Published in final edited form as:

IEEE Trans Image Process. 2008 June ; 17(6): 847–856. doi:10.1109/TIP.2008.920795.

Dynamic Denoising of Tracking Sequences

Oleg Michailovich [Member, IEEE]

School of Electrical and Computer Engineering, University of Waterloo, Waterloo ON N2L 3G1, Canada (olegm@ece.uwaterloo.ca)

Allen Tannenbaum [Member, IEEE]

Schools of Electrical and Computer and Biomedical Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA

Department of Electrical Engineering, The Technion—Israel Institute of Technology, Haifa, Israel (tannenba@ece.gatech.edu)

Abstract

In this paper, we describe an approach to the problem of simultaneously enhancing image sequences and tracking the objects of interest represented by the latter. The enhancement part of the algorithm is based on Bayesian wavelet denoising, which has been chosen due to its exceptional ability to incorporate diverse *a priori* information into the process of image recovery. In particular, we demonstrate that, in dynamic settings, useful statistical priors can come both from some reasonable assumptions on the properties of the image to be enhanced as well as from the images that have already been observed before the current scene. Using such priors forms the main contribution of the present paper which is the proposal of the *dynamic denoising* as a tool for simultaneously enhancing and tracking image sequences. Within the proposed framework, the previous observations of a dynamic scene are employed to enhance its present observation. The mechanism that allows the fusion of the information within successive image frames is Bayesian estimation, while transferring the useful information between the images is governed by a Kalman filter that is used for both prediction and estimation of the dynamics of tracked objects. Therefore, in this methodology, the processes of target tracking and image enhancement “collaborate” in an interlacing manner, rather than being applied separately. The dynamic denoising is demonstrated on several examples of SAR imagery. The results demonstrated in this paper indicate a number of advantages of the proposed dynamic denoising over “static” approaches, in which the tracking images are enhanced independently of each other.

Keywords

Bayesian estimation; Kalman filtering; predictive tracking; wavelet denoising

I. Introduction

In this paper, we consider the standard problem in which a sequence of images of a maneuvering target is given with the goal being to estimate the target dynamics as precisely as possible. Such an estimation task typically consists of establishing a functional relationship between the appearance of the target along consecutive image frames. In particular, depending on a specific application at hand, the apparent motion of the target can be perceived through a deformation of its boundary, which can be described by means of either active contours [1] or active

polygons [2]. Another possibility for describing the target motion is to represent the latter by apparent displacements of a set of control points pertaining to target's image [3]. The method of optical flows constitutes another practically important solution to the problem of modeling the apparent motion [4].

Given a specific description of the apparent motion of an object of interest (as well as a relation of this motion to the actual 3-D motion of the object in space), the problem of predictive tracking could be readily solved by fitting (in a proper sense) a dynamic model to available observations [5, Ch.2]. Such a direct approach, however, may not be satisfactory in the settings when the levels of measurement noises are particularly high. In such cases, applying an image enhancement procedure should be considered as an important prerequisite which could substantially improve the robustness of tracking [6]. On the other hand, applying a denoising procedure to each image of the sequence *independently* seems to be a suboptimal strategy, as well, since it discards the information which has been observed up to the current scene.

A more accurate solution to the aforementioned problem seems to be possible using a method that takes advantage of all the useful information about the tracked object which is contained in tracking images up to the present time point. Moreover, such information should be used for *both* enhancing the current image and updating the dynamic parameters of the target. In such a case, a basic processing step may include enhancing the current image using all the available estimates of the previous images and of their associated dynamic parameters, followed by updating the dynamic parameters pertaining the current image using its *enhanced* version.

An algorithm that uses the principle described above is proposed in the present paper. The enhancement part of the proposed method is based on wavelet denoising, initially proposed in [7] as a powerful method of recovering nonstationary signals. A particularly useful feature of wavelet denoising consists of its ability to recover spatially nonhomogeneous signals without oversmoothing their details, while barely exceeding the computational demands of ordinary linear filtering. In this paper, the wavelet denoising is performed within the Bayesian framework [8], which allows incorporating any *a priori* information on the signal to be recovered. In the current implementation, such *a priori* knowledge comes from the statistical model assumed for the wavelet coefficients of tracking images, as well as from their already estimated predecessors.

The dynamic part of the algorithm is based on modeling the target motion using an *affine* optical flow, whose time evolution is governed by a Kalman filter [9]. More specifically, a two-step Kalman filter is applied in order to estimate and predict the apparent target motion. It is worthwhile noting that this is the particular structure of the Kalman filtering, which includes prediction as its intrinsic stage, is what allows the wavelet denoising and tracking to be effectively combined together. This construction will be described in much greater details in the sections that follow.

Finally, in the experimental part of this study, the applicability of the proposed dynamic denoising to synthetic aperture radar (SAR) imagery is demonstrated. It is worthwhile noting that the use of SAR imaging for the purpose of tracking is currently considered to be a problem of substantial importance. This is mainly due to the belief that the SAR imagers may come as an attractive alternative to existing optical devices, primarily because of the exceptional ability of the former to be all-weather operational. Unfortunately, target tracking using SAR data still seems to pose a formidable challenge to both system and image processing engineers, mainly because of the relatively high levels of noises contaminating this type of imagery data. This is why the SAR imaging has been chosen as an example of image modality that, we believe, could benefit from the methodology proposed in this paper.

The present paper is organized as follows. Section II introduces the Bayesian framework for wavelet denoising. The way through which the information on previously enhanced images is “injected” into the current image estimate is detailed in Section III. Section IV introduces the dynamical model for affine optical flows. The overall algorithm structure is defined in Section V. Finally, some relevant reconstruction examples are demonstrated in Section VI of the paper, while Section VII finalizes the paper with a discussion and conclusions.

II. Bayesian Wavelet Denoising

A. Diagonal Estimation and Wavelet Shrinkage

In this section, the problem of enhancement of a single image $g(x)$ (with $x \in \Omega \subseteq \mathbb{R}^2$) is considered first. To this end, $g(x)$ is assumed to be an observation of the original scene $f(x)$ contaminated by additive white Gaussian noise. Consequently, $f(x)$ can be estimated via applying to $g(x)$ a *diagonal operator* of the following form [10, Ch.X]

$$f(x) \simeq \sum_{k \in \mathbb{Z}} a_k \langle g, \varphi_k \rangle \varphi_k(x) \quad (1)$$

with $\{\varphi_k(x)\}_{k \in \mathbb{Z}}$ being an orthonormal basis in $\mathbb{L}_2(\Omega)$, $\langle \cdot, \cdot \rangle$ denoting the standard inner product, and $\{a_k\}_{k \in \mathbb{Z}}$ being a sequence of scalars satisfying $|a_k| \leq 1$ for all k . It is of importance to note that, if the above basis provides a *sparse representation* of the image of interest, then the diagonal estimator (1) can be shown to be nearly optimal (i.e., resulting in the smallest variance) among all nonlinear estimators [10, Ch.X]. As a result, wavelet orthonormal bases have become here the most preferable choice, as providing sparse representations for bounded-variation signals—the functional class, to which most of real-life images are generally assumed to belong. It should also be noted that there exists a number of Bayesian approaches to the problem of image denoising which eventually lead to the estimation (1). By virtue of the property of sequence $\{a_k\}_{k \in \mathbb{Z}}$ to be positive and bounded by unity, such methods are commonly referred to as *shrinkage* estimators [11].

Most of the Bayesian wavelet shrinkage methods proposed so far can be divided into two main groups based on the way they treat dependencies between the wavelet coefficients of $f(x)$. The first group encompasses the methods that assume the coefficients to be nearly independent due to the decorrelation property of wavelet transforms [12]–[17]. It is worthwhile noting that, in this case, the assumption of independency allows one to estimate each wavelet coefficient separately of the others. On the other hand, the methods of the second group attempt to employ the information contained in the joint behavior of wavelet coefficients. For instance, [18]–[20] perform the shrinkage adaptively to each coefficient using the concept of an activity function of a wavelet coefficient defined over a local neighborhood of the latter. The inter-scale dependencies between wavelet coefficients can be taken into account using either the fusion procedure of [21] or the approaches of [22], [23] which are based on the theory of Markov random fields.

Although the methods of the second group are more general and, hence, can provide more accurate denoising results as compared with the methods of the first group, the latter may still be preferable in the situations when computational complexity is of utmost importance. Since this is almost always the case in tracking applications, we assume the wavelet coefficients of $f(x)$ to be mutually independent—an assumption whose modification, removal, or mitigation can be derived in subsequent extensions of the method proposed below. Moreover, as a basis for our developments, we will use the wavelet shrinkage method exploited in [15] and [16].

B. MMSE Solution for Wavelet Shrinkage

Let us consider the 1-D case first. A 1-D wavelet transform provides a representation of the signal of interest in terms of the basis functions $\psi_{j,k}(x) \stackrel{\text{def}}{=} 2^{-j/2} \psi(2^{-j}x - k)$, which are dilations at scale j and translations by $2^j k$ of the mother wavelet function $\psi(x)$. The *discrete wavelet transform* (DWT) of a 1-D signal is implemented by two-channel subband filtering followed by downsampling by a factor of 2 [10, Ch.VII]. The transformation results in a sequence of wavelet coefficients (doubly indexed by the pair (k,j)) which represent the energetic contents of the signal in certain locations and resolutions.

Let $G_{j,k}$, $F_{j,k}$, and $U_{j,k}$ denote the wavelet coefficients of the observed signal $g(x)$, the original signal $f(x)$, and the noise process $u(x)$, respectively. Then, due to the linearity of the DWT, we have

$$G_{j,k} = F_{j,k} + U_{j,k}. \quad (2)$$

In order to construct an estimator for $F_{j,k}$ using the Bayesian framework, the probabilistic models for the signals in (2) should be specified first. To this end, the noise samples are assumed to be normal with mean 0 and variance σ^2 . Namely, the conditional *probability density function* (pdf) of $U_{j,k}$ given $F_{j,k}$ and σ^2 is assumed to obey

$$p(U_{j,k} | F_{j,k}, \sigma^2) = \mathcal{N}(F_{j,k}, \sigma^2), \quad \forall k, j. \quad (3)$$

At the same time, the signal coefficients $F_{j,k}$ are assumed to be independently distributed according to

$$p(F_{j,k} | \gamma_{j,k}, \sigma^2) = \mathcal{N}(0, \gamma_{j,k} c_j \sigma^2), \quad \forall k, j \quad (4)$$

where $c_j > 0$ is a scale dependent constant, and $\gamma_{j,k}$ are independent random variables that obey

$$p(\gamma_{j,k}) = \text{Bernoulli}(\pi_j), \quad \forall k, j. \quad (5)$$

We note that the model above implies that the signal coefficients are normally distributed with variance $c_j \sigma^2$ when $\gamma_{j,k} = 1$ and degenerate to zero, when $\gamma_{j,k} = 0$. The practical meaning of $\gamma_{j,k}$ is quite intuitive: the set of indices where it is equal to 1 designates significant coefficients of the useful signal $f(x)$, whereas the complementary set is associated with the noise alone. Note that such a model well agrees with the generally parsimonious representation of signals in the wavelet domain, where only a relatively small number of signal coefficients bear the most significant portion of its energy.

Let $p(G_{j,k} | \gamma_{j,k} = 1)$ and $p(G_{j,k} | \gamma_{j,k} = 0)$ denote the pdf of $G_{j,k}$ *conditional* on the presence and absence of the useful signal in the coefficient indexed by (j,k) , respectively. The *posterior probability* $\pi(\gamma_{j,k} = 1 | G_{j,k})$ that the coefficient contains the useful signal can be expressed as a function of the posterior odds $O_{j,k}$ given by

$$\pi(\gamma_{j,k} = 1 | G_{j,k}) = \frac{O_{j,k}}{1 + O_{j,k}} \quad (6)$$

where

$$\begin{aligned} O_{j,k} &= \frac{p(\gamma_{j,k}=1)p(G_{j,k}|\gamma_{j,k}=1)}{p(\gamma_{j,k}=0)p(G_{j,k}|\gamma_{j,k}=0)} \\ &= \frac{\pi_j}{1-\pi_j} \sqrt{\frac{c_j}{1+c_j}} \exp \left\{ \frac{1}{2\sigma^2} \frac{c_j}{1+c_j} G_{j,k}^2 \right\}. \end{aligned} \quad (7)$$

Consequently, the *posterior mean* of $F_{j,k}$, which is nothing else but its *minimum mean square estimate* (MMSE) $\hat{F}_{j,k}$, can be defined as

$$\hat{F}_{j,k} \stackrel{\text{def}}{=} \mathcal{E} \{ F_{j,k} | G_{j,k} \} = \left[\frac{c_j}{1+c_j} \pi(\gamma_{j,k}=1 | G_{j,k}) \right] G_{j,k} = \left[\frac{c_j}{1+c_j} \cdot \frac{O_{j,k}}{1+O_{j,k}} \right] G_{j,k}. \quad (8)$$

Denoting by $a_{j,k} = c_j(1+c_j)^{-1}O_{j,k}(1+O_{j,k})^{-1}$ (so that $\hat{F}_{j,k} = a_{j,k}G_{j,k}$), one can see that the MMS estimate (8) is indeed of the diagonal form (1).

Finally, we note that, whenever the noise standard deviation σ is unknown, it can be estimated from the available data as $\sigma \approx \text{median}_k |G_{1,k}| / 0.6745$ [7].

C. Inclusion of Spatial Priors

It is interesting to note that $\hat{F}_{j,k}$ in (8) can be alternatively derived as a solution to the problem of *simultaneously detecting and estimating* signals in a noisy environment as it was originally addressed by Middleton and Esposito in [24]. This perspective provides a very useful insight into the structure of $\hat{F}_{j,k}$ in (8), which appears as a product of the MMS estimate of $\hat{F}_{j,k}$ conditional on $\gamma_{j,k} = 1$ (which is the “estimation part” of the combined estimate) and the *Bayesian detection term* $O_{j,k}/(1+O_{j,k})$, with $O_{j,k}$ being the familiar *generalized likelihood ratio*.

From the discussion above, it follows that if one was given additional information as to the presence of the useful signal $F_{j,k}$ in $G_{j,k}$ then this information could be incorporated in the “detection part” of (8). This can be done by replacing the *unconditional* prior probabilities $p(\gamma_{j,k} = 1)$ and $p(\gamma_{j,k} = 0)$ in (7) with the priors $p(\gamma_{j,k} = 1 | I_{j,k})$, and $p(\gamma_{j,k} = 0 | I_{j,k})$ respectively, which are *conditional* on this additional information $I_{j,k}$. In this case, according to the Bayes rule, the *prior ratio* (to be substituted in (7) in place of $\pi_j/(1-\pi_j)$) is given by

$$\begin{aligned} r_{j,k} &= \frac{p(\gamma_{j,k}=1)p(I_{j,k}|\gamma_{j,k}=1)}{p(\gamma_{j,k}=0)p(I_{j,k}|\gamma_{j,k}=0)} \\ &= \frac{\pi_{j,k}}{1-\pi_{j,k}} \frac{p(I_{j,k}|\gamma_{j,k}=1)}{p(I_{j,k}|\gamma_{j,k}=0)}. \end{aligned} \quad (9)$$

In the present work, $I_{j,k}$ is defined to represent the information about the positions of the signal coefficients $F_{j,k}$, so that the conditional densities $p(I_{j,k} | \gamma_{j,k} = 1)$ and $p(I_{j,k} | \gamma_{j,k} = 0)$ (which are defined over the index set (j,k)) quantify the probability of either a signal or noise wavelet coefficient to be present at a specific scale/location indexed by j and k , respectively. In this case, to proceed quantitatively, the question to be addressed next is: Where are the above “position probabilities” coming from? The noise part is the easiest. Since the noise is assumed to be white and Gaussian, its presence at any data coefficient indexed by (j,k) is equiprobable, implying that $p(I_{j,k} | \gamma_{j,k} = 0) = \| (j,k) \|^{-1}$, where $\| (j,k) \|$ is the cardinality of the set (j,k) . The situation with $p(I_{j,k} | \gamma_{j,k} = 1)$, however, is not that simple, and this is where the dynamics comes into play. The Section III-A provides a simple recipe for constructing this probability for an image, based on its *time-delayed* estimate.

III. Dynamic Denoising

A. Two-Dimensional Separable DWT

A 2-D separable DWT decomposes an image into a hierarchy of four subbands. At each scale j (also referred to as a decomposition level), the subbands consist of an approximation subband LL_j and three detail subbands LH_j , HL_j , and HH_j . While the approximation subband LL_j contains the low-frequency portion of the original image, the detail subbands LH_j , HL_j , and HH_j capture the image details extending in the horizontal, vertical, and diagonal directions, respectively. In the course of the decomposition, subband LL_j is used as an input for the decomposition level $j + 1$, with LL_0 representing the original image.

In the case of a separable DWT, there is a “canonical” way to organize the wavelet coefficients as it is shown in Fig. 1(A) for the case of three decomposition levels [Fig. 1(B) exemplifies this structure by showing the DWT coefficients of the standard “Lena” image]. In the course of wavelet-based image enhancement, the approximation coefficients are usually not processed. On the other hand, the detail coefficients of the image can be processed using the wavelet shrinkage procedure (8) under replacing the scale index j by a double index (j, m) , with m accounting for a specific orientation.

B. Estimation of “Location” Probabilities

In what follows, the images under consideration are assumed to be square $N \times N$ images with N defined to be an integral power of 2. Consequently, according to the “canonical” structure of Fig. 1(A), the wavelet coefficients of image f at scale j and orientation m can be viewed as a $2^{-j}N \times 2^{-j}N$ matrix. Let this matrix be denoted by $\mathbf{F}_{j,m}$, with its (k_1, k_2) th element (with $k_1, k_2 = 0, 1, \dots, 2^{-j}N$) being equal to the wavelet coefficient of f corresponding to scale j , orientation m , and position (k_1, k_2) .

Additionally, let $\mathbf{P}_{j,m}^1$ and $\mathbf{P}_{j,m}^0$ be two square matrices of the *same* size as $\mathbf{F}_{j,m}$. These matrices will contain the “location probabilities” $p(I_{j,k} | \gamma_{j,k} = 1)$ and $p(I_{j,k} | \gamma_{j,k} = 0)$, respectively, as discussed in Section II-C. In order to specify the above matrices, let us first define an $2^{-j}N \times 2^{-j}N$ indicator matrix $\mathbf{I}_{j,m}$ as

$$\mathbf{I}_{j,m}(k_1, k_2) = \begin{cases} 1, & \text{if } |\mathbf{F}_{j,m}(k_1, k_2)| \geq \tau \\ 0, & \text{if } |\mathbf{F}_{j,m}(k_1, k_2)| < \tau \end{cases} \quad (10)$$

whose nonzero entries indicate the positions of the *significant wavelet coefficient*, the absolute values of which exceed a predefined threshold $\tau > 0$. Subsequently, given an $2^{-j}N \times 2^{-j}N$ matrix $\mathbf{K}_j$ of *discrete* values of the isotropic Gaussian density function

$$K_j(x, y) = \frac{1}{\sqrt{2\pi}h_j} \exp \left\{ -\frac{1}{2} \frac{(x^2 + y^2)}{h_j^2} \right\} \quad (11)$$

the matrix $\mathbf{P}_{j,m}^1$ can be defined as

$$\mathbf{P}_{j,m}^1 = (\#\mathbf{I}_{j,m})^{-1} (\mathbf{I}_{j,m} * \mathbf{K}_j) \quad (12)$$

where $\#\mathbf{I}_{j,m} = \sum_{k_1, k_2} \mathbf{I}_{j,m}(k_1, k_2)$, and $*$ stands for the convolution operator. It deserves noting that the above construction is, in fact, a *kernel-based estimation* of the probability of

location of the significant wavelet coefficients within the chosen subband [25]. We also note that the convolution in (12) may be cyclic if the periodized, interval-adapted DWT of [26] is used to compute the wavelet coefficients (as it is implemented by the WaveLab® package of

Donoho). In this case, the normalization guarantees that $\sum_{k_1, k_2} \mathbf{P}_{j,m}^1(k_1, k_2) = 1$. Moreover, for the smoothing effect imposed by the convolution with $\mathbf{K}_j$ to be similar for different DWT levels, the band-width parameter h_j of the Gaussian function should be defined as

$$h_j = 2^{-j} h_0 \quad (13)$$

where the bandwidth $h_0 > 0$ can be precomputed using any of the standard methods described, e.g., in [25].

It goes without saying that the local dependencies between adjacent wavelet coefficients could be better accounted for, if the kernel function in (11) was chosen to be anisotropic [27]. In this paper, however, this possibility has not been explored for its being more computationally involved as compared to the simple choice in (11), which seems to be preferable for real-time tracking applications.

The “location” distribution matrix $\mathbf{P}_{j,m}^0$ pertaining to the noise process should be defined next. Since the noise is assumed to be Gaussian and white, the *uniform* distribution is the only reasonable choice here. Formally

$$\mathbf{P}_{j,m}^0(k_1, k_2) = 4^j N^{-2}, \quad \forall k_1, k_2 \quad (14)$$

so that $\sum_{k_1, k_2} \mathbf{P}_{j,m}^0(k_1, k_2) = 1$.

Finally, we note that the matrices $\mathbf{P}_{j,m}^1$ and $\mathbf{P}_{j,m}^0$ defined above play the role of the conditional probabilities $p(I_{j,k} | \gamma_{j,k} = 1)$ and $p(I_{j,k} | \gamma_{j,k} = 0)$ in (9), respectively. Consequently, once these matrices have been computed, the corresponding prior ratio [as defined by (9)] can be substituted in (7) instead of $\pi_j/(1 - \pi_j)$ to compute the MMSE (8). Fig. 2 shows an interim example that compares the performance of the Bayesian shrinkage without and with the “location” priors. For the reason which will become clear shortly, we refer to the above estimates as “static” and “dynamic” MMSE, respectively. These estimates are shown in Fig. 2(C) and (D), whereas Fig. 2(A) and (B) shows the original test image and its noise-contaminated version (SNR = 11.8 dB), respectively. Note that, in this example, the indicator functions (10) corresponded to about 20% of the largest wavelet coefficients of “Lena” for each j and m . One can see that using the additional information allows considerably improving the SNR from 17.6 to 20.1 dB.

C. Dynamically Learned Priors

The Bayesian estimation described in the preceding section would be only possible if we knew the locations of (significant) wavelet coefficients of the noise-free image f , which, of course, makes this approach rather impractical. However, in tracking applications, a noisy measurement of f is normally available along with its *time-delayed* version. Suppose now that the latter has already been enhanced. Then, using *predicted* dynamic parameters of the target motion, this time-delayed estimate of f could be *warped-forward* to obtain a *prediction* of the true image f at the present time. Needless to say, this prediction can be used for neither averaging it with the present observation of f nor for any other kind of “compounding.” Because of the error in estimating the target dynamics, the features (e.g., edges) in the predicted image

can never be expected to be perfectly aligned with those in f . Consequently, the averaging would have unavoidably smeared the resulting estimates. However, the above prediction of f seems to be a reasonably good candidate that one can use for computing the indicator functions (10) and the related “location” probability matrices $\mathbf{P}_{j,m}^1$ and $\mathbf{P}_{j,m}^0$, and this is how these matrices are computed in this study.

It should be noted that precaution should be taken in regard to using the above estimation approach. As it was noted before, because of the errors in estimating the target motion, the features of the current image f and its prediction are never perfectly aligned. Moreover, the prediction is prone to errors, since it is derived from a time-delayed, enhanced version of f , which is, by itself, an estimate. However, all these potential sources of inaccuracies have already been implicitly taken into account by the algorithm. First, the “location” probabilities are computed based only on the *significant* wavelet coefficients of the prediction (i.e., those exceeding a predefined threshold τ), which can be reasonably assumed to contain the useful signal. This makes the algorithm resistant to the errors in estimating the time-delayed version of f , based on which the prediction is computed. Second, to compute $\mathbf{P}_{j,m}^1$ and $\mathbf{P}_{j,m}^0$, the indicator functions are smoothed by the kernel functions $\mathbf{K}_j$, thereby accounting for the *uncertainty* as to the location of the wavelet coefficients of f —the uncertainty that has to be expected because of the inaccuracies in recovering the associated dynamics. Moreover, to achieve a more stable convergence of the algorithm, we replace the time-invariant bandwidth h_0 in (13) by a time-dependent bandwidth $h(t)$ defined as

$$h(t) = h_0 \left((A - 1) \alpha^t + 1 \right) \quad (15)$$

where $A \geq 1$, $0 < \alpha < 1$, and t stands for the time (or, equivalently, time-index) of a given image. In this case, when $t = 0$, the bandwidth of $\mathbf{K}_j$ is equal to $A 2^{-j} h_0$, and, as a result, if A is much greater than 1, the resulting probabilities $\mathbf{P}_{j,m}^1$ will be substantially *oversmoothed*.

Consequently, the ratios $\mathbf{P}_{j,m}^1 / \mathbf{P}_{j,m}^0$ will change relatively slowly as functions of (k_1, k_2) , implying that the algorithm “ignores” the locational priors, which may be rather unreliable in the beginning of convergence (it should be noted that, at this stage, the algorithm enhances the tracking images almost *independently of each other*). As time goes on, the estimates of the target dynamics converge to their optimal values, implying that the prediction by “warping forward” is now performed with a relatively high accuracy. This fact is reflected in the behavior of $h(t)$ that approaches h_0 as $t \rightarrow \infty$.

Apparently, the optimal values of A and α in (15) should depend on the *posterior error covariance* of the dynamic parameters as computed by, e.g., a Kalman filter (see the discussion below). Finding such an analytical relationship represents an interesting direction of our future research. In the present paper, however, the parameters A and α are left to be user-defined.

IV. Motion Estimation by Kalman Filtering

From the above discussion, it follows that estimating the target motion constitutes an integral part of the image enhancement process. It is true as well that the enhancement allows one to estimate the dynamics with a greater precision. Thus, in this section, we briefly describe the method used for the motion prediction and estimation.

In the current paper, the estimation of (apparent) target motion is based on the assumption that any subsequent image of a tracking sequence can be approximated as a locally translated version of its “predecessor,” where the translation is governed by a displacement field, also known as an *optical flow* [4]. Moreover, whenever the tracking of *rigid* targets (e.g., tanks,

armored troop-carriers) is of concern, assuming *affine* displacement fields may be quite reasonable. In this case, denoting by $\mathbf{c}_t \in \mathbb{R}^6$ the column vector of affine parameters corresponding to the image g_t observed at time t , and by Δg_t and $\partial_t g_t$ the gradient and the time derivative of the latter, respectively, it can be shown that, on the assumption of small displacements, it holds that [28]

$$\mathbf{A}_t^T(x, y) \mathbf{c}_t \simeq -\partial_t g_t(x, y), \quad \forall (x, y) \in \Omega \quad (16)$$

where

$$\mathbf{A}_t^T(x, y) = \begin{bmatrix} 1 & x & y & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & x & y \end{bmatrix}^T \Delta g_t(x, y) \quad (17)$$

with the gradient $\Delta g_t(x, y)$ being viewed as a column vector in $\mathbb{R}^2$ at every $(x, y) \in \Omega$. Note that (16) holds at every point of the domain of image definition Ω , while the vector $\mathbf{c}_t$ is position independent. Therefore, finding the optimal solution $\mathbf{c}_t^*$ for $\mathbf{c}_t$ amounts to solving an overcompleted system of equation, which can be done using the standard Moore–Penrose pseudoinverse.¹ Moreover, for the case when the model errors in (16) can be assumed to be Gaussian and white, this solution can be rigorously proven to be the MMS estimate of $\mathbf{c}_t$ [28].

The optimal solution $\mathbf{c}_t^*$ derived above provides a *momentous* estimate of the target dynamics, which does not take into account the time-coherence of target motion. To account for the time evolution of the affine parameters, the *true* vector $\mathbf{c}_t$ can be assumed to be a noisy version of $\mathbf{c}_t^*$, viz.

$$\mathbf{c}_t = \mathbf{c}_t^* + \mathbf{u}_t \quad (18)$$

where $\mathbf{u}_t$ is a noise process. Moreover, using the first-order Markov assumption, $\mathbf{c}_t$ itself can be modeled as a noisy version of its “predecessor” $\mathbf{c}_{t-1}$, namely

$$\mathbf{c}_t = \mathbf{c}_{t-1} + \mathbf{v}_t \quad (19)$$

where $\mathbf{v}_t$ is an another noise process, different from $\mathbf{u}_t$.

Equations (18) and (19) form the system of equations, which can be efficiently solved by means of Kalman filtering [9]. It is important to emphasize that the Kalman filter always estimates the state $\mathbf{c}_t$ in two steps. First, a *prediction* of the state at time t is computed based on its estimate at time $t - 1$ and (19). Second, this prediction is *updated* using the information brought in with a new observation $\mathbf{c}_t^*$ as described by (18). Thus, prediction of the affine parameters is inherently integrated in the filter structure, appearing explicitly as one of its stages.

V. Overall Diagram of DDN

In this section, some necessary details on the overall organization of the proposed dynamic denoising (DDN) are provided. As before, it is assumed that any image g_t of the tracking sequence $\{g_t\}_{t \geq 0}$ represents a noise-free image f_t contaminated by white Gaussian noise. The

¹Note that, in this case, the pseudo-inverse involves inversion of only a 6×6 matrix.

algorithm is applied recursively, and its input at time t is formed by: 1) the noisy data image g_t ; 2) an estimate $\hat{f}_{t-1}$ of f_{t-1} ; and 3) an estimate of the optical flow parameters $\tilde{\mathbf{c}}_{t-1}$ pertaining to time $t - 1$. At the output, the algorithm returns: 1) an estimate $\hat{f}_t$ of f_t and 2) an estimate $\tilde{\mathbf{c}}_t$ of the optical flow parameters at time t .

One recursion of the algorithm is carried out according to the steps shown in Fig. 3. First, the estimated affine parameters $\tilde{\mathbf{c}}_{t-1}$ at time $t - 1$ are used to compute a prediction $\hat{\mathbf{c}}_t$ of the parameters for time t according to the first stage of Kalman filtering, followed by using this prediction to obtain a prediction $\hat{f}_t$ of f_t via warping forward the previously estimated image $\hat{f}_{t-1}$. Subsequently, the wavelet coefficients of the prediction $\hat{f}_t$ are used for computing the “location” probabilities $\mathbf{P}_{j,m}^1$ and $\mathbf{P}_{j,m}^0$ as detailed in Section III. These probabilities along with the data image g_t are then passed on to the Bayesian shrinkage procedure of Sections II, resulting in the estimate $\hat{f}_t$ of f_t , which is the first output of the algorithm. Finally, the estimates $\hat{f}_{t-1}$ and $\hat{f}_t$ are used to compute the MMS estimate of the affine parameters (see Section IV), which is used by the Kalman filter to update the prediction $\hat{\mathbf{c}}_t$ resulting in a new estimate of the optical flow parameters $\tilde{\mathbf{c}}_t$. The latter forms the second output of the algorithm.

At the beginning of the recursion, the “location” probabilities can be defined to be uniform, which can be viewed as using *noninformative priors* as to the position of signal's coefficients. The parameters c_j and π_j of the probability densities in (4) and (5) are supposed to be either learned through training or estimated directly from data by means of the maximum-likelihood algorithm as detailed, e.g., in [13].

Finally, we note that once the parameters c_j and π_j (or estimates thereof) are available, each recursion of the DDN algorithm can be performed with logarithmic complexity in the number of image samples. Indeed, the most computationally expensive procedures which are involved in the estimation according to the block-diagram of Fig. 3 are warping forward, computing the “location” priors according to (12), and wavelet transformation. While the last two of the aforementioned procedures require only a few convolution operations, the complexity of warping depends on the type of image interpolation used. Fortunately, there exist computationally efficient methods which allow one to perform this procedure in logarithmic time, as well [29]. Consequently, the overall complexity of DDN remains logarithmic, which implies the possibility of its real-time implementation.

VI. Experimental Results

A. Data Preprocessing

In the experimental part of this study, the performance of the proposed DDN algorithm is tested using a set of real-world synthetic aperture radar (SAR) images of military targets. Before demonstrating some specific reconstruction examples, it should be noted that, due to the bandlimitedness of radars' transfer functions as well as because of the coherent nature of SAR image formation, such images are normally contaminated by *speckle noise* [30]. The statistical properties of this noise has been addressed in numerous studies, with the Rayleigh, K-, generalized gamma, and Nakagami distributions being among the noise models proposed hitherto.

Unfortunately, none of the existing statistical models for speckle noise suggests that the latter can be approximated by white Gaussian noise, and, hence, the results of the preceding sections do not seem to be applicable to the SAR data. To overcome this difficulty, it is common to use the property of the speckle noise of being a *multiplicative* noise, which allows one to convert the multiplicative noise into additive by means of the logarithmic transformation.

Subsequently, the log-transformed speckle noise can be rejected via a filtering procedure,

followed by the exponential transformation that takes the result thus obtained to the original domain.

The procedure explained above is known as *homomorphic despeckling* [31], and it has long been used to improve the quality of coherent-type imagers, including SAR units [32] and ultrasound scanners [33]. The main deficiency of the algorithms based on this technique, however, results from their assuming the log-transformed speckle noise to be Gaussian and white, while, in practice, this assumption seems to be rarely valid. In particular, the log-transformed speckle noise can be shown to be correlative and obeying a bi-exponential, Fisher–Tippett type distribution [34]. Since the pdf of such a distribution is heavy tailed, the corresponding noise is necessarily of a *spiky* type. Needless to say, that such a noise may not be effectively rejected by an algorithm that has been designed under the assumption of noise's “whiteness” and “Gaussianity.”

Fortunately, the log-transformed speckle noise can be “gaussianized” by means of the two step procedure proposed in [33]. At the first step of this procedure, the (complex-valued) data images are subjected to blind equalization that makes the image samples nearly uncorrelated. At the second step, the log-transformed magnitudes of the resulting images are subjected outlier-shrinkage that suppresses the spiky component of the noise. It can be shown that this preprocessing is capable of transforming the log-transformed speckle noise into nearly white Gaussian noise, while retaining the underlying image virtually intact. Moreover, although in [33] the above method was demonstrated for ultrasound images, it is also applicable to SAR imaging due to the similarity of the image formation mechanisms exploited by these imaging modalities. This fact is further supported by the example shown in Fig. 4(A1) and (B1) of which demonstrate a section of the log-domain images of a military target (a T-72 tank) before and after applying the above preprocessing, respectively.

The effect of suppressing the “spiky” component of the log-transformed speckle noise can be appreciated through comparing the corresponding histograms of the images² which are shown in Fig. 4(A2) and (B2). One can see that the histogram of the preprocessed image has the shape of a Gaussian probability density as opposed to the “heavy-tailed” behavior of the histogram of the original image. Additionally, the auto-correlation function of the preprocessed image [Fig. 4(B3)] converges to zero considerably faster as compared to the auto-correlation function of the original image [Fig. 4(A3)]. This implies that the samples of the preprocessed image are much less correlated than in the case of the original image.

In Section VI-B, all the images are enhanced in the log-transform domain after their are preprocessed by the procedure explained above. After the enhancement is completed, the resulting images may be back-transformed to the original (i.e., linear) domain, if it is required by a specific application at hand.

B. Tracking Military Targets

The upper row of Fig. 5 shows a subset of tracking images of a military target, *viz.* a T-72 tank. The images are shown in the log-transform domain after they have been subjected to the “gaussianization” procedure discussed in the previous subsection. The estimated SNR of these data images was found to be about 5 dB.

As the first step, the data images were enhanced independently of each other by the classical soft-thresholding in the wavelet domain as initially proposed in [7]. The uniform threshold was defined to be equal to $\sqrt{2\log N} \tilde{\sigma}$, with N being the number of image samples and $\tilde{\sigma}$ being the

²Only homogeneous portions of the backgrounds of the images were used to compute the histograms.

standard deviation of the noise estimated as mentioned at the end of Section II-B. The results of this wavelet denoising are shown in the second row of Fig. 5. One can see that this method is capable of effectively suppressing the noise. Unfortunately, it seems to happen on the account of substantially oversmoothing the target details (e.g., edges). On the other hand, the Bayesian shrinkage of Section II-B is capable of considerably better preserving the target's details as shown by the third row of Fig. 5. However, in this case, the preservation of the target's details comes on the account of letting a significant portion of the noise “survive” the shrinkage. Moreover, since this type of Bayesian shrinkage is applied to each of the tracking images *independently*, all the estimates appear to be equally noisy.

It should be noted that both above-mentioned results were obtained using the same type of orthogonal wavelet analysis based on the nearly symmetric wavelets of I. Daubechies with four vanishing moments [35, Ch.6]. The wavelet decomposition level was set to be equal to be $L = 4$. In the case of Bayesian shrinkage, the hyperparameters of the signal and noise pdf's were estimated directly from the data by means of the expectation-maximization algorithm as detailed in [13].

As the next step, the same decomposition and distributional parameters were used to enhance the tracking images by means of the proposed DDN method, whose results are demonstrated in the lowest row of Fig. 5. In this case, the learning parameters A and α in (15) were set to be equal to 3 and 0.85, respectively. The measurement and model noises $\mathbf{u}_t$ and $\mathbf{v}_t$ in (18) and (19) were assumed to be mutually independent, white, Gaussian noises with zero means and covariances equal to $10^{-2} \times \mathbf{I}_6$ and $10^{-3} \times \mathbf{I}_6$, respectively (where $\mathbf{I}_6$ denotes a 6×6 identity matrix). The initial value of the posterior error covariance was set to be equal to $0.5 \times \mathbf{I}_6$.

Analyzing the last row of Fig. 5 one can see that, in the beginning of convergence, the overall quality of the DDN-enhanced images is virtually identical to that of the images processed by the “static” Bayesian denoising. However, the noise contamination of the DDN-estimates goes and weakens as the frame index grows and a constantly increasing amount of information on the target is accumulated by the algorithm. In particular, one can see that, by the frame number 25, the noise decreases to the level comparable to that of the soft-thresholded images, while the edge preservation of the DDN method is similar to the case of the “static” Bayesian shrinkage. For qualitative assessment of the above results, the measure of *equivalent number of looks* (ENL) [32] was used in this paper. This measure is conventionally used in SAR imagery to quantify the contamination of images by speckle noise, and it is defined as a ratio of the squared mean and the variance of a SAR image.³ Note that, in general, higher values of ENL indicates lower levels of the noise contamination. Thus, for instance, the average ENL for the soft-thresholding and the “static” Bayesian methods were found to be equal to 9.6 and 7.2, respectively, which implies that the former approach suppresses the noise more effectively as compared to the latter (it has to be kept in mind, however, that ENL is a global criterion and, hence, incapable of assessing the preservation of image details). On the other hand, in the course of convergence, the average ENL obtained by DDN improved from 7.2 to 9.46, thereby closely reaching the smoothness of the soft-thresholded images, while notably better preserving the target's details.

VII. Discussion and Conclusion

Whenever reasonable assumptions on the statistical properties of an image of interest can be made, the Bayesian estimation framework often allows one to derive a considerably more informative estimate of the image as compared to the case when the estimation is performed based on observed data alone. This is why Bayesian estimation has long become a preferred

³For analogous purposes, a square root of ENL is used in ultrasound imagery, where it is known as speckle-SNR.

method of reconstructing signals and images, especially in the situations when *a priori* assumptions are necessary to either recover lost information or fuse multiple informational sources [36]. Unfortunately, defining the priors is well-known to be a very delicate step that, when performed incorrectly, could produce rather misleading results [37]. The central idea of the present study has been to show conceptually and experimentally that, in dynamic scenarios, some useful priors can be extracted from previous observations of the dynamic scene that needs to be enhanced. As a result, an image enhancement procedure—referred to as DDN—was introduced as a method for simultaneously enhancing sequences of images and tracking the objects of interest represented by the latter.

The results of image enhancement by means of the DDN method were compared with the results obtained using the standard soft-thresholding procedure of [7] as well as the “static” Bayesian shrinkage detailed in Section II-B. In both cases, the performance of the reference approaches was found to be inferior to that of the proposed method. In particular, *in the steady state*, the DDN algorithm outperformed the soft-thresholding in terms of edge-preservation, while resulting in substantially more effective noise cancellation as compared with the “static” Bayesian approach.

The above results, however, can be only regarded as preliminary, as a formal proof of the convergence properties of the DDN algorithm still needs to be derived. Specifically, a dependency between the “annealing” parameters A and α in (15) and the posterior error covariance of the Kalman filter should be found. We note that this dependency appears to be central to the convergence properties of the proposed method. Indeed, in the beginning of the recursion, when the Kalman filter has not yet converged, and, hence, the error covariance is relatively large, the “location” priors $\mathbf{P}_{j,m}^1$ and $\mathbf{P}_{j,m}^0$ *have to be* oversmoothed. This is necessary to make the algorithm discard the information coming from the past observations as not sufficiently certain. Subsequently, as the filter converges (implying that the error covariance decreases), the amount of smoothing can be adequately diminished. Thus, in some sense, the degree of smoothing in estimating the “location” priors is to be inversely proportional to the posterior error covariance of the Kalman filter.

The wavelet shrinkage procedure used in this paper was based on orthogonal wavelet transforms. While being preferable from many perspectives, the latter have been observed to produce Gibbs-like artifacts when used in denoising applications. On the other hand, this undesirable phenomenon appears to be much less pronounced, when nonorthogonal *undecimated* wavelet decompositions are used instead [38]. Consequently, modifying the DDN method to take advantage of the above property of undecimated wavelet transforms may be a possible development of the present study.

In this paper, the apparent motion of tracked objects was described by affine optical flows. Needless to say that whenever the motion of a nonrigid object has to be estimated and/or the scenario involves tracking of more than one target, the affine model may not be appropriate for the task at hand. In this case, using the displacement models of higher complexity, such as the local affine [39] or spline-based [40] models, can be considered instead. Moreover, the linear Kalman filter employed in the present study can be replaced by either by a more general *extended* Kalman filter [41] when (18) and (19) are known to be nonlinear or by a particle filter [42] when the motion distribution is, e.g., multimodal (as it would be the case in multiobject tracking). All the above modifications may also be considered as possible developments of the method proposed in this paper.

Acknowledgments

The authors would like to thank Dr. A. Betser for very helpful conversations on Kalman filtering. The authors would also like to thank all the anonymous reviewers, whose input allowed substantial improvement of the quality of this paper.

This work was supported in part by the National Science Foundation, in part by the the Air Force Office of Sponsored Research, in part by the Army Research Office, in part by MURI, in part by MRI-HEL, in part by the National Institutes of Health (NAC P41 RR-13218), and in part by the Natural Sciences and Engineering Research Council of Canada under a Discovery grant. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Pier Luigi Dragotti.

Biographies

Oleg Michailovich (M'02) was born in Saratov, Russia, in 1972. He received the M.Sc. degree in electrical engineering from the Saratov State University in 1994 and the M.Sc. and Ph.D. degrees in biomedical engineering from The Technion—Israel Institute of Technology, Haifa, in 2003.

He is currently with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada. His research interests include the application of image

processing to various problems of image reconstruction, segmentation, inverse problems, nonparametric estimations, approximation theory, and multiresolution analysis.

Allen Tannenbaum (M'93) was born in New York in 1953. He received the Ph.D. degree in mathematics from Harvard University, Cambridge, MA, in 1976.

He has held faculty positions at the Weizmann Institute of Science, Rehovot, Israel; McGill University, Montreal, QC, Canada; ETH, Zurich, Switzerland; The Technion—Israel Institute of Technology, Haifa; the Ben-Gurion University of the Negev, Israel; and the University of Minnesota, Minneapolis. He is presently the Julian Hightower Professor of Electrical and Biomedical Engineering, Georgia Institute of Technology, Atlanta, and Emory University, Atlanta. He has done research in image processing, medical imaging, computer vision, robust control, systems theory, robotics, semiconductor process control, operator theory, functional analysis, cryptography, algebraic geometry, and invariant theory.

References

- [1]. Malladi R, Sethian J, Vemuri B. Shape modeling with front propagation: A level set approach. *IEEE Trans. Pattern Anal. Mach. Intell* Feb.;1995 17(2):1 158–175.
- [2]. Unal G, Krim H, Yezzi A. Fast incorporation of optical flow into active polygons. *IEEE Trans. Image Process* Jun.;2005 14(6):745–759. [PubMed: 15971774]
- [3]. Blake, A.; Yuille, A. *Active Vision*. Mass. Inst. Technol.; Cambridge, MA: 1992.
- [4]. Horn P, Schunck B. Determining optical flow. *Artif. Intell* Aug.;1981 17:185–203.
- [5]. Raol, J.; Girija, G.; Singh, J. *Modeling and Parameter Estimation of Dynamic Systems*. IET; London, U.K.: 2004.
- [6]. Kornprobst P, Deriche R, Aubert G. Image sequence analysis via partial differential equations. *J. Math. Imag. Vis* Sep.;1999 11(1):5–26.
- [7]. Donoho D. De-noising by soft-thresholding. *IEEE Trans. Inf. Theory* May;1995 41(5):613–627.
- [8]. Muller, P.; Vidakovic, B. *Lecture Notes in Statistics*. Vol. vol. 141. Springer-Verlag; New York: 1999. Bayesian inference in wavelet-based models.
- [9]. Sorenson, H. *Kalman Filtering: Theory and Application*. IEEE; Piscataway, NJ: 1985.
- [10]. Mallat, S. *A Wavelet Tour of Signal Processing*. Academic; New York: 1998.
- [11]. Gruber, M. *Improving Efficiency by Shrinkage: The James–Stein and Ridge Regression Estimators*. Marcel Dekker; New York: 1998.
- [12]. Moulin P, Liu J. Analysis of multiresolution image de-noising schemes using generalized Gaussian and complexity priors. *IEEE Trans. Inf. Theory* Apr.;1999 45(3):909–919.
- [13]. Johnstone, M.; Silverman, BW. *Empirical Bayes selection of wavelet threshold*. Stanford Univ.; Oxford Univ.; Stanford, CA: Oxford, U.K.: 2004.
- [14]. Chipman HA, Kolaczyk E, McCulloch R. Adaptive Bayesian wavelet shrinkage. *J. Amer. Statist. Assoc* Dec.;1997 92(440):1413–1421.
- [15]. Clyde M, Parmigiani G, Vidakovic B. Multiple shrinkage and subset selection in wavelets. *Biometrika* 1998;85:391–402.

- [16]. Abramovich, F.; Sapatinas, T. Bayesian approach to wavelet decomposition and shrinkage. In: Muller, P.; Vidakovic, B., editors. *Bayesian Inference in Wavelet-Based Models*. Springer; New York: p. 33-50.ch. 3
- [17]. Vidakovic B. Nonlinear wavelet shrinkage with Bayes rules and Bayes factors. *J. Amer. Statist. Assoc* Mar.;1998 93(441):173–179.
- [18]. Chang G, Yu B, Vetterli M. Spatially adaptive wavelet thresholding with context modeling for image de-noising. *IEEE Trans. Image Process* Sep.;2000 9(9):1532–1546. [PubMed: 18262991]
- [19]. Simoncelli, E. Bayesian de-noising of visual images in the wavelet domain. In: Muller, P.; Vidakovic, B., editors. *Bayesian Inference in Wavelet-Based Models*. Springer; New York: p. 291-308.ch. 18
- [20]. Pizurica, Philips W, Lemahieu I, Acheroiy M. A versatile wavelet domain noise filtration technique for medical imaging. *IEEE Trans. Med. Imag* Mar.;2003 22(3):323–331.
- [21]. Scharcanski J, Jung C, Clarke R. Adaptive image de-noising using scale and space consistency. *IEEE Trans. Image Process* Sep.;2002 11(9):1092–1100. [PubMed: 18249730]
- [22]. Crouse M, Nowak RD, Baraniou RG. Wavelet-based statistical signal processing using hidden Markov models. *IEEE Trans. Image Process* Apr.;1998 46(4):886–902.
- [23]. Pizurica, Philips W, Lemahieu I, Acheroiy M. A joint inter- and intra-scale statistical model for Bayesian wavelet based image de-noising. *IEEE Trans. Image Process* May;2002 11(5):545–557. [PubMed: 18244654]
- [24]. Middleton D, Esposito R. Simultaneous optimum detection and estimation of signals in noise. *IEEE Trans. Inf. Theory* May;1968 14(5):434–444.
- [25]. Silverman, BW. *Density Estimation for Statistics and Data Analysis*. CRC; Boca Raton, FL: 1986.
- [26]. Cohen A, Daubechies I, Vial P. Wavelets on the interval and fast wavelet transform. *Appl. Comput. Harmon. Anal* 1993;1(1):54–81.
- [27]. Breiman L, Meisel W, Purcell E. Variable kernel estimates of multivariate densities. *Technometrics* 1977;19:135–144.
- [28]. Robinson D, Milanfar P. Fast local and global projection-based methods for affine motion estimation. *J. Math. Imag. Vis* 2003;18:35–54.
- [29]. Unser M, Thevenaz P, Yaroslavsky L. Convolution-based interpolation for fast, high-quality rotation of images. *IEEE Trans. Image Process* Oct.;1995 4(10):1371–1381. [PubMed: 18291969]
- [30]. Goodman JW. Some fundamental properties of speckle. *J. Opt. Soc. Amer* Nov.;1976 66(11):1145–1150.
- [31]. Jain, AK. *Fundamentals of Digital Image Processing*. Prentice-Hall; Englewood Cliffs, NJ: 1989.
- [32]. Guo H, Odegard JE, Lang M, Gopinath RA, Selesnick IW, Burrus CS. Wavelet based speckle reduction with application to SAR based ATD/R. *Proc. Int. Conf. Image Processing* 1994;1:75–79.
- [33]. Michailovich O, Tannenbaum A. De-speckling of medical ultrasound images. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr* Jan.;2006 53(1):64–78.
- [34]. Arsenault HH, April G. Properties of speckle integrated with a finite aperture and logarithmically transformed. *J. Opt. Soc. Amer* Nov.;1976 66(11):1160–1163.
- [35]. Daubechies, I. *Ten Lectures on Wavelets*. SIAM; Philadelphia, PA: 1992.
- [36]. Maskell S. A Bayesian approach to fusing uncertain, imprecise and conflicting information. *Inf. Fusion*. 2007
- [37]. Carlin, BP.; Louis, TA. *Bayes and Empirical Bayes Methods for Data Analysis*. CRC; Boca Raton, FL: 2000.
- [38]. Gnanadurai D, Sadasivamb V. Undecimated wavelet based speckle reduction for SAR images. *Pattern Recognit. Lett* May;2005 25(6):793–800.
- [39]. Periaswamy S, Farid H. Elastic registration in the presence of intensity variations. *IEEE Trans. Med. Imag* Jul.;2003 22:865–874.
- [40]. Kybic J, Thevenaz P, Nirkko A, Unser M. Unwarping of unidirectionally distorted EPI images. *IEEE Trans. Med. Imag* Feb.;2000 19(2):80–93.
- [41]. Ljung L. Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems. *IEEE Trans. Autom. Contr* Feb.;1979 24(1):36–50.

- [42]. Risfic, B.; Arulampalam, S.; Gord, N. Beyond the Kalman Filter: Particle Filters for Tracking Applications. Artech House; Boston, MA: 2004.

Fig. 1.
 (A) "Canonical" arrangement of (orthogonal) wavelet coefficients of an image. (B) DWT decomposition of the Lena image.

Fig. 2.

(A) Original image of “Lena”; (B) noisy image of “Lena” (SNR = 11.8 dB); (C) image estimation using the Bayesian shrinkage *without* the “location” priors (SNR = 17.6 dB); (D) image estimation using the Bayesian shrinkage *with* the “location” priors (SNR = 20.1 dB).

Fig. 3.
Block-diagram of the DDN algorithm.

Fig. 4. (A1)–(A3) Original image of a T-72 target in the log-domain, its auto-correlation and histogram, respectively. (B1)–(B3) Corresponding preprocessed image, its auto-correlation and histogram, respectively.

Fig. 5. (First row) Original images of the T-72 target; (second row) reconstruction by the uniform soft-thresholding; (third row) reconstruction by the “static” Bayesian shrinkage; (fourth row) reconstruction by the DDN method.