Published in final edited form as:

IEEE Trans Ultrason Ferroelectr Freq Control. 2006 January ; 53(1): 64–78.

Despeckling of Medical Ultrasound Images

Oleg V. Michailovich and Allen Tannenbaum [Member, IEEE]

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

Abstract

Speckle noise is an inherent property of medical ultrasound imaging, and it generally tends to reduce the image resolution and contrast, thereby reducing the diagnostic value of this imaging modality. As a result, speckle noise reduction is an important prerequisite, whenever ultrasound imaging is used for tissue characterization. Among the many methods that have been proposed to perform this task, there exists a class of approaches that use a multiplicative model of speckled image formation and take advantage of the logarithmical transformation in order to convert multiplicative speckle noise into additive noise. The common assumption made in a dominant number of such studies is that the samples of the additive noise are mutually uncorrelated and obey a Gaussian distribution. The present study shows conceptually and experimentally that this assumption is oversimplified and unnatural. Moreover, it may lead to inadequate performance of the speckle reduction methods. The study introduces a simple preprocessing procedure, which modifies the acquired radio-frequency images (without affecting the anatomical information they contain), so that the noise in the log-transformation domain becomes very close in its behavior to a white Gaussian noise. As a result, the preprocessing allows filtering methods based on assuming the noise to be white and Gaussian, to perform in nearly optimal conditions. The study evaluates performances of three different, nonlinear filters—wavelet denoising, total variation filtering, and anisotropic diffusion—and demonstrates that, in all these cases, the proposed preprocessing significantly improves the quality of resultant images. Our numerical tests include a series of computer-simulated and *in vivo* experiments.

I. Introduction

Among the currently available medical imaging modalities, ultrasound imaging is considered to be noninvasive, practically harmless to the human body, portable, accurate, and cost effective. These features have made the ultrasound imaging the most prevalent diagnostic tool in nearly all hospitals around the world. Unfortunately, the quality of medical ultrasound (as defined by image resolution and contrast) is generally limited due to a number of factors, which originate both from physical phenomena underlying the image acquisition and imperfections of the imaging system design. Whereas the latter is remaining a challenge for design engineers, the undesirable physical effects should be compensated by using efficient signal processing tools. As a result, in the past few decades considerable efforts in the field of ultrasound imaging have been directed at development of signal processing techniques intended to combat the main foe of this imaging modality—speckle noise.

Speckle noise is a phenomenon that accompanies all coherent imaging modalities in which images are produced by interfering echoes of a transmitted waveform that emanate from

heterogeneities of the studied objects. The superposition of acoustical echoes coming with random phases and amplitudes tends to produce an intricate interference pattern, known as speckle noise that scales from zero to a maximum, depending on whether the interference is destructive or constructive. Demonstrating little relationship to the macroscopic properties of studied biological tissues, speckle noise tends to obscure and mask diagnostically important details, thereby distracting the diagnosis.

Although speckle noise is a random process, it is not devoid of information. The statistics of the speckle, which generally depend on the microstructure of tissue parenchyma, can be useful for differentiating between either different tissue compositions or types [1], [2]. However, there is no consensus on a unified way to interpret and use this information. However, it is well-known that speckle noise tends to reduce the image contrast, obscure and blur image details, thereby decreasing the quality and reliability of medical ultrasound. As a result, image processing methods for suppressing the speckle noise (which for brevity, will be referred to as despeckling methods) have proven useful for enhancing image quality and increasing the diagnostic potential of medical ultrasound.

Comprehensive analysis of statistical properties of the speckle noise was given a major impetus by the seminal paper of Goodman [3], in which the statistical mechanism of laser speckle formation was first presented. Besides providing basic theoretical results, the study advocates the necessity of rejecting speckle noise via linear filtering to improve the perceptual quality of the images. The results of [3] were subsequently revised in [1], [4], and [5] to account for the specificity of ultrasound imaging. It was recognized that the linear filtering (as it was initially proposed in [3], [5]) is far from being an optimal tool to be used for suppressing the speckle noise because it tends to suppress the noise at the expense of overly smoothing the image details. To perform the filtering, while preserving the anatomical content of the images, adaptive median filters were proposed in [6], [7] (for a comparative analysis of the linear and median filtering, see also [8]). Although these filters are capable of effectively suppressing the speckle pattern, they still seem to remove fine details being actually filters with a low-pass characteristic.

The multiplicative nature of the speckle noise formation was explicitly used in [9] in which the author proposes an algorithm that first converts the multiplicative speckle noise into an additive noise by applying the logarithmic transformation to a speckled image. Subsequently, Wiener filtering is used in order to reject the resultant additive noise, followed by the exponential transformation. The structure of this algorithm is general in the sense that it allows further modification by replacing the linear Wiener filter with other filtering schemes. In particular, discovery of the wavelet transform [10] and fast wavelet decomposition methods [11] had led to wavelet denoising [12], [13] as a powerful method of recovering nonstationary signals. The application of wavelet denoising to the despeckling problem in medical ultrasound imaging was reported in [14]–[16]. The methods, which are based on a multiplicative model of the speckle noise and use the logarithmic transformation to convert the multiplicative noise into an additive one, followed by wavelet denoising, are referred to as the homomorphic wavelet despeckling (HWDS) methods.

Note that the HWDS methods were initially considered in the field of synthetic aperture radar (SAR) imaging in which the first work on this subject seems to be [17]. Since then, many of this type of methods have migrated from the field of SAR imaging to the field of medical ultrasound imaging, using the similarity between the processes of producing the SAR and ultrasound images. A comparative study between the HWDS method and several standard speckle reduction methods, which are largely used by the SAR imaging community, was presented in [18], demonstrating that the former is among the best for

speckle removal. (As to the ultrasound imaging, an analogous conclusion was drawn in [15].)

Despite the impressive results of using the HWDS methods (as reported in [14]–[16]), the present study demonstrates that the performance of such algorithms can be considerably improved via an accurate analysis of statistical properties of the noise to be rejected, and subsequent adjustment of the despeckling scheme. It is shown that the main drawback of the above algorithms stems from considering the log-transformed noise to be white Gaussian noise (WGN). The current study provides evidence that such an assumption is generally oversimplified and leads to inadequate performance of the despeckling. The study shows that, in the most typical practical setting, the noise is neither Gaussian nor white, and it is likely to obey the Fisher-Tippett distribution, implying that the noise is spiky in nature. If not properly treated, the spiky component of such a noise can be processed by a denoising algorithm as a part of the useful signal that needs to be recovered. It allows a significant portion of the noise to be preserved by denoising, thereby considerably decreasing the efficiency of HWDS. Moreover, the present study indicates that the performance of HWDS for ultrasound imaging does not depend on the refinement of a specific wavelet denoising scheme used (that can be achieved via using either an ad hoc thresholding scheme [14] or “fine-spun” statistical priors assumed for the signal's wavelet coefficients [16]), to the same extent as it depends upon the degree to which the denoising scheme is adapted to the particular noise model at hand. As a result, a modified HWDS algorithm is proposed here. The modification is accomplished by adding to the standard HWDS structure a preprocessing stage, which is intended to alter the noise statistics without changing the anatomical content of the image.

The preprocessing consists of two steps. First, a radio-frequency (RF) image is subjected to a spectrum equalization procedure intended to decorrelate the image samples. Second, the log-transformed envelope image is passed through a nonlinear outlier-shrinkage procedure, whose purpose is to suppress the spiky component of the log-transformed speckle noise. It is shown that the preprocessing causes the noise to change in such a way that its behavior becomes similar to that of a WGN model for which the performance of most wavelet denoising schemes is guaranteed to be optimal. By means of comparing the despeckling results, obtained in a series of *in silico* and *in vivo* experiments, the present study demonstrates that the proposed preprocessing procedure results in a considerable improvement of the quality of despeckled images.

In addition, similar results were obtained for two alternative homomorphic despeckling methods derived from the above approach via replacing the wavelet denoising step by total-variation filtering [19] and anisotropic diffusion [20]. Therefore, besides presenting a method for improving the performance of HWDS approaches, the study also demonstrates the applicability of some alternative filtering methods to the despeckling problem in medical ultrasound imaging.

This paper is organized as follows. Section II analyzes the correlation properties of ultrasound images and introduces a whitening procedure for reducing the correlation between the image samples. Some basic properties of speckle noise, as well as the multiplicative model of speckled image formation, are discussed in Section III. The outlier-shrinkage procedure for suppressing a spiky component of the noise in the log-transform domain is also presented in Section III. Section IV provides a brief overview of the despeckling method proposed in the present study. Experimental results are summarized in Section V, and Section VI concludes the paper.

II. RF-Image Modeling and Equalization

A. Image Formation Model in Medical Ultrasound

In order to construct a decorrelation operator for ultra-sound images, an image formation model should be specified first. Assuming linear wave propagation and weak scattering, the back-scattered signal and the tissue reflectivity function¹ are well-known to obey a simple Fourier transform relationship with respect to each other [21]. In this case, an RF-image is considered to be a result of the convolution of the point-spread function (PSF) of the imaging system with the tissue reflectivity function. Denoting by $g(n, m)$, $f(n, m)$, and $h(n, m)$ the RF-image, the tissue reflectivity function, and the PSF, respectively, the convolution model is formally given by:

$$g(n, m) = f(n, m) * h(n, m) + u(n, m), \quad (1)$$

where n and m denote the axial and lateral (or radial and angular, for B-scan sector images) indices of the image samples. The term $u(n, m)$ is added to describe measurement noises as well as all the physical phenomena, which are not accounted for by the convolution model.

It should be noted that the convolution model only approximates the real signal-tissue interaction. Although linear wave propagation is almost always the case, when a moderate amount of acoustical energy is transmitted and tissues with significant attenuation are interrogated, weak scattering generally is not a norm for biological tissues. For instance, in the vicinity of organ boundaries, in which the reflections are typically strong, acoustical reverberations can take place. The latter produce spurious reflectors, thus introducing an error in the definition of the true reflectivity function. However, considering the fact that the regions occupied by strong reflectors are not numerous in regular ultrasound images, the convolution model is known to approximate very closely the real image formation process. This fact has been widely used in numerous methods of ultrasound image reconstruction by deconvolution [22], [23].

The model (1) assumes that the PSF is spatially invariant; an assumption which generally does not hold in practice. Although the variability of the PSF along the lateral direction is primarily due to the changes in the spatial-impulse response of the transducer aperture [24], along the axial direction it also results from the frequency-dependent attenuation [25] and a number of other factors (e.g., nonuniformity of transmission focus, phase aberrations, etc.). Note that, in many cases, the lateral variability of the PSF can be safely neglected, as its effect is largely reduced in modern scanners through dynamically apodizing the transducer aperture. Along the axial direction, however, the PSF variability cannot be compensated by similar means.

Perhaps the simplest way to overcome the problem of the PSF variability, while preserving the translation invariance of the model, is to divide the whole image into a number of (possibly overlapping) segments. If the PSF dependency on the spatial coordinate is sufficiently smooth, one can reasonably assume, that each image segment is formed by convolving the corresponding fragment of the reflectivity function with a local PSF. (Note that such a segmentation can be thought of as an approximation of the spatial dependency of the PSF by a piecewise constant function.) Consequently, the image segments can be processed separately using the model (1) with corresponding local PSF, and, subsequently, the entire image is recovered by combining together the local results obtained in this

¹The tissue reflectivity function accounts for the heterogeneity of the tissue due to density and propagation velocity perturbations, which give rise to the scattered signal. It describes overall reflections in a tissue via defining relative strengths of acoustic reflectors and scatterers as a function of spatial coordinates.

manner. In spite of its simplicity, the above method has proven sufficiently accurate [22], [23], [26], and particularly useful in cases in which fast processing is desired (for more discussions regarding this segmentation, see [26]). Without any loss of generality, the discussion below is focused on an arbitrary image segment, which, with a slight abuse of notation, also will be referred to as an RF-image.

B. Decorrelation of Ultrasound Images

Let $P_g(\omega_1, \omega_2)$, $P_f(\omega_1, \omega_2)$, and $P_u(\omega_1, \omega_2)$ denote the power spectral densities of the RF-image, the tissue reflectivity function, and the additive noise, respectively. Also, let $H(\omega_1, \omega_2)$ denote the Fourier transform of the PSF. Then, provided that the samples of the reflectivity function are independent of the noise, the following spectral relationship takes place:

$$P_g(\omega_1, \omega_2) = P_f(\omega_1, \omega_2) |H(\omega_1, \omega_2)|^2 + P_u(\omega_1, \omega_2). \quad (2)$$

Due to the natural intricacy of most biological tissues and the fact that tissue heterogeneity is generally formed by numerous small “independent” structures, the samples of the reflectivity function can reasonably be assumed to be uncorrelated [22], [23]. In this case,

$P_f(\omega_1, \omega_2)$ is a constant function with its amplitude equal to the variance σ_f^2 of $f(n, m)$. Moreover, in many cases of practical interest, the samples of the additive noise $u(n, m)$ can be reasonably assumed to be uncorrelated as well, so that the power spectral density $P_u(\omega_1, \omega_2)$ is constant and equal to the noise variance σ_u^2 . The above assumptions allow simplifying (2) to the following form:

$$P_g(\omega_1, \omega_2) = \sigma_f^2 |H(\omega_1, \omega_2)|^2 + \sigma_u^2. \quad (3)$$

One can see that the power spectral density of $g(n, m)$ is nothing else, but the power spectrum of the PSF plus the noise term, and, hence, the autocorrelation of the RF-image is completely defined by the autocorrelation of the PSF. The latter is well-known to have a non-negligible support. The spatial extent of the PSF along the lateral direction is defined by a non-negligible width of the acoustical beam, and along the axial direction, it is defined by the Q-factor of the ultrasound transducer. Non-negligibility of the correlation between the speckle noise samples is further illustrated in Fig. 1. The upper subplot of the figure shows a segment of a B-scan image of the liver of an adult volunteer with normal liver functions. The segment has been cropped in such a way that it does not contain visible organ structures, thereby presenting an almost homogeneous field of stationary speckle noise. In this case, the autocorrelation of the image represents the correlation of the speckle noise field. The axial and lateral profiles of the autocorrelation function are shown in the lower two sub-plots of Fig. 1. The fact that the autocorrelation function has non-negligible support along both axes indicates considerable correlation between the speckle noise samples.

The above considerations provide evidence that ultrasound speckle noise cannot be assumed and, subsequently, dealt with as a white noise process; and its correlation properties must be properly taken into account [27], [28]. Alternatively, one can try to find an operator that can transform a received RF-image into *another* RF-image, whose samples correlate less than those of the original. Following [29], we define this operator to be a linear filter $L(n, m)$, whose transfer function is given by:

$$L(\omega_1, \omega_2) = \left(|H(\omega_1, \omega_2)|^2 + \sigma_u^2 / \sigma_f^2 \right)^{-1/2}. \quad (4)$$

One easily can see that applying (4) to an RF-image, the power spectral density of which is defined by (3), results in “flattening” the latter and, therefore, reduces the correlation between the RF-image samples. The constant $\varepsilon = \sigma_u^2 / \sigma_f^2$ can be thought of as a tunable parameter of the decorrelation that controls amplification of the “out-of-band” frequencies of $g(n, m)$, which have been damped due to the band-limitedness of $H(\omega_1, \omega_2)$. In order to compute an optimal value of ε , the variances of the reflectivity function and of the noise need to be estimated first, and this could be done using, for instance, the methods reported in [30] and [31], respectively. In practice, however, we found it quite acceptable to set this parameter empirically, so that its optimal value would result in maximal decorrelation while avoiding any undesirable artifacts caused by “overamplification” of the high frequencies.

Although the decorrelation filter (4) has been defined for RF-images, it also is applicable to demodulated, in-phase/quadrature (I/Q) images (i.e., the images whose absolute value is what eventually viewed on the displays of most ultrasound systems). Due to the linearity of the frequency demodulation process, the I/Q-image still can be modeled as a convolution mixture of a complex reflectivity function with a complex PSF. Therefore, (4) can be used for decorrelating the I/Q-images with $|H(\omega_1, \omega_2)|$ being the amplitude of the Fourier transform of the complex PSF. Moreover, as the demodulation is typically followed by an anti-aliasing filtering and down-sampling, processing the I/Q-images is advantageous due to reduction in both noise level and number of samples.

In order to implement the spectrum equalization using the filter (4), the power spectrum of the PSF needs to be estimated first. In the current study, the estimation is performed using the method of [26]. Below, a brief overview of this method is provided so as to render the presentation self-contained.

C. Estimation of the PSF Spectrum

Let $G(\omega_1, \omega_2)$, $F(\omega_1, \omega_2)$, and $H(\omega_1, \omega_2)$ denote the log-magnitude of the Fourier transforms of the I/Q-image, the complex tissue reflectivity function, and the complex PSF, respectively. Disregarding for the moment the noise term in (1), the convolution model implies:

$$G(\omega_1, \omega_2) = H(\omega_1, \omega_2) + F(\omega_1, \omega_2), \quad (5)$$

which suggests that $H(\omega_1, \omega_2)$ could be estimated from $G(\omega_1, \omega_2)$ by rejecting the “noise” $F(\omega_1, \omega_2)$. Thus, the problem of estimating the power spectrum of the PSF is basically a filtering problem. However, before proposing a specific filtering method, the statistical properties of the noise should be carefully analyzed.

It was demonstrated in [32] that, when samples of the reflectivity function behaves as a WGN, the samples of $F(\omega_1, \omega_2)$ are *i.i.d.* and obey the Fisher-Tippett distribution, whose *pdf* is given by:

$$p_Y(y) = 2 \exp \left\{ \left(2y - \ln 2\sigma_f^2 \right) - \exp \left\{ 2y - \ln 2\sigma_f^2 \right\} \right\}, \quad (6)$$

here, as before, σ_f^2 denotes the variance of the reflectivity function. Unfortunately, the above statistical description cannot in general be applied to characterize the reflectivity function corresponding to whole RF-image. Yet, it is rarely a problem to detect within a given RF-image a smaller fragment, within which the reflectivity function behaves like a WGN. Such detection can be done using, e.g., the Kolmogorov-Smirnov testing as proposed in [30].

It is interesting to note that the Fisher-Tippett distribution possesses an approximant in the form of a Gaussian *pdf*. The latter is obtained by replacing the inner exponent in (6) by the first three terms of its series expansion, resulting in:

$$p_f(y) \simeq 2e^{-1} \exp \left\{ -\frac{1}{2} \left(\frac{y - \ln \sqrt{2}\sigma_f}{0.5} \right)^2 \right\}. \quad (7)$$

An interesting fact about the approximation above is that it has a constant variance of 0.25, implying that the additive noise in (5) may be roughly viewed as WGN with a fixed variance. Thus, from the viewpoint of filtering, the problem can now be stated as the requirement to cancel “almost” WGN, whose variance never changes.

Unfortunately, the approximation (7) is acceptable only in close proximity of the mean value of the original *pdf* (6). Moreover, as compared to the Gaussian, the *pdf* of the Fisher-Tippett distribution is asymmetric (its skewness is equal to $12 \cdot 6\zeta/\pi^3$, where ζ is the Apery's constant) and leptokurtic (its kurtosis is equal to 12/5). Fig. 2 exemplifies the *pdf* given by (6) for the case $\sigma_f^2=1$ (upper subplot) and a realization of corresponding noise (lower subplot). One can see that, although the right-hand side of the *pdf* has a form similar to that of a Gaussian *pdf*, its left-hand side has a long, heavy tail. It implies that white noise produced by the Fisher-Tippett distribution will be similar to WGN, except for a relatively small number of relatively large-amplitude samples, which appear to project from the main ensemble (see the lower subplot of Fig. 2). As a result, such a noise may be viewed as a WGN contaminated by occasional transients or outliers.

The above considerations imply that the noise to be rejected is of spiky type. It is generally known that spiky noise is difficult to deal with, and many methods, which exploit the concept of L_2 -projections (e.g., wavelet denoising), often fail to reject such a noise in a satisfactory manner. The main reason for this is the fact that the noise outliers are recognized by such a filtering scheme as features of the signal to be recovered and, as a result, are preserved.

In order to overcome the difficulty of rejecting the outliers, it was proposed in [32] to “Gaussianize” the noise via estimating and subsequently subtracting its spiky component. The latter can be estimated as robust residuals of $G(n, m)$ computed according to:

$$R(n, m) = \text{sign}(\Delta G(n, m)) (|\Delta G(n, m)| - \lambda)_+, \quad (8)$$

here ΔG denotes difference between G and its median-filtered version, λ is a predefined threshold, and the operator $(x)_+$ returns x if $x > 0$ and zero otherwise. It was observed that, in most cases, the robust residuals R correspond to the outliers of the spiky noise, when the size of the median filter is set to be 3×3 (or 5×5) and the threshold λ is set to a level such that 93–95% of the differences $|\Delta G|$ do not exceed the predefined threshold λ . In this case, subtracting R from G results in suppressing the spiky component of the Fisher-Tippett noise. Moreover, it was demonstrated in [32] that the noise contaminating the difference signal ($G - R$) behaves very similarly to WGN, and the desired signal H remains practically unchanged. Note that the above-described procedure of computing and subsequently subtracting the robust residuals is known as outlier-shrinkage, and it was originally proposed as a part of outlier-resistant wavelet denoising in [33].

Once the spiky component of F in (5) has been rejected, the signal ($G - R$) can be filtered in order to estimate H . Following [26], [32] the latter is recovered using wavelet denoising

[12]. This method was chosen due to its very nice characteristic of being capable of rejecting WGN without oversmoothing the resulted estimates. In this paper, a separable wavelet transform [11] based on the nearly symmetric wavelet [10] with six vanishing moments was used to perform the denoising. The thresholding rule was chosen to be soft thresholding (see [12] or the formal definition in the section that follows), with the threshold defined using the theoretically predicted variance of the Gaussian approximation in (7), i.e., 0.25.

Having estimated the power spectrum of the PSF, the decorrelation filter (4) can be readily computed and applied to acquired I/Q-images. Throughout the following sections it will be tacitly assumed that the I/Q-images have been subjected to the decorrelation procedure, and, consequently, samples of the corresponding envelope images are nearly uncorrelated.

III. Speckle Noise and Outlier Shrinkage

A. Generalized Model of Speckled Images

The most critical part of developing a method for recovering a signal from its noisy measurement consists in choosing a reasonable (either statistical or analytical) description of the physical phenomena underlying the data formation process. The extent to which the chosen model succeeds to account for the observed physical effect often defines the reliability of the reconstruction algorithm as well as the precision with which the signal is recovered. Consequently, the availability of an accurate and reliable model of speckle noise formation is a prerequisite for the development of a useful despeckling algorithm.

In ultrasound imaging, however, a universally agreed upon definition of such a model still seems to be lacking. Nevertheless, a number of possible formulations, whose feasibility was verified via their practical use, exist. A generalized model of the speckle imaging as proposed in [9] and used, e.g., in [14], [16] is given by:

$$g(n, m) = f(n, m) u(n, m) + \xi(n, m), \quad (9)$$

where g , f , u , and ξ stand for the observed envelope (not I/Q or RF) image, original image, multiplicative and additive components of the speckle noise, respectively. Here the indices n and m denote the axial and lateral indices of the image samples (or, alternatively, the angular and range indices for sector images).

Despite its possible theoretical shortcomings [34], the model (9) has been successfully used both in ultrasound and SAR imaging. Moreover, evidence exists that, when applied to ultrasound images, only the multiplicative component u of the noise must be reckoned with, and, hence, (9) can be considerably simplified by disregarding the additive noise term. This leads to the following simplified model:

$$g(n, m) \simeq f(n, m) u(n, m). \quad (10)$$

Note that there exists an alternative model, as proposed in [6] and used in [15], [35], describing the speckle noise as an additive noise, with its amplitude proportional to square root of the true image. However, this model was proposed to account for the speckle pattern, as it appears “on screen,” i.e., after a sequence of standard processing steps performed by a typical ultrasound scanner (e.g., nonlinear amplification, dynamic-range adjustment via log-compression, etc.). Consequently, adopting (10) as the basic model, it is assumed that the image $g(n, m)$ is observed before the system processing is applied.

Homomorphic despeckling methods take advantage of the logarithmic transformation that, when applied to both parts of (10), converts multiplicative noise into additive noise.

Denoting the logarithms of g , f , and u by g_l , f_l , and u_l , respectively, the measurement model in the log-transform domain becomes:

$$g_l(n, m) = f_l(n, m) + u_l(n, m). \quad (11)$$

At this stage, the problem of despeckling is reduced to the problem of rejecting an additive noise, and a variety of noise suppression techniques could be invoked in order to perform this task. However, before proposing a specific denoising method, it is instructive to take a closer look at the properties of the noise term in (11). Note that most of the homomorphic despeckling methods proposed so far simply assume the noise to be WGN. However, such an assumption may be a serious limitation, as shown by the considerations that follow.

B. Statistics of Log-Transformed Speckle Noise

The mechanism of the speckle formation in ultrasound imaging is similar to the laser imaging [3], and the statistical description of the speckle noise generally depends upon tissue composition and type. For the case in which the resolution cell consists of a relatively large number of independent scatterers (more than 10, normally), the image amplitude is widely recognized as possessing a Rayleigh distribution [1]. However, when either the number of scatterers is low or their spatial locations are not independent, the statistics are likely to deviate from the Rayleigh model [36]. In order to account for the non-Rayleigh scattering, a number of distributions have been proposed, including the K-distribution [37], [38], the Nakagami distribution [39], the Weibull distribution [40], and Generalized Gamma (GG) distribution [41], [42]. When a structure of specular reflectors, which produce the coherent portion of backscattered energy, is superimposed on a background of relatively weak diffuse scatterers, the resulted images seem to obey a Rician distribution [43].

For the reasons of space, it is not possible to address all of the above cases. Thus, in order to keep the discussion as general as possible, the speckle noise $u(n, m)$ in (10) is assumed to obey the GG distribution, whose *pdf* is given by:

$$p_z(z) = \frac{\gamma z^{(\gamma\nu-1)}}{\alpha^{\gamma\nu} \Gamma(\gamma)} \exp\left\{-\left(\frac{z}{\alpha}\right)^\gamma\right\}, \quad z \geq 0, \alpha, \nu, \gamma > 0. \quad (12)$$

The GG distribution is especially attractive because it contains several distributions as special cases, viz. Rayleigh ($\nu = 1$, $\gamma = 2$), exponential ($\nu = 1$, $\gamma = 1$), Nakagami ($\gamma = 2$), Weibull ($\nu = 1$), and log-normal ($\nu \rightarrow \infty$). Assuming the GG distribution implies that the noise samples u_l in (11) are distributed with *pdf* given by:

$$p_Y(y) = \frac{\gamma}{\Gamma(\gamma)} \exp\{\gamma\nu(y - \ln \alpha) - \exp\{\gamma(y - \ln \alpha)\}\}. \quad (13)$$

One can see that the distribution in (13) is of the two-exponential type, being very close in form to the Fisher-Tippett distribution [44] as given by (6).² Fig. 3 depicts a number of *pdf* given by (13), computed for different parameters of the GG distribution, which were set according to the experimental results of [42]. One can see that, for a fairly wide range of parameters, the shape of the *pdf* in (13) resembles that of the Fisher-Tippett *pdf*, with all the

²In fact, the pdf in (13) attains the precise analytical form of the Fisher-Tippett *pdf*, when ν approaches 1 (the case of the Rayleigh distributed speckle noise).

implications discussed in Section II-C. In particular, it implies that the noise in (11) is likely to be similar to a WGN contaminated by a relatively small number of outliers.

In order to overcome the nontrivial problem of rejecting the spiky noise u_k , it is possible to make it be more amenable to filtering methods, which are based on the assumption that the noise is WGN, using the same outlier-shrinkage procedure introduced in Section II-C in connection with the PSF spectrum estimation. Specifically, we propose to subtract from $g(n, m)$ its robust residuals [computed according to (8)], before a filter is applied to $g(n, m)$ to reject $u(n, m)$. Together with the decorrelation procedure of Section II, the “Gaussianization” of the log-transformed speckle noise $u(n, m)$ using the robust-shrinkage forms the core idea of the modified homomorphic despeckling, the overall structure of which is summarized below.

C. Modified Homomorphic Despeckling

The overall block-diagram of the proposed processing scheme is depicted in Fig. 4. It suggests that the logarithm of the absolute value of a decorrelated I/Q-image is passed through the outlier-shrinkage stage before it is subjected to a certain noise-reduction procedure. Note that the term denoising is used here as a substitute for arbitrary filtering. It is important to emphasize that the homomorphic despeckling methods proposed so far do not perform the decorrelation and the outlier-shrinkage of the log-envelope image, and thus their performance could be described by the block-diagram similar to that in Fig. 4, but devoid of the above preprocessing stages.

Before proceeding to the discussion on possible filtering approaches and demonstrating some experimental results, it would be instructive to assess the effect of the proposed preprocessing via an example. The leftmost sub-plot of Fig. 5 shows the original envelope image of a human (right) kidney acquired by a VIVID3 (GE Medical Ultrasound, Inc., Tirat Carmel, Israel) commercial ultra-sound scanner. Note that the image was acquired using a curved array transducer and, hence, a scan conversion procedure should have preceded the visualization to preserve morphological consistency of the image. However, the scan conversion stage was omitted here (and, thus, the image is shown in the “angle-depth” coordinate system), as it could have altered the image correlation properties, which are central for purposes of the example. The middle subplot of Fig. 5 shows the standard envelope image after applying the log-transformation. Note that such an image would be considered as an input to the denoising stage by most of the existing homomorphic despeckling algorithms. The rightmost subplot of Fig. 5 shows the log-envelope³ of the same image after it is passed through all stages of the block-diagram in Fig. 4. One can see that, although the standard and the preprocessed log-envelopes are similar from the viewpoint of anatomical structures they present, the speckle pattern of the preprocessed image is much finer than that of the standard image, implying significant loss of the correlation between the image samples. This fact is further verified via comparing the autocorrelation functions of the standard and of the preprocessed log-envelopes, which are shown on subplots A and B of Fig. 6, respectively. Note that the autocorrelation functions were computed using homogenous (“structure-free”) regions of the images to exclude the effect of inter-pixel dependencies produced by the structured image content. One can see that the autocorrelation of the preprocessed log-envelope decays in a much faster rate, as compared to the autocorrelation of the standard log-envelope. This fact indicates considerable loss of the correlation between the samples of the preprocessed image.

³The term “log-envelope” here and hereof is used as a shorthand substitute for “logarithmically transformed envelope image”.

Moreover, one can see that the preprocessed log-envelope possesses considerably better contrast as against the standard log-envelope. As prior to the visualization, the values of both log-envelopes were normalized to lie within the same range (viz. $[0, 1]$), it is reasonable to assume that the contrast improvement has been caused by suppression of the noise outliers. In order to confirm this assumption, the histograms of both log-envelopes were computed and compared. The histogram of the standard log-envelope is shown on subplot C of Fig. 6. One can see that this histogram agrees well with the theoretical model of (13). Furthermore, its heavy left tail implies that the standard log-envelope is contaminated by a noise having a spiky behavior. However, the histogram of the preprocessed log-envelope, which is shown on subplot D of Fig. 6, is shaped more like a Gaussian *pdf*, thereby indicating that the spiky component of the noise has been effectively rejected. The implications of these results are demonstrated via the reconstruction examples given below.

IV. Filtering Methods

The original impetus for the present study was the desire to improve performance of HWDS via introducing the preprocessing method described above. However, in addition to wavelet denoising, there exist a number of alternative filtering methods that may result in reconstructions of comparable quality. In this section, a number of alternative despeckling approaches are presented, which are obtained via substituting different filters at the “denoising” stage of the block-diagram in Fig. 4.

A. Wavelet Denoising by Soft-Thresholding

In [12], a very simple thresholding procedure for the recovery of functions from noisy data was proposed. It consists of three steps: the signal is transformed into an orthogonal domain, using a discrete wavelet transform producing empirical wavelet coefficients. The empirical wavelet coefficients are subjected to nonlinear soft-thresholding $\eta_\lambda(y) = \text{sign}(y) (|y| - \lambda)$, with a threshold $\lambda = \sqrt{2 \log(n)} \sigma$, where σ is the standard deviation of the white noise and n is the data length. The thresholded wavelet coefficients are inversely transformed, supplying an estimation of the true signal. The above scheme is known as wavelet denoising, and the thresholding rule is known as the uniform soft thresholding.

The three-step reconstruction procedure mentioned above was shown to minimize the estimation error (which, in this case, achieves almost the minimax error for a magnitude of important smoothness classes) subject to an additional constraint requiring that the estimate is at least as smooth as the function to be recovered. As a result, denoised images generally are much less oversmoothed, in comparison with the images denoised by, e.g., linear filtering. Note that uniform soft thresholding is not the only way to suppress the wavelet coefficients of the noise, and a multitude of various methods have been proposed based, for example, on principles of Bayesian estimation and detection theory [45]. In most cases, these methods were shown to outperform the soft-thresholding. However, because the purpose of this paper is not really concerned with comparing various thresholding schemes, the original approach of [12] is used here. The noise variance needed for definition of the threshold was estimated by assuming, that most empirical wavelet coefficients at the finest level of the decomposition are induced by the noise, and, thus, the median absolute deviation of wavelet coefficients at this level accurately reflects the noise size [12].

B. Total Variation Filter

Let $\mathcal{I}(n, m)$ denote a noise-contaminated version of the original image $\mathcal{I}(n, m)$. Also, let $\mathbf{D}_x$ and $\mathbf{D}_y$ denote the approximations to the first order partial derivative operators. Then, for a predefined $\lambda > 0$, a discrete version of the total variation filter, as originally specified in [19], recovers $\mathcal{I}(n, m)$ by solving:

$$\min_I \left\{ E(I) \triangleq \|I - J\|_F^2 + \lambda \sum_{n,m} \left(|\mathbf{D}_x \{I(n, m)\}|^2 + |\mathbf{D}_y \{I(n, m)\}|^2 \right)^{1/2} \right\}, \quad (14)$$

where $\|\cdot\|_F$ stands for the Frobenius matrix norm. The filter (14) is now considered to be among the most successful methods for image restoration and edge enhancement. It is mainly because of its capability of filtering out the noise without blurring the most universal and crucial features of images: edges. Because of the specific form of the regularizing (second) term of the functional in (14), the total-variation filter is especially useful for recovering piecewise constant signals.

Further, note that the regularization parameter λ controls the balance between noise removal and smoothing. Too large a value for λ tends to yield a smooth solution of poor fidelity to the data, and too small a value provides results that could be insufficiently smooth. In this paper, the optimal value for λ was set experimentally to achieve the most visually pleasing result.

The total-variation filter was implemented by solving the minimization problem (14) using the conjugate gradient algorithm [46]. In order to overcome the problem of nondifferentiability of $E(I)$, a small positive number (usually of the order 10^{-4}) was added under the square root of the second term in (14), and, subsequently, the resultant functional was minimized. Note that the total-variation filter also can be implemented as a signal-dependent filter, as described in [47] [such a formulation stems from an explicit discretization of the Euler-Lagrange equation corresponding to continuous version of (14)].

C. Anisotropic Diffusion

Another approach to the filtering problem takes advantage of the locality and anisotropy of certain partial differential equations. Among all differential operators, the diffusion class is the most widely applied in current image analysis. Because linear homogeneous diffusion may smooth out noise successfully only at the expense of overly blurring out significant sharp details of images (e.g., edges), anisotropic diffusion has attracted much attention [48]. It was found that operators of this class are capable of smoothing images without blurring the boundaries between their homogeneous regions. One choice is to use the following affine invariant anisotropic smoothing filter [20]:

$$\frac{\partial I(x, y, t)}{\partial t} = \left[\left(\frac{\partial I}{\partial y} \right)^2 \frac{\partial^2 I}{\partial x^2} - 2 \frac{\partial I}{\partial x} \frac{\partial I}{\partial y} \frac{\partial^2 I}{\partial x \partial y} + \left(\frac{\partial I}{\partial x} \right)^2 \frac{\partial^2 I}{\partial y^2} \right]^{1/3}, \quad (15)$$

here $I(x, y, t)$ represents the image to be filtered, which is now considered to be a function of two spatial coordinates x and y , as well as of time t . It can be shown, that the above equation involving only the first and second order spatial derivatives of the image I defines the affine geometric heat flow, under which the level sets of I undergo affine curve shortening. Moreover, such a diffusion process has the desirable characteristics of preserving edges while exhibiting numerical stability and straightforward computation [20]. The time discretization step and the number of iterations were used as parameters of the nonlinear smoothing, and they were adjusted to achieve the best possible visual results.

V. Experimental Results

A. In Silico Experiments

Simulation studies are usually the first validation step used to examine the performance of an estimation method in a quantitative way. In the current study, two-dimensional (2-D) RF-

images were simulated according to the model of (1) using the PSF that was measured by imaging a point-target (viz., a thin steel wire in a water tank) using a single-element, 3.5 MHz-transducer (Pana-metrics V383, Waltham, MA) for both transmission and reception. The lateral scanning of the target was carried out mechanically with a lateral resolution of 0.4 mm, and the acquired RF-lines were sampled at a rate of 25 MHz.

The tissue reflectivity functions were generated as 2-D WGN fields weighted by the amplitude profile shown on the left subplot of Fig. 7. Note that the resulted reflectivity functions had been “designed” to mimic a fragment of homogeneous tissue containing four round regions with relatively low reflectivity (e.g., cysts or blood vessels). In this case, the despeckling is expected to reconstruct the original tissue profile, i.e., the piecewise constant function shown on the left subplot of Fig. 7. An example of the simulated (original) envelope image is shown on the right subplot of Fig. 7.

The despeckling was first performed using the wavelet denoising as a method for canceling the additive noise term in (11). As before, this approach is referred to as HWDS. Each RF-image was processed twice, viz. with and without the decorrelation and the outlier-shrinkage stages, as depicted in the block-diagram in Fig. 4. In the first case, the processing is referred to below as modified, whereas in the second case as standard. The wavelet denoising was implemented using the WaveLab® package (Department of Statistics, Stanford University) of D. Donoho (<http://www-stat.stanford.edu/~wavelab/>). Four-level wavelet decomposition based on nearly symmetric wavelets of I. Daubechies [10] with six vanishing moments was used to this end.

The next set of results was obtained using the total-variation filter (14) at the denoising stage of the despeckling. The corresponding despeckling method is referred to below as the total variation despeckling (TVDS) (both standard and modified, by analogy to the previous case). The regularization parameter λ of TVDS was set to be 1.2.

Finally, the denoising was performed using the standard and modified versions of the anisotropic diffusion despeckling (ADDS) algorithm that was obtained by using the anisotropic diffusion filter of Section IV-C at the denoising stage. Diffusion filtering in the ADDS approach was implemented using 50 iterations with the time discretization step of 0.1.

A representative result is demonstrated in Fig. 8 in which subplots A1, A2, and A3 show the reconstructions obtained by the standard versions of the HWDS, TVDS, and ADDS algorithms, respectively, whereas the corresponding images in the right column (i.e., subplots B1–B3) were obtained using their modified versions. One can see that, in the case of the standard despeckling, all the recovered images poorly represent the homogeneity structure of the underlying “tissue”, and they are hardly more informative than the standard envelope image shown in Fig. 7. However, the modified despeckling provides useful reconstructions, representing the original homogeneity structure in a considerably better way.

The difference in performances of the standard and modified despeckling algorithms can be further accentuated by comparing the intensity values of the despeckled images taken along a line segment, with corresponding true homogeneity profile. The upper subplot of Fig. 9 shows the intensity values of an original envelope image taken along the line passing through the center of the upper “cyst” in the axial direction (solid line) together with the corresponding true homogeneity profile (dotted line). However, subplots A1–A3 show the same intensity profile obtained after processing the image using the standard versions of the HWDS, TVDS, and ADDS algorithms, respectively. The corresponding intensity profiles in the right column (i.e., subplots B1–B3) were obtained using the modified versions of the

algorithms. The composition of Fig. 10 is identical to that of Fig. 9 with the only difference being that the line segment here passes through the center of the upper “cyst” in the lateral direction. One can see that in all the cases, the modified despeckling succeeds in recovering the true homogeneity profiles almost perfectly; but in the case of the standard despeckling, the reconstructed profiles barely resemble the true ones.

For the quantitative assessment, a number of performance measures were used to compare the despeckling methods. The first measure was the normalized mean-squared error (NMSE). Denoting by I_{org} and I_{est} the original image to be recovered and its estimate, respectively, the NMSE can be defined as:

$$NMSE = E \left\{ \frac{\|I_{org} - I_{est}\|_F}{\|I_{org}\|_F} \right\}, \quad (16)$$

where $\|\cdot\|_F$ stands for the Frobenius matrix norm. It is worthwhile noting that both NMSE and its reciprocal (that is also known as the signal-to-MSE ratio [18]) are commonly used in the coherent imaging in which the standard definition of the signal-to-noise ratio might be inadequate, because of the multiplicative nature of speckle noise.

Additionally, in order to evaluate the degree of contamination of the images by speckle noise, the conventional speckle-SNR was used, which is defined as a ratio of the mean to the standard deviation of speckled images. Note that, in the case of fully developed (Rayleigh distributed) speckle noise, this ratio is known to be approximately equal to 1.91 [1]. After applying a speckle reduction algorithm, the ratio is expected to increase.

The last measure was indented to assess the ability of despeckling methods to preserve sharp details of the images. If ΔI_{org} and ΔI_{est} denote approximations of the Laplacians of the original image and its estimate, respectively, then this performance measure is given by [15]:

$$\beta = E \left\{ \frac{\langle \Delta I_{org}, \Delta I_{est} \rangle}{\|\Delta I_{org}\|_F \|\Delta I_{est}\|_F} \right\}, \quad (17)$$

where $\langle \cdot, \cdot \rangle$ denote the standard inner product. The closer the index β is to 1, the better is the ability of despeckling to preserve the image edges.

Table I summarizes the quantitative results obtained in the simulation study. Note that the expectations in [16] and [17] were estimated by corresponding sample means based on results of 100 independent trials. One can see that the proposed preprocessing results in reducing the NMSE by the factor of 1.76, 1.71, and 1.53 in the case of HMDS, TVDS, and ADDS, respectively. In addition, the images, which were despeckled using the proposed preprocessing procedure, have the speckle-SNR that is, on average, two times higher than that of the images despeckled without the preprocessing. Analyzing the β index indicates that the preprocessed images better represent the edges of the original scene.

B. In Vivo Experiments

In vivo data were acquired next in order to evaluate the performance of the despeckling methods. A set of RF-images was recorded from adult volunteers using the VIVID3 scanner. The scanning was performed using linear transducer array with a central frequency in the vicinity of 3.5 MHz. A set of 10 different images of the carotid arteries of the volunteers was used for evaluating the algorithms. The images were acquired with a single transmission focal point, localized in the center of the field of view. All the RF-images were composed of

256 RF-lines, each of 1024 points in length. The sampling rate and resolution were 20 MHz and 16 bits, respectively. Each image was divided into three quasi-stationary segments along the axial direction, according to the methodology of Section II-A. Parameters of the despeckling algorithms were chosen to be the same as in the simulation study of the preceding subsection.

An example of the (original) longitudinal view of a fragment of the carotid artery is shown in the upper subplot of Fig. 11⁴. However, the right column of subplots of the figure shows this image after applying to it the standard versions of the (from top to bottom) HWDS, TVDS, and ADDS algorithms. The corresponding images in the left column were obtained using the modified versions of these algorithms. Comparing the images, it seems that each despeckling method does a reasonable job of enhancing the structure of interest. However, the modified methods seem to provide much more noiseless results as compared to their standard versions, while preserving all the fine structures in each case. Moreover, the modified methods better represent homogeneous regions of the underlying tissue that appear here more uniform. Additionally, one can see that most of the edge-like structures (e.g., the intima of the carotid artery) are better represented in the case of modified solutions, which have obviously better contrast, thereby better representing the overall structure of the tissue.

The difference from method to method lies in the nature of the smoothing. For example, in the case of the total variation filtering one gets a more piecewise constant effect compared to a wavelet-based de-noising as expected from standard theorems in the literature (see [49] and the references therein)⁵. Because in each case the modified methods bring out the desired structure clearly, we suppose it is a problem-dependent matter, which smoothing methodology should be chosen. Generally, one can see that all anatomical structures in the images processed using the modified despeckling appear considerably less noisy than in the “standard” images.

The ability of different despeckling methods to recover the homogeneous areas of tissue can be further appreciated via analyzing the results shown in Fig. 12, whose composition is analogous to that of Fig. 11. The subplots of Fig. 12 demonstrate a fragment of a human bladder imaged using the VIVID-3 scanner with a curved array transducer. Although each kind of despeckling has its own “fashion” to reject the speckle noise (with the “burnishing,” “fragmentizing,” and the “worm-effect” peculiar to HDWS, TVDS, and ADDS, respectively), one can see that, in all the cases, the modified processing provides much smoother estimates, without overly smoothing their morphological structures (e.g., the organ boundaries, blood vessels, etc.).

Unfortunately, in the *in vivo* case, quantitatively assessing the performance of the despeckling methods is problematic because of the absence of corresponding original images. As a result, only two performance measures were used here for comparison. The first measure was the speckle-SNR defined in the preceding subsection, and the second measure was defined as a ratio of the number of pixels of the image autocorrelation function, which exceed 75% of its maximum value to the total number of pixels. Note that the latter measure (denoted below by α) is typically used in order to evaluate the resolution in ultrasound imaging [23].

The quantitative results obtained during the *in vivo* study are summarized in Table II. One can see the significant improvement in the speckle-SNR after applying the proposed

⁴Note that all images in the figure were normalized and subsequently compressed for visualization in 8-bit resolution, so that they have the same dynamic range.

⁵Note that in order to suppress the “Gibbs-like” artifacts, which frequently take place in the case of denoising using undecimated wavelet transforms, the cycle-spinning scheme of [50] was used here.

preprocessing, implying that the preprocessed images are superior in representing the homogeneous structure of the studied tissues. At the same time, the resolution of despeckled images is considerably worse in comparison with that of the nonprocessed images. Yet, the immediate conclusion that despeckling tends to deteriorate the resolution would not be quite right. It is because of the fact that a relatively low correlation within speckled images is primarily due to their noisy nature. Having been subjected to a despeckling procedure, the ultrasound images become considerably less noisy; and, as a result, their autocorrelation functions begin to represent the correlation within the original image rather than the correlation between the noise samples. Therefore, the measure α in Table II is likely to represent the degree of speckle noise contamination, with higher values of α corresponding to more clear images.

VI. Discussion and Conclusions

A new method for improving the performance of homomorphic despeckling methods has been presented. The fundamental idea underpinning this class of speckle reduction techniques consists in using the log-transformation in order to convert multiplicative speckle noise into an additive noise process, followed by suppressing the latter using certain filtering procedures. The present study has demonstrated conceptually and experimentally that assuming the additive noise to be a WGN (as it is done in most cases) can lead to inadequate performance for a number of despeckling algorithms of this kind.

Consequently, a simple preprocessing procedure was proposed in the present study. Its distinctive feature lies in the fact that it does not modify the structure of a specific filtering method, but rather alters the noise in such a way that it becomes very similar in behavior to WGN. Because a number of powerful filtering methods exist, which are based on assuming the noise to be a WGN, the proposed “noise-correction” procedure allows them to perform under nearly optimal conditions.

The proposed preprocessing procedure can be viewed as an “add-on” for existing homomorphic despeckling schemes. It consists of two simple stages: a received I/Q-image is passed through a spectral equalization stage that is intended to reduce the correlation between the image samples; the log-envelope of the decorrelated I/Q image is subjected to the outlier-shrinkage process suppressing the spiky component of the additive noise. After that, any filtering can be applied for the noise rejection. Although only three denoising methods have been examined in the present study, it is believed that the proposed preprocessing may be beneficial for many other filtering methods as well.

The performance of three homomorphic despeckling methods, as defined by three different denoising techniques (viz. wavelet denoising, total-variation filtering, and anisotropic diffusion filtering) was examined in the current study. In all the cases, the despeckling results obtained with and without the preprocessing were compared. It was shown in a series of computer-simulated and *in vivo* experiments that in all the cases, the proposed preprocessing results in remarkable improvement in the quality of resulted despeckled images. Compared to the standard despeckling approaches, the results obtained using the proposed preprocessing procedure are shown to be significantly less noisy, and to have higher contrast, thereby better representing the anatomical structures of interrogated tissue (see the results summarized in Tables I and II).

As the primary purpose of the study is to present a method for enhancing the performances of currently existing homomorphic despeckling techniques, no comparison with more classical speckle reduction methods (e.g., median filtering) was done. Such comparison can be found in virtually all the studies on homomorphic despeckling, which recognizes the latter to be among the best “performers”.

Note that no attempt was made to compare the results obtained by different despeckling methods. Such a comparison, including evaluation of these methods from a different number of viewpoints (e.g., computational efficiency, reliability of recovering different anatomical structures, and different tissue morphologies) well deserves a future study.

Acknowledgments

The authors would like to thank Prof. Dan Adam from the Technion-Israel Institute of Technology for the ultrasound data kindly provided to them. The authors also would like to thank the anonymous reviewers, whose valuable comments have helped to considerably improve the quality of the material presented in the paper.

This work was partially funded by grants from AFOSR, NIH, ARO, MRI-HEL, and NIH (NAC and NAMIC through Brigham and Women's Hospital).

Biography

Oleg Michailovich was born in Saratov, Russia, 1972. He received the B.Sc. degree and M.Sc. degree in electrical engineering from the University of Saratov, Russia, in 1994, and the M.S. degree and Ph.D. degrees in biomedical engineering from the Technion-Israel Institute of Technology in 2003. He is currently a research scientist in the School of Electrical and Computer Engineering at the Georgia Institute of Technology. His research interests include the application of image processing to various problems of image reconstruction, in verse problems, non-parametric estimations, approximation theory and multiresolution analysis.

Allen Tannenbaum (M'93) was born in New York City in 1953. He received his Ph.D. degree in mathematics from Harvard in 1976. He has held faculty positions at the Weizmann Institute of Science, McGill University, ETH in Zurich, Technion, Ben-Gurion University of the Negev, and University of Minnesota. He is presently Julian Hightower Professor of Electrical and Biomedical Engineering at the Georgia Institute of Technology and Emory University. 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. Wagner RF, Smith SW, Sandrik JM, Lopez H. Statistics of speckle in ultrasound B-scans. *IEEE Trans. Sonics Ultrason.* May.1983 30:156–163.
2. Sehgal CM. Quantitative relationship between tissue composition and scattering of ultrasound. *J. Acoust. Soc. Amer.* Oct.1993 94:1944–1952. [PubMed: 8227740]
3. Goodman, JW. Statistical properties of laser speckle patterns. In: Dainty, JC., editor. *Laser Speckle and Related Phenomena*. Springer-Verlag; Berlin: 1977. p. 9-77.
4. Burckhardt CB. Speckle in ultrasound B-mode scans. *IEEE Trans. Sonics Ultrason.* Jan.1978 SU-25:1–6.
5. Abbott JG, Thurstone FL. Acoustic speckle: Theory and experimental analysis. *Ultrason. Imag.* 1979; 1:303–324.
6. Loupas T, McDicken WN, Allan PL. An adaptive weighted median filter for speckle suppression in medical ultrasound images. *IEEE Trans. Circuits Syst.* Jan.1989 36:129–135.
7. Karaman M, Kutay MA, Bozdagi G. An adaptive speckle suppression filter for medical ultrasound imaging. *IEEE Trans. Med. Imag.* Jun.1995 14:283–292.
8. Verhoeven JTM, Thijssen JM. Improvement of lesion detectability by speckle reduction filtering: A quantitative study. *Ultrason. Imag.* 1993; 15:181–204.
9. Jain, AK. *Fundamental of Digital Image Processing*. Prentice-Hall; Englewood Cliffs, NJ: 1989.
10. Daubechies, I. *Ten Lectures on Wavelets*. SIAM; Philadelphia, PA: 1992.
11. Mallat, SG. *A Wavelet Tour of Signal Processing*. Academic; New York: 1998.
12. Donoho DL. De-noising by Soft-Thresholding. *IEEE Trans. Inform. Theory.* May.1995 41:613–627.
13. Donoho DL, Johnstone I. Adapting to unknown smoothness via wavelet shrinkage. *J. Amer. Stat. Assoc.* 1995; 90:1200–1224.
14. Zong X, Laine AF, Geiser EA. Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing. *IEEE Trans. Med. Imag.* Aug.1998 17:532–540.
15. Hao X, Gao S, Gao X. A novel multiscale nonlinear thresholding method for ultrasound speckle suppressing. *IEEE Trans. Med. Imag.* Sep.1999 18:787–794.
16. Achim A, Bezerianos A, Tsakalides P. Novel Bayesian multiscale method for speckle removal in medical ultrasound images. *IEEE Trans. Med. Imag.* Aug.2001 20:772–783.
17. Guo H, Odegard JE, Lang M, Gopinath RA, Selesnick IW, Burrus CS. Wavelet based speckle reduction with application to SAR based ATD/R. *Proc. ICIP.* 1994:75–79.
18. Gagnon L, Jouan A. Speckle filtering of SAR images—A comparative study between complex-wavelet based and standard filters. *SPIE.* 1997; 3169:80–91.
19. Rudin L, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. *Phys. D.* 1992; 60:259–268.
20. Sapiro G, Tannenbaum A. On invariant curve evolution and image analysis. *Indiana Univ. J. Math.* 1993; 42:985–1009.
21. Angelsen, BAJ. *Ultrasound Imaging: Waves, Signals, and Signal Processing*. Emantec; Trondheim, Norway: 2000.
22. Jensen JA, Leeman S. Nonparametric estimation of ultrasound pulses. *IEEE Trans. Biomed. Eng.* Nov.1994 41:929–936. [PubMed: 7959799]
23. Taxt T. Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* Jul.1995 42:543–554.
24. Jensen J. A model for the propagation and scattering of ultrasound in tissue. *J. Acoust. Soc. Amer.* Jan.1991 89:182–191. [PubMed: 2002167]
25. Wear KA. The effects of frequency-dependent attenuation and dispersion on sound speed measurements: Application in human trabecular bone. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* Jan.2000 47:265–273.
26. Michailovich O, Adam D. A novel approach to the 2-D blind deconvolution problem in medical ultrasound. *IEEE Trans. Med. Imag.* Jan.2005 24:86–104.

27. Czerwinski RN, Jones DL, O'Brien WD. Line and boundary detection in speckle images. *IEEE Trans. Image Processing*. Dec.1998 7:1700–1714.
28. Bovik AC. On detecting edges in speckle imaginary. *IEEE Trans. Signal Processing*. Oct.1988 46:1618–1627.
29. Iracà D, Landini L, Verrazzani L. Power spectrum equalization for ultrasonic image restoration. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* Mar.1989 36:216–222.
30. Georgiou G, Cohen F. Statistical characterization of diffuse scattering in ultrasound images. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* Jan.1998 45:57–64.
31. Roychowdhury J, Demir A. Estimating noise in RF systems. *Proc. ICCAD*. Nov.1998 :199–202.
32. Michailovich O, Adam D. Robust estimation of ultrasound pulses using outlier-resistant de-noising. *IEEE Trans. Med. Imag.* Mar.2003 22:368–392.
33. Bruce AG, Donoho DL, Gao HY, Martin RD. Denoising and robust non-linear wavelet analysis. *SPIE*. 1994; 2242:325–336.
34. Tur M, Chin KC, Goodman JW. When is speckle noise multiplicative? *Applied Optics*. Apr.1982 21:1157–1159. [PubMed: 20389823]
35. Rakotomamonjy A, Deforge P, Marche P. Wavelet-based speckle noise reduction in ultrasound B-scan images. *Ultrason. Imag.* Apr.2000 22:73–94.
36. Tuthill TA, Sperry RH, Parker KJ. Deviations from Rayleigh statistics in ultrasonic speckle. *Ultrason. Imag.* Apr.1988 10:81–89.
37. Jakeman E, Tough RJA. Generalized K distribution: A statistical model for weak scattering. *J. Opt. Soc. Amer.* Sep.1987 4:1764–1772.
38. Dutt V, Greenleaf JF. Ultrasound echo envelope analysis using a homodyned K distribution signal model. *Ultrason. Imag.* 1994; 16:265–287.
39. Shankar PM. A general statistical model for ultrasonic backscattering from tissues. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* May.2000 47:727–736.
40. Fernandes D, Sekine M. Suppression of Weibull radar clutter. *IEICE Trans. Comm.* 1993; E76-B: 1231–1235.
41. Stacy EW. A generalization of the gamma distribution. *Annal. Math. Stat.* 1962; 33:1187–1192.
42. Raju BI, Srinivasan MA. Statistics of envelope of high-frequency ultrasonic backscatter from human skin in vivo. *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* Jul.2002 49:871–882.
43. Insana MF, Wagner RF, Garra BS, Brown DG, Shawker TH. Analysis of ultrasound image texture via generalized Rician statistics. *Opt. Eng.* 1986; 25(6):743–748.
44. Arsenault HH, April G. Properties of speckle integrated with a finite aperture and logarithmically transformed. *J. Opt. Soc. Amer.* 1976; 66:1160–1163.
45. Muller, P.; Vidakovic, B., editors. *Bayesian Inference in Wavelet Based Models, Lecture Notes in Statistics*. Vol. 141. Springer-Verlag; New York: 1999.
46. Bertsekas, D. *Nonlinear Programming*. Athena Scientific; Belmont, MA: 1999.
47. Chan T, Osher S, Shen J. The digital TV filter and nonlinear de-noising. *IEEE Trans. Image Processing*. Feb.2001 10:231–241.
48. Perona P, Malik J. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* Jul.1990 12:629–639.
49. You Y, Xu W, Tannenbaum A, Kaveh M. Behavioral analysis of anisotropic diffusion in image processing. *IEEE Trans. Image Processing*. Nov.1996 5:1539–1553.
50. Donoho, DL.; Coifman, RR. *Translation-Invariant Denoising*. Department of Statistics, Stanford University; Stanford, CA: May. 1995 Technical Report 475

Fig. 1. (Upper plot) Fragment of ultrasound image of a human liver. (Lower left plot) Axial profile of the image autocorrelation function. (Lower right plot) Lateral profile of the image autocorrelation function.

Fig. 2.

(Upper plot) The Fisher-Tippett *pdf* as given by (6) for the case $\sigma_f^2=1$. (Lower plot) A realization of the corresponding “noise”.

Fig. 3.

Examples of the pdf function as given by (13) for different parameters of the corresponding GG distribution.

Fig. 4.

Block diagram of the proposed modified homomorphic despecking algorithm (note that at the “denoising” stage of the block diagram an arbitrary filter can be substituted).

Fig. 5.

(Left panel) Standard envelope image of a human kidney. (Center panel) Standard envelope image after the log-transformation. (Right panel) The same log-transformed envelope image after applying the decorrelation and the outlier shrinkage according to the block diagram of Fig. 4.

Fig. 6. (Subplot A) Autocorrelation function of the standard log-envelope. (Subplot B) Autocorrelation function of the preprocessed log-envelope. (Subplot C) Histogram of the standard log-envelope. (Subplot D) Histogram of the preprocessed log envelope.

Fig. 7. (Left subplot) Amplitude profile modulating the WGN samples of simulated reflectivity functions *in silico* experiments. (Right Subplot) Example of a simulated (original) envelope image.

Fig. 8. (Subplots A1–A3) Images despeckled using the standard versions of the [from top to bottom] HWDS, TVDS, and ADDS algorithms. (Subplots B1–B3) images despeckled using the modified versions of the [from top to bottom] HWDS, TVDS, and ADDS algorithms.

Fig. 9.

(Upper subplot) Intensity values of an original envelope image taken along the line segment passing through the center of the largest “cyst” in the axial direction (solid line) versus the corresponding true homogeneity profile (dotted line). (Left column of subplots) Intensity values of the image taken along the same line segment after applying the standard versions of the HWDS (A1), TVDS (A2), and ADDS (A3) algorithms. (Right column of subplots) Intensity values of the image after applying the modified versions of the HWDS (B1), TVDS (B2), and ADDS (B3) algorithms.

Fig. 10.

(Upper subplot) Intensity values of a nonprocessed envelope image taken along the line segment passing through the center of the largest “cyst” in the lateral direction (solid line) versus the corresponding true homogeneity profile (dotted line). (Left column of subplots) Intensity values of the image taken along the same line segment after applying the standard versions of the HWDS (A1), TVDS (A2), and ADDS (A3) algorithms. (Right column of subplots) Intensity values of the image after applying the modified versions of the HWDS (B1), TVDS (B2), and ADDS (B3) algorithms.

Fig. 11.

(Upper image) Ordinal envelope image of a fragment of a carotid artery. (Left column of images) The image despeckled using the standard versions of the [from top to bottom] HWDS, TVDS, and ADDS algorithms. (Right column of images) The same image after despeckling using the modified versions of the [from top to bottom] HWDS, TVDS, and ADDS algorithms.

Fig. 12.

(Upper image) Ordinal envelope image of a fragment of a urinary bladder. (Left column of images) The image despeckled using the standard versions of the [from top to bottom] HWDS, TVDS, and ADDS algorithms. (Right column of images) The same image after despeckling using the modified versions of the [from top to bottom] HWDS, TVDS, and ADDS algorithms.

TABLE I
Performance Measures Computed for the Results of the Simulation Study in Section V-A Obtained Using Different Despeckling Approaches.

	HWDS			TVDS			ADDS		
	NMSE	Speckle-SNR	β	NMSE	Speckle-SNR	β	NMSE	Speckle-SNR	β
Original	0.687	1.908	0.185	0.687	1.908	0.185	0.687	1.908	0.185
Standard	0.532	2.755	0.277	0.392	4.923	0.324	0.471	3.246	0.298
Modified	0.302	7.906	0.375	0.229	10.412	0.487	0.307	5.903	0.381

TABLE II
Performance Measures Obtained in the In Vivo Study of Section V-B Using Different Despeckling Approaches.

	HWDS		TVDS		ADDS	
	Speckle-SNR	α	Speckle-SNR	α	Speckle-SNR	α
Original	0.716	0.0027	0.716	0.0027	0.716	0.0027
Standard	0.838	0.0107	0.972	0.253	0.915	0.0183
Modified	1.212	0.0207	1.433	0.287	1.356	0.0218