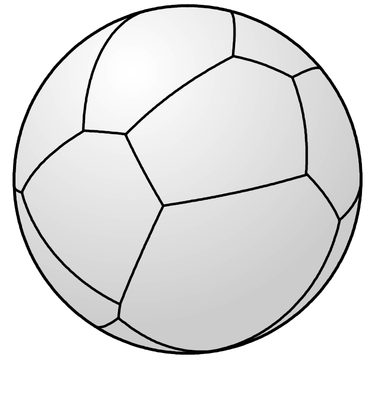

Fig. 2. Example of a covering of the sphere with non-overlapping bodies (Voronoi tiling of the sphere).

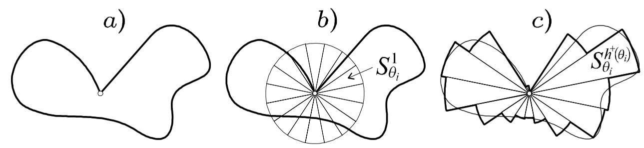

Fig. 3. A sectorial approximation of the estimation neighborhood: a) the best estimation set, b) the unit ball segmentation, c) sectorial approximation with cones of adaptive length.

following generic form

$$\hat{y}(x) = \sum\nolimits_i \lambda_i \hat{y}_{h,i}(x), \quad \hat{y}_{h,i}(x) = (g_{h,i} \circledast z)(x), \quad (1)$$

where $\lambda_i \geq 0$, $\sum_i \lambda_i = 1$, and the directional kernel $g_{h,i}(x)$ satisfies vanishing moment conditions:

$$(g_{h,i} \circledast 1)(0) = 1, \ \ (g_{h,i} \circledast x^t)(0) = 0, \ \ 0 \leq t \leq m, \ |t| \neq 0.$$

Here and in what follows a compact multi-index notation is used. A multi-index t is a *triplet* of non-negative integers t_j, $j = 1, 2, 3$, $t = (t_1, t_2, t_3)$, where $t_j \geq 0$ and $|t|$ is used to denote the length $\sum_{j=1}^{3} t_j$. Then $x^t = x_1^{t_1} x_2^{t_2} x_3^{t_3}$ for $x \in \mathbb{R}^3$, and $0 \leq t \leq m$ means $0 \leq t_j \leq m_j$, $j = 1, 2, 3$.

The $\hat{y}_{h,i}(x)$ in (1) is the directional estimate of $y(x)$ using the observations z from the sector $S^h_{\theta_i,\varphi_i}$. Optimization of h for each sector estimates gives the adaptive scales h_i^+ depending on (θ_i, φ_i). The union of the supports $g_{h_i^+,i}$, $\bigcup_i \operatorname{supp} g_{h_i^+,i}$, can be therefore considered as an approximation of the best neighborhood in which the estimation model fits the data.

Figure 3 illustrates this concept, for the two dimensional case, and shows sequentially: a local best estimation neighborhood, a sectorial segmentation of the unit ball and the sectorial approximation of the local best estimation neighborhood using the adaptive scales h_i^+ defining the length of the corresponding sectors. Varying size sectors enable to get a good approximation of any neighborhood provided that it is starshaped body.

Formula (1) makes clear our basic intentions. We introduce the directional estimates $\hat{y}_{h,i}(x)$, optimize the scalar scale parameter for each of the directions (sectors) and fuse these directional estimates into the final one $\hat{y}(x)$, using the weights λ_i. Two points are of the main importance here. First, we are able to find good approximations of estimation supports which can be of a complex shape. Second, this approximation is composed from the univariate scale optimizations on h, thus the complexity is proportional to the number of sectors.

What follows mainly concerns applied aspects of the approach and includes:

- Design of the discrete directional kernels $g_{h,(\theta_i,\varphi_i)}$;
- Application of the *ICI* rule for the adaptive varying scale selection for each direction;
- Fusing of the directional estimates into the final one using the data-driven weights λ_i;
- Application examples proving a good performance of the presented technique.

3. DIRECTIONAL *LPA* KERNEL DESIGN

Let us start from the standard *LPA* technique. Introduce three-dimensional polynomials $x^k/k! = x_1^{k_1}/k_1! \cdots x_3^{k_3}/k_3!$, $k_1 = 0,\ldots,m_1$, $k_2 = 0,\ldots,m_2$, and $k_3 = 0,\ldots,m_3$. Thus the polynomials $x^k/k!$ have m_1, m_2, m_3 as the maximum powers with respect to the variables x_1, x_2, x_3, respectively. The vector $\phi(x)$ is composed of these polynomials starting from the zero order term 1. The observations z are given on the three-dimensional grid $\{\tilde{x}_s\}$ and the estimates are needed for a desired x. Then, the weighted least squares criteria

$$J_h(C) = \sum_s w_h(x - \tilde{x}_s)(z(\tilde{x}_s) - \bar{y}_h(x - \tilde{x}_s))^2$$

is commonly used to calculate the non-parametric estimation kernels (e.g. [1]). Here: $\bar{y}_h(x) = C^T\phi_h(x)$, $\phi_h(x) = \phi(x/h)$, $h \in \mathbb{R}^+$ is the scale parameter, w is a window function used for localization of the estimates and $w_h(x) = w(x/h)/h^3$. Thus, we produce a fit of the observations z by the model $C^T\phi_h(x)$ with unknown C. According to the idea of the *LPA*, minimizing $J_h(C)$ on C gives the vector $\hat{C}$, whose elements correspond to the estimates of the function $\hat{y}_h^{(0)}(x)$ and its derivatives $\hat{y}_h^{(r)}(x)$, $r = (r_1, r_2, r_3)$, in the form

$$\hat{y}_h^{(r)}(x) = \hat{C}^T\phi^{(r)}(0)(-1)^{|r|}/h^r,\ \phi^{(r)} = \partial^{(r)}\phi.$$

Assuming that the grids $\{\tilde{x}_s\}$ and $\{x\}$ are regular, identical and unrestricted, these estimates can be given in the convolution form (e.g. [5, 6])

$$\begin{aligned} \hat{y}_h^{(r)}(x) &= (z \circledast g_h^{(r)})(x), \qquad (2)\\ g_h^{(r)}(x) &= \frac{(-1)^{|r|}}{h^r} w_h(x)\phi_h^T(x)\Phi_h^{-1}\phi^{(r)}(0),\\ \Phi_h &= \sum_x w_h(x)\phi_h(x)\phi_h^T(x). \end{aligned}$$

where $g_h^{(r)}(x)$ are the estimation kernels.

There are two ways how the directional kernels can be obtained. First, a *basic kernel*, $g_h^{(r)}$ is designed for a main direction, along the axis x_1 ($\varphi = \pi/2$ and $\theta = 0$), and then it is rotated to the desired direction (θ_i, φ_i) in order to obtain $g_{h,(\theta_i,\varphi_i)}^{(r)}$. In practice, for discrete data, this rotation assumes interpolation of the kernel values $g_{h,(\theta_i,\varphi_i)}^{(r)}$ and can be efficient only for sufficiently smooth (not sharp) kernels.

Alternatively, the kernels are specially designed for each desired direction (θ_i, φ_i). This design is not restricted by requirements of smoothness of the kernel and can be applied in order to obtain very sharp kernels. The *LPA* method was modified to the directional *LPA* in order to design this sort of sharp, well directed kernels. It comprises three independent steps. First, the support of the basic window w is selected. This particular support is oriented in some basic direction, say $\varphi_0 = \pi/2$ and $\theta_0 = 0$ (in our $3D$ model, along the purely spatial direction parallel to x_1 coordinate); it is finite, non-symmetric, elongated and well oriented in the main direction. Second, the basic support of w is rotated to the desired direction (θ, φ) so to obtain a support for this new direction. This rotation is here expressed by the $3D$ rotation matrix $U_{\theta,\varphi}$. In practical applications, it can be performed by any discrete rotation method such as nearest neighborhood or linear interpolation. Third, the standard *LPA* procedure is applied on the rotated support in the truly rotated variables $U_{\theta,\varphi}x$. Finally the directional kernel has a form

$$\begin{aligned} g_{h,(\varphi,\theta)}^{(r)}(x) &= \frac{(-1)^{|r|}}{h^r} w_h(u)\phi_h^T(u)\Phi_h^{-1}\phi^{(r)}(0),\\ \Phi_h &= \sum_{x \in X_\theta} w_h(u)\phi_h(u)\phi_h^T(u), \qquad u = U_{\theta,\varphi}x. \end{aligned}$$

What makes this procedure different from any attempt to interpolate the kernels (2) to the desired directions is that the directional *LPA* preserves the normalization and the polynomial smoothness of the kernels (vanishing moment conditions) as well as the directionality of the kernel support. The directional *LPA* kernels enable the estimates to be accurate with respect to smooth polynomial components of the signal, found along a specific direction.

- Unlike many other transforms which start from the continuous domain and then pass to the discrete one, this technique works directly in the discrete spatial domain and is applicable for data of any dimension;
- The designed kernels are truly multivariable, non-separable and anisotropic with arbitrary widths, lengths and orientations;
- The desirable smoothness of the kernels along and across the main direction is enabled by the corresponding vanishing moment conditions;
- The kernel support can be flexibly shaped to any desirable geometry in order to capture geometrical structure and pictorial information. In this way a special design can be done for complex form objects;
- The smoothing and corresponding differentiating directional kernels can be designed.

4. ICI ADAPTIVE SCALE SELECTION ALGORITHM

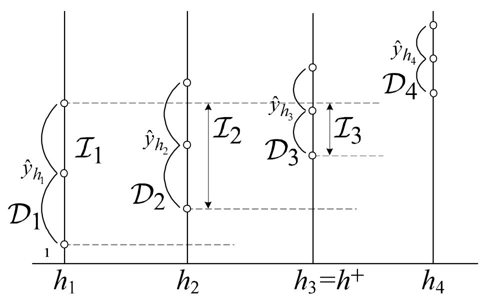

Fig. 4. Intersection of Confidence Intervals (*ICI*) rule.

Let us remind the *ICI* rule technique [6, 5].

Given a finite set of the ordered scales $H = \{h_1 < h_2 < \cdots < h_J\}$ and their corresponding varying scale kernel estimates $\{\hat{y}_{h_j}(x)\}_{j=1}^{J}$, with decreasing standard deviations $\sigma_{\hat{y}_{h_1}} > \cdots > \sigma_{\hat{y}_{h_J}}$, we determine a sequence of confidence intervals

$$\mathcal{D}_j = \left[\hat{y}_{h_j}(x) - \Gamma\sigma_{\hat{y}_{h_j}}, \hat{y}_{h_j}(x) + \Gamma\sigma_{\hat{y}_{h_j}}\right],$$

where $\Gamma > 0$ is a threshold parameter. The *ICI* rule can be stated as follows (see Figure 4):

Consider the interesection of confidence intervals $\mathcal{I}_j = \bigcap_{i=1}^{j} \mathcal{D}_i$ and let j^+ be the largest of the indexes j for which $\mathcal{I}_j$ is non-empty, $\mathcal{I}_{j^+} \neq \varnothing$ and $\mathcal{I}_{j^++1} = \varnothing$. The optimal scale h^+ is defined as $h^+ = h_{j^+}$ and the optimal scale kernel estimate is $\hat{y}_{h^+}(x)$.

This is a procedure for a fixed x. It is produced for all $x \in X$ and in this way we obtain varying adaptive scale $h^+(x)$. The *ICI* procedure requires to know only the estimates for different scales as well as the corresponding variances of these estimates.

5. ADAPTIVE ALGORITHM

In the developed algorithm we treat the scale parameter h as the length size of the directional kernel and apply the *ICI* rule, for data driven selection of its values for each of the directional estimates $\hat{y}_{h,i}(x)$ in (1), i.e. for each direction (θ_i, φ_i), and for each x. It gives the adaptive scales $h_i^+(x)$, which define the length of the supports for different directional estimates. All these supports are then combined to form the adaptive estimation neighborhood. Let $\hat{y}_{h_i^+}(x)$ be the adaptive directional estimate and $\sigma_i^2(x)$ be the variance of this estimate; then these directional sectorial estimates can be fused according to (1) in the final one as follows:

$$\hat{y}(x) = \sum_i \lambda_i \hat{y}_{h_i^+}(x), \quad \lambda_i = \sigma_i^{-2}(x) / \sum_j \sigma_j^{-2}(x) \quad (3)$$

We use a linear fusing of the estimates with the inverse variances of the estimates as weights. This rule can be derived from the maximum likelihood method assuming that the estimates to fuse are Gaussian, unbiased and independent. If the directional kernels' supports are non-overlapping, independency is ensured; however the *LPA* estimates are usually biased.

The final estimate can be written also in the following, equivalent way:

$$\hat{y}(x) = \int g_x^+(x - v) z(v)\, dv, \quad (4)$$

where g_x^+ is the anisotropic adaptive kernel. In the case of formula (3) this kernel is given by $\sum_i \lambda_i^+ g_{h_i^+(x)}$. Observe that if the directional kernels are uniform over their supports and non-overlapping, then g_x^+ is also uniform.

Concerning the algorithm complexity we note that it depends on the fast convolution operations. The calculation of the estimate $\hat{y}_{h,i}(x)$ for a given set of scales h_j, $j = 1, \ldots, J$, is a linear convolution requiring $N_{conv} \sim n \log n$ where n is the size of the signal. This procedure is repeated JK times, where K is the number of the $3D$ sectors in the estimator and J is the number of the used scales h_j.

6. SIMULATIONS

As an illustrative application, we wish to recover the *Akiyo* video sequence y from its noisy observation $z = y + n$, where n is an additive white Gaussian noise with zero mean and σ=20. We implemented the proposed method in the simplest possible way, where the directional kernels are uniform over 1-pixel-width segments oriented along the twenty six directions originating from the center of a cube to the eight vertices, to the middle of the twelve sides, to the center of the six faces. These kernels are *LPA* kernels of order zero, $m = [0, 0, 0]$. The following set of scales was used, $H = \{1, 2, 3, 5, 7, 10\}$. Exploiting the directional nature of the kernel supports, we improve the adaptive scale selection (and thus the signal estimate) using a larger threshold Γ on the purely temporal directions, taking advantage of the high temporal correlation between frames.

When the data are discrete, it is impossible to have non-overlapping supports for the directional kernels, and, in practice, they are all overlapping in the origin voxel. Although in a number of applications ([6, 7, 8, 2]) formula (3) has been used also for origin-overlapping kernels, the larger number of directional estimates makes (3) unsuitable for the proposed algorithm. Thus, in our scenario, we use a slightly different fusing formula:

$$\hat{y}(x)=\frac{\left(\sigma^{-2}z(x)+\sum_i\left(\tilde{\lambda}_i(x)\hat{y}_{h_i^+,(\theta_i,\varphi_i)}-\sigma^{-2}z(x)\right)\right)}{\left(\tilde{\lambda}^{\Sigma}(x)+\sigma^{-2}\left(1-\sum_i 1\right)\right)}$$

where $\tilde{\lambda}_i(x) = \sigma_i^{-2}(x)$ and $\tilde{\lambda}^{\Sigma}(x) = \sum_i \tilde{\lambda}_i(x)$. This formula allows to obtain a final estimate in the form (4) where the anisotropic kernel g_x^+ is uniform over the adaptive anisotropic neighborhood discussed in section 2.2.

A performance comparison of the proposed algorithm over the two-dimensional version, working on single frames, has been done. Table 1 presents results for this comparison: for the $2D$ case, Γ=0.9 gives an average PSNR of 30.32dB, while for the $3D$ case, using Γ=0.7 for all directions but the temporal ones (Γ=1.2) an average PSNR of 33.86dB is reached. Experimental results show that not only the $3D$ method outperforms the classical $2D$ version, but also that despite the very simple structure of the used kernels, this basic implementation of the anisotropic $3D$ filtering yields a very good performance.

In the performed tests, a smaller value for Γ is considered, with respect to the usual $2D$ algorithm, since a larger number of directional estimators are taken into account in the fusing. However, along the purely temporal directions, bigger values for Γ can be chosen, taking advantage of high stationary areas in the frames.

Figure 5 shows (a) the original version of the 41st frame of the test sequence and (b) the corrupted one. Figure 6 shows the same frame, restored applying (a) the $2D$ version of the method, obtaining PSNR=30.32dB and (b) the proposed method ($3D$ algorithm), reaching a PSNR of 33.75dB. Visual inspection shows that edges and salient points of the video frame are preserved better in Figure 6(b) than in (a).

The *ICI* adaptive scales $h_i^+(x)$ represent the distribution of image features across the direction (θ_i, φ_i). Figure

	average	min	min*	max
noisy	22.11	21.99	21.99	22.24
$2D$	30.32	30.02	30.02	30.64
$3D$	**33.86**	**32.97**	**33.45**	**34.38**

Table 1. PSNR (dB) values of the noisy (σ=20) and restored *Akiyo* sequence. Filtering is performed using the $2D$ and the $3D$ *LPA-ICI* estimator. Average, minimum and maximum PSNR values are calculated frame by frame on the whole (300 frames) sequence; min* is the minimum value of PSNR obtained on the *trimmed* sequence from frame 10 to 291.

(a) Original (b) Noisy

Fig. 5. Frame 41 from the *Akiyo* sequence: (a) original and (b) noisy observation.

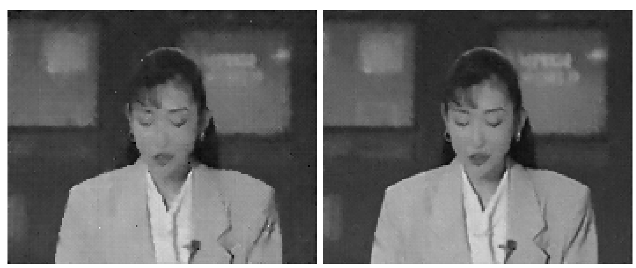

(a) $2D$ (b) $3D$

Fig. 6. Frame 41 from the *Akiyo* sequence: (a) restored with $2D$ algorithm, (b) restored with the proposed $3D$ algorithm.

7 shows these adaptive scales and the corresponding directional estimates for three different directions. In particular, Figure 7(a)-(b) are obtained for a purely spatial direction (left, θ=0, φ=$\pi/2$); Figure 7(c)-(d) are obtained for a spatio-temporal direction (θ=$3\pi/4$, φ=$3\pi/4$); Figure 7(e)-(f) are obtained for the purely temporal direction in the future (φ=π). It is remarkable how the temporal directions can give important information on the motion in the video sequence, selecting larger adaptive scales (white areas in the figure) for points that show slow motion or no motion at all from frame to frame, and smaller scales (dark areas) for points that move from frame to frame.

7. ACKNOWLEDGMENTS

The work of Dr. Vladimir Katkovnik was supported in part by *Visiting Fellow* grant from Nokia Foundation.

8. REFERENCES

[1] Fan J. and Gijbels I., *Local polynomial modelling and its application*. London: Chapman and Hall, 1996.

[2] Foi, A., Katkovnik, V., Egiazarian, K. & Astola, J., "A novel anisotropic local polynomial estimator based on directional multiscale optimizations", *Proc. of the 6th IMA Int. Conf. Math. in Signal Processing*, Cirencester (UK), pp. 79-82, 2004.

[3] Goldenshluger, A. and Nemirovski A., "On spatial adaptive estimation of nonparametric regression," *Math. Meth. Statistics,* vol. 6, pp. 135-170, 1997.

[4] Katkovnik V., "A new method for varying adaptive bandwidth selection," *IEEE Trans. on Signal Processing,* vol. 47, no. 9, pp. 2567-2571, 1999.

[5] Katkovnik V., Egiazarian K., and Astola J., "Adaptive window size image de-noising based on intersection of confidence intervals (*ICI*) rule", *Journal of Mathematical Imaging and Vision*, vol. 16, N. 3, pp. 223-235, 2002.

[6] Katkovnik V., Egiazarian K., and Astola J., "*Adaptive varying scale methods in image processing*". Tampere International Center for Signal Processing, TICSP Series, N. 19, Tampere, TTY, Monistamo, 2003.

[7] Katkovnik V., Foi A., Egiazarian K., and Astola J., "Directional varying scale approximations for anisotropic signal processing," *EUSIPCO 2004*, Wien, 2004.

[8] Katkovnik V., Foi A., Egiazarian K., and Astola J., "Anisotropic Local Likelihood Approximations: Theory, Algorithms, Applications", *Proc. of El. Imaging 2005: Image Processing: Algorithm and Systems IV*, 5672-19, 2005.

[9] Lepski, O., Mammen, E. and Spokoiny, V. , "Optimal spatial adaptation to inhomogeneous smoothness: an approach based on kernel estimates with variable bandwidth selection," *Annals of Statistic*s, vol. 25, no. 3, 929–947, 1997.

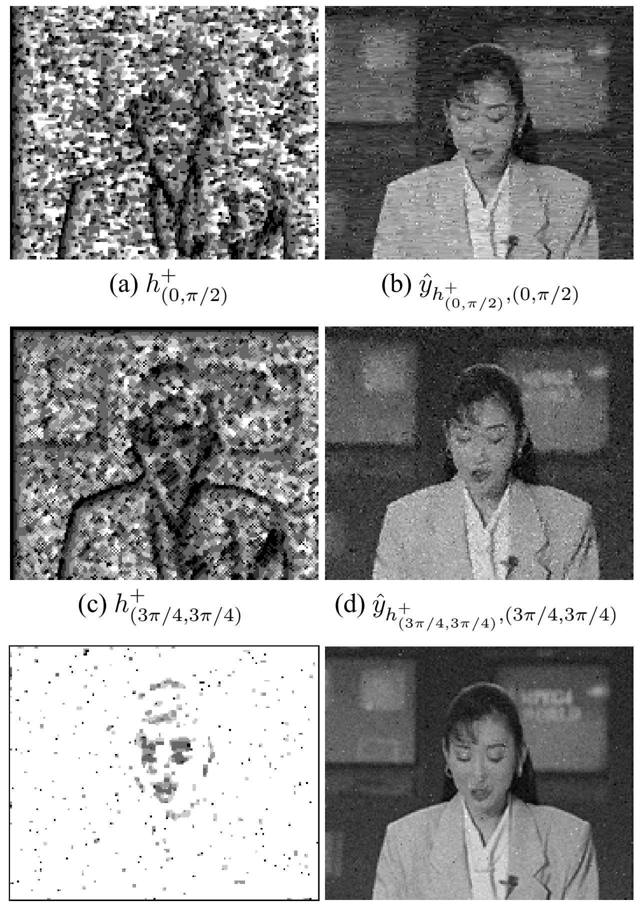

(a) $h^+_{(0,\pi/2)}$ (b) $\hat{y}_{h^+_{(0,\pi/2)},(0,\pi/2)}$

(c) $h^+_{(3\pi/4,3\pi/4)}$ (d) $\hat{y}_{h^+_{(3\pi/4,3\pi/4)},(3\pi/4,3\pi/4)}$

(e) $h^+_{(\varphi=0)}$ (f) $\hat{y}_{h^+_{(\varphi=0)},(\varphi=0)}$

Fig. 7. Adaptive scales h^+_i and directional estimates $\hat{y}_{h^+_i,(\theta_i,\varphi_i)}$ for the 41st frame from the *Akiyo* sequence: (a)-(b) purely spatial direction ($\theta = 0, \varphi = \pi/2$); (c)-(d) spatio-temporal direction ($\theta = 3\pi/4$, $\varphi = 3\pi/4$); (e)-(f) purely temporal direction in the future ($\varphi = \pi$). Darker colour is used in the left column to represent smaller scales. PSNR (dB) values for the directional estimates shown in the right column are, from top to bottom, 23.69, 23.91 and 29.56, respectively.