Novel Bayesian Multiscale Method for Speckle Removal in Medical Ultrasound Images

Alin Achim*, *Student Member, IEEE*, Anastasios Bezerianos, *Member, IEEE*, and Panagiotis Tsakalides, *Member, IEEE*

***Abstract*—A novel speckle suppression method for medical ultrasound images is presented. First, the logarithmic transform of the original image is analyzed into the multiscale wavelet domain. We show that the subband decompositions of ultrasound images have significantly non-Gaussian statistics that are best described by families of heavy-tailed distributions such as the alpha-stable. Then, we design a Bayesian estimator that exploits these statistics. We use the alpha-stable model to develop a blind noise-removal processor that performs a nonlinear operation on the data. Finally, we compare our technique with current state-of-the-art soft and hard thresholding methods applied on actual ultrasound medical images and we quantify the achieved performance improvement.**

***Index Terms*—Alpha-stable distributions, Bayesian estimation, ultrasound speckle, wavelet decomposition.**

I. Introduction

FOR MORE than two decades, ultrasonography has been considered as one of the most powerful techniques for imaging organs and soft tissue structures in the human body. Today, it is being used at an ever-increasing rate in the field of medical diagnostic technology. Ultrasonography is often preferred over other medical imaging modalities because it is noninvasive, portable, and versatile, it does not use ionizing radiations, and it is relatively low-cost. The images produced by commercial ultrasound systems are usually optimized for visual interpretation, since they are mostly used in real-time diagnostic situations. However, the main disadvantage of medical ultrasonography is the poor quality of images, which are affected by multiplicative speckle noise [1].

Imaging speckle is a phenomenon that occurs when a coherent source and a noncoherent detector are used to interrogate a medium, which is rough on the scale of the wavelength. Speckle occurs especially in images of the liver and kidney whose underlying structures are too small to be resolved by large wavelength ultrasound. The presence of speckle is undesirable since it degrades image quality and it affects the tasks of human interpretation and diagnosis. As a result, speckle filtering is a critical pre-processing step for feature extraction, analysis, and recognition from medical imagery measurements.

Current speckle reduction methods are based on temporal averaging [2], [3], median filtering [4], [5], and Wiener filtering. The adaptive weighted median filter, first introduced in [6], can effectively suppress speckle but it fails to preserve many useful details, being merely a low-pass filter. The classical Wiener filter, which utilizes the second-order statistics of the Fourier decomposition, is not adequate for removing speckle since it is designed mainly for additive noise suppression. To address the multiplicative nature of speckle noise, Jain developed a homomorphic approach, which by taking the logarithm of the image, converts the multiplicative into additive noise, and consequently applies the Wiener filter [1].

Recently, there has been considerably interest in using the wavelet transform as a powerful tool for recovering signals from noisy data [7]–[12]. The main reason for the choice of multiscale bases of decompositions is that the statistics of many natural signals, when decomposed in such bases, are significantly simplified. More specifically, methods based on multiscale decompositions consist of three main steps: First, the raw data are analyzed by means of the wavelet transform, then the empirical wavelet coefficients are shrunk, and finally, the denoised signal is synthesized from the processed wavelet coefficients through the inverse wavelet transform. These methods are generally referred to as *wavelet shrinkage techniques.* In [9], Zong *et al.* use a logarithmic transform to separate the noise from the original image. They adopt regularized soft thresholding (wavelet shrinkage) to remove noise energy within the finer scales and nonlinear processing of feature energy for contrast enhancement. A similar approach applied to synthetic aperture radar (SAR) images is presented in [10]. The authors perform a comparative study between a complex wavelet coefficient shrinkage filter and several standard speckle filters that are largely used by SAR imaging scientists, and show that the wavelet-based approach is among the best for speckle removal.

Thresholding methods have two main drawbacks: 1) the choice of the threshold, arguably the most important design parameter, is done in an *ad hoc* manner; and 2) the specific distributions of the signal and noise may not be well matched at different scales. To address these disadvantages, Simoncelli *et al.* developed nonlinear estimators, based on formal Bayesian theory, which outperform classical linear processors and simple thresholding estimators in removing noise from visual images [8], [13]. They used a generalized Laplacian model for the

Manuscript received October 25, 2000; revised June 8, 2001. The work of A. Achim was supported by a grant from the State Scholarship Foundation of Greece. The work of P. Tsakalides was supported by the Greek General Secretariat for Research and Technology under Program ΕΠΕΤ II, Code 97ΕΛ - 152. The Associate Editor responsiblle for coordinating the review of this paper and recommending its publication was M. Insana. *Asterisk indicates corresponding author.*

*A. Achim is with the Biosignal Processing Group, Medical Physics Department, University of Patras, 265 00 Rio, Greece (e-mail: amarian@heart.med.upatras.gr).

A. Bezerianos is with the Biosignal Processing Group, Medical Physics Department, University of Patras, 265 00 Rio, Greece.

P. Tsakalides is with the VLSI Design Laboratory, Department of Electrical and Computer Engineering, University of Patras, 261 10 Rio, Greece.

Publisher Item Identifier S 0278-0062(01)07083-5.

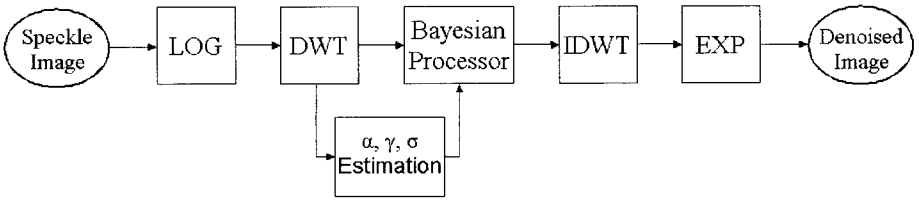

Fig. 1. Block diagram of the proposed multiscale homomorphic Bayesian-based algorithm for speckle suppression. Our proposed novel wavelet coefficient statistical characterization and Bayesian processing modules result in a more accurate ultrasound image reconstruction.

subband statistics of the signal and developed a noise-removal algorithm, which performs a "coring" operation to the data. The term "coring" refers to a widely used technique for noise suppression, which preserves high-amplitude observations while suppressing low-amplitude values from the high-pass bands of a signal decomposition.

It is recognized that parametric Bayesian processing presupposes proper modeling for the prior probability density function (PDF) of the signal. In a recent work, Tsakalides *et al.* showed that alpha-stable distributions, a family of heavy-tailed densities, are sufficiently flexible and rich to appropriately model wavelet coefficients of images in coding applications [14]. In this paper, we present a novel speckle suppression method for medical ultrasound images. The proposed processor consists of two major modules: 1) a subband representation function that utilizes the wavelet transform; and 2) a Bayesian denoising algorithm based on an alpha-stable *prior* for the signal. First, the original image is logarithmically transformed to change multiplicative speckle to additive white noise. Then, the transformed image is analyzed into a multiscale wavelet domain. We show that the subband decompositions of actual ultrasound images have significantly non-Gaussian statistics that are best described by families of heavy-tailed distributions like the alpha-stable. Motivated by our modeling results, we design a Bayesian estimator that exploits these statistics. We use the alpha-stable model to develop a blind speckle-suppression processor that performs a nonlinear operation on the data, and we relate this nonlinearity to the degree of non-Gaussianity of the data.

The paper is organized as follows: In Section II, we define the ultrasound speckle suppression problem by outlining the speckle noise model and common preprocessing steps such as the logarithmic and the wavelet transforms. Section III provides some necessary preliminaries on alpha-stable processes and presents results on the modeling of the subband coefficients of actual medical ultrasound images indicating their heavy-tailed nature. The design of a Bayesian estimator that exploits the signal alpha-stable statistics is described in Section IV. Section V compares the performance of our proposed algorithm with the performance of current denoising methods applied to actual ultrasound images and quantifies the achieved performance improvement. Finally, conclusions and future work directions are drawn in Section VI.

II. Problem Formulation

Speckle noise affects all coherent imaging systems including laser, SAR imagery, and ultrasound. Speckle may appear distinct in different imaging systems but it is always manifested in a granular pattern due to image formation under coherent waves. The basic properties of speckle are described by Goodman in [2] while the main differences between ultrasound and laser speckle are discussed in [3]. A general model for speckle noise proposed by Jain [1] was also used by Zong [9]. In the following, we formulate the ultrasound speckle removal problem starting with a brief essential overview of the speckle model.

Denote by $I(x, y)$ a noisy observation (i.e., the recorded ultrasound image) of the two-dimensional (2-D) function $S(x, y)$ (i.e., the noise-free image that has to be recovered) and by $\eta_m(x, y)$ and $\eta_a(x, y)$ the corrupting multiplicative and additive speckle noise components, respectively. One can write

$$I(x, y)=S(x, y) \cdot \eta_m(x, y)+\eta_a(x, y), \qquad (x, y)\in \mathbf{Z}^2. \quad (1)$$

Generally, the effect of the additive component of the speckle in ultrasound images is less significant than the effect of the multiplicative component. Thus, ignoring the term $\eta_a(x, y)$, one can rewrite (1) as

$$I(x, y) = S(x, y) \cdot \eta_m(x, y). \quad (2)$$

To transform the multiplicative noise model into an additive one, we apply the logarithmic function on both sides of (2)

$$\log I(x, y) = \log S(x, y) + \log \eta_m(x, y). \quad (3)$$

Expression (3) can be rewritten as

$$f(x, y) = g(x, y) + \epsilon(x, y) \quad (4)$$

where $f(\cdot)$, $g(\cdot)$, and $\epsilon(\cdot)$ are the logarithms of $I(\cdot)$, $S(\cdot)$, and $\eta_m(\cdot)$, respectively. In fact, this logarithmic transform constitutes the first preprocessing step of our proposed algorithm as shown in the block diagram depicted in Fig. 1.

At this stage, one can consider $\epsilon(x, y)$ to be white noise and subsequently apply any conventional additive noise suppression technique, such as Wiener filtering. However, it is recognized that standard noise filtering methods often result in blurred image features. Indeed, single-scale representations of signals, either in time or in frequency, are often inadequate when attempting to separate signals from noisy data. The wavelet transform has been proposed as a useful processing tool for signal recovery [15]–[17].

The wavelet transform expands a signal using a set of basis functions, which are obtained from a single prototype function called the "mother wavelet." The result of the expansion is a sequence of signal approximations at successively coarser resolutions. The so-called "detail signal" is the difference in information between approximations at two consecutive resolu-

tions, and it can be represented by another series expansion. If we consider an original 2-D signal of size $N \times N$, N usually being a power of 2 ($N = 2^J$), such a decomposition scheme is mathematically referred to as the *dyadic wavelet transform* (DWT). In image processing applications, the above scheme is applied along both the abscissa and the ordinate. Thus, the DWT decomposes images with a multiresolution scale factor of two, providing at each resolution level one low-resolution approximation and three spatially oriented wavelet details [17], [18], which are referred as *image subbands.*

The wavelet transform is a linear operation. Consequently, after applying the DWT to (4) we get, in each of the three directions, sets of noisy wavelet coefficients written as the sum of the transformations of the signal and of the noise

$$d^i_{j,k} = s^i_{j,k} + \xi^i_{j,k} \tag{5}$$

where $k = 0, \ldots, 2^{J+j} - 1$ and $-1 < j < -J$ refer to the decomposition level or scale and $i = 1, 2, 3$ refers to the three spatial orientations. In Fig. 2, we show an example of a three-scale decomposition of an ultrasound image.

Current state-of-the-art multiscale techniques for ultrasonic speckle suppression are based on various thresholding schemes [9], [19]. These methods try to address the inability of the original soft thresholding technique to balance between speckle suppression and signal detail preservation. In principle, a successful ultrasound imaging algorithm should achieve both noise reduction and feature preservation if it takes into consideration the true statistics of the signal and noise components. Previous studies related to wavelet shrinkage using Bayesian theory have underlined the need for a *prior* model that accurately approximates the probability density function of the signal and noise wavelet coefficients [8], [11], [20], [21]. For example, (1)–(4) suggest the use of a multiplicative random field as speckle noise model. It has been shown that, if the number of scatterers per resolution cell is large, a fully developed speckle pattern can be modeled as the magnitude of a complex Gaussian field with independent and identically distributed (i.i.d.) real and imaginary components (see [2], [22], and references therein). In order to generate spatially correlated speckle noise for use in simulations, one can lowpass filter a complex Gaussian random field and take the magnitude of the filtered output [23].

The Bayesian approach for ultrasound speckle noise removal, which we propose in this paper, is based on a novel heavy-tailed family of distributions that better models the prior statistics of the signal component. In Section III, we present the fundamental properties of the model and we justify its use by fitting actual ultrasonic signals.

III. Alpha-Stable Modeling of Ultrasound Wavelet Coefficients

This section is intended to provide an introduction on the alpha-stable statistical model used to characterize the wavelet subband coefficients of logarithmic transforms of actual ultrasound images. The model is suitable for describing signals that have highly non-Gaussian statistics and its parameters can be estimated from noisy observations. The appearance of alpha-stable models in the context of ultrasound images has been already noticed in [24] and [25], but they were used to process the ultrasound RF echoes, rather than the recorded images. A review of the state of the art on stable processes from a statistical point of view is provided by a collection of papers edited by Cambanis, Samorodnitsky and Taqqu [26], while textbooks in the area have been written by Samorodnitsky and Taqqu [27], and by Nikias and Shao [28].

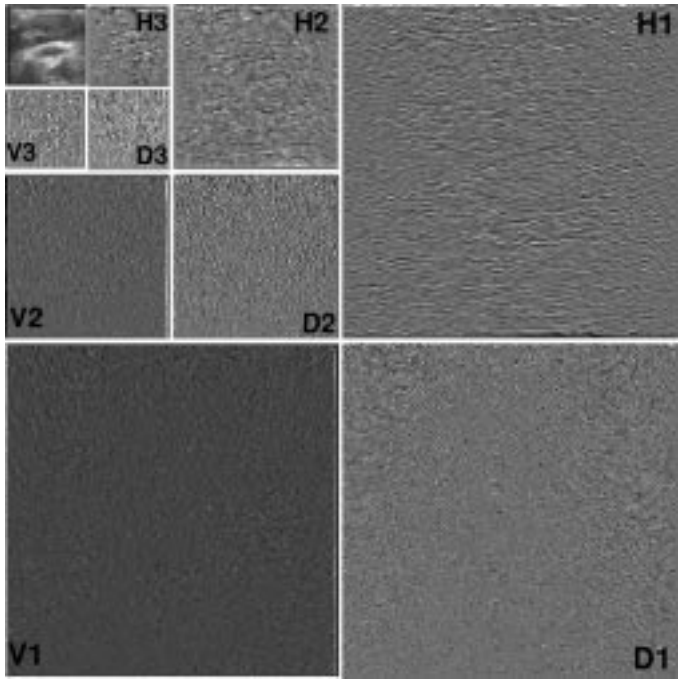

Fig. 2. Three-scale decomposition of an ultrasound image. The horizontal (H), vertical (V), and diagonal (D) details are shown in each scale. The upper-left image represents the approximation at level 3.

A. Basic Properties of the Alpha-Stable Family

The appeal of symmetric alpha-stable ($S\alpha S$) distributions as a statistical model for signals derives from some important theoretical and empirical reasons. First, stable random variables satisfy the stability property which states that linear combinations of jointly stable variables are indeed stable. The word *stable* is used because the shape of the distribution is unchanged (or stable) under such linear combinations. Second, stable processes arise as limiting processes of sums of i.i.d. random variables via the generalized central limit theorem. Actually, the *only* possible nontrivial limit of normalized sums of i.i.d. terms is stable. On the other hand, strong empirical evidence suggests that many data sets in several physical and economic systems exhibit heavy tail features that justify the use of stable models [29].

The $S\alpha S$ distribution is best defined by its characteristic function

$$\varphi(\omega) = \exp(j\delta\omega - \gamma|\omega|^\alpha) \tag{6}$$

where

α	*characteristic exponent,* taking values $0 < \alpha \leq 2$;
δ ($-\infty < \delta < \infty$)	*location parameter*;
γ ($\gamma > 0$)	*dispersion* of the distribution.

For values of α in the interval $(1, 2]$, the location parameter δ corresponds to the mean of the $S\alpha S$ distribution, while for $0 < \alpha \leq 1$, δ corresponds to its median. The dispersion parameter γ determines the spread of the distribution around its location parameter δ, similar to the variance of the Gaussian distribution.

The characteristic exponent α is the most important parameter of the $S\alpha S$ distribution and it determines the shape of the distribution. The smaller the characteristic exponent α is, the heavier the tails of the $S\alpha S$ density. This implies that random variables following $S\alpha S$ distributions with small characteristic exponents are highly impulsive. Gaussian processes are stable processes with $\alpha = 2$ while Cauchy processes result when $\alpha = 1$. In fact, no closed-form expressions for the general $S\alpha S$ PDF are known except for the Gaussian and the Cauchy members.

Although the $S\alpha S$ density behaves approximately like a Gaussian density near the origin, its tails decay at a lower rate than the Gaussian density tails [27]. Indeed, let X be a non-Gaussian $S\alpha S$ random variable. Then, as $x \to \infty$

$$P(X > x) \sim c_\alpha x^{-\alpha} \tag{7}$$

where $c_\alpha = \Gamma(\alpha)(\sin(\pi\alpha/2))/\pi$, $\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}\,dt$ is the Gamma function, and the statement $h(x) \sim g(x)$ as $x \to \infty$ means that $\lim_{x\to\infty} h(x)/g(x) = 1$. Hence, the tail probabilities are asymptotically power laws. In other words, while the Gaussian density has exponential tails, the stable densities have algebraic tails. Fig. 3 shows the tail behavior of several $S\alpha S$ densities including the Cauchy and the Gaussian. We should note that because expression (7) gives exactly the tail probability of the Pareto distribution, the term "*stable Paretian laws*" is used to distinguish between the fast decay of the Gaussian law and the Pareto like tail behavior when $\alpha < 2$.

The alpha-stable tail power law provided one of the earliest approaches in estimating the stability index α of real measurements [27]. The empirical distribution of the data, plotted on a log-log scale, should approach a straight line with slope $-\alpha$ if the data is stable. Another approach is based on quantiles [30]. Maximum likelihood (ML) methods developed by DuMouchel [31] and by Brorsen and Yang [32] are asymptotically efficient but were considered difficult to compute. Recently, Nolan showed that ML estimation of stable parameters is feasible by designing an efficient program [33].

One consequence of heavy tails is that only moments of order less than α exist for the non-Gaussian alpha-stable family members, i.e.,

$$E|X|^p < \infty \qquad \text{for } p < \alpha. \tag{8}$$

As a result, stable Paretian laws have infinite variance. In the past, the infinite variance property of the $S\alpha S$ family has caused skeptics to dismiss the stable model. With the same reasoning, one could argue that the routinely used Gaussian distribution, which has infinite support, should also be dismissed as a model of bounded measurements. In practice, one should remember that it is important to capture the shape of the distribution and that the variance is only one measure of the spread of a density [33].

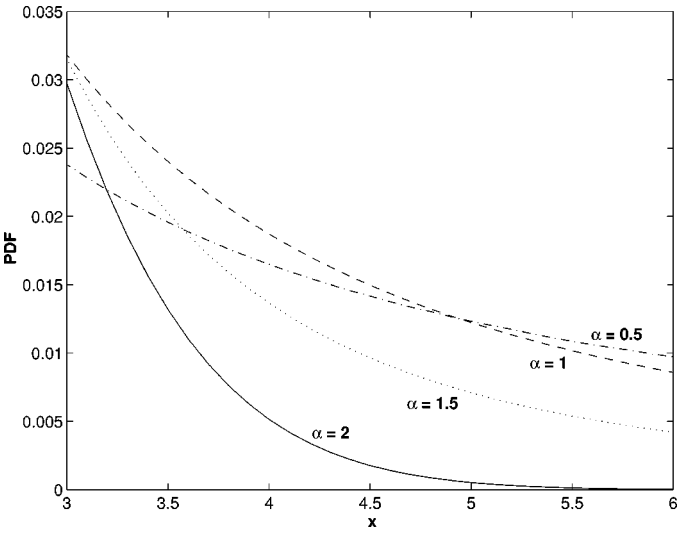

Fig. 3. Tail behavior of the $S\alpha S$ probability density functions for $\alpha = 0.5$, 1.0 (Cauchy), 1.5, and 2.0 (Gaussian). The dispersion parameter is kept constant at $\gamma = 1$.

B. Modeling Results of Ultrasound Image Subbands

In the past, several authors have pointed out that, in a subband representation of images, histograms of wavelet coefficients have heavier tails and more sharply peaked modes at zero than what is assumed by the Gaussian distribution [13], [14], [18]. In this section, we study whether the stable family provides a flexible and appropriate tool for modeling the coefficients within the framework of multiscale wavelet analysis of logarithmically transformed ultrasound images.

Two sets of test images, obtained from two different sources, are included in this research. The first set consists of a series of 44 abdominal ultrasound images (DICOM format) including liver, kidney, gall bladder, and pancreas images. These images were acquired from the same patient with a 4-MHz transducer frequency on a GE LOGIQ 500 system. They have been made available to us by the IT Lab at the Medical University of South Carolina. The second set of test data comes from a directory containing example DICOM image files that were donated by various vendors for the DICOM demonstrations held at the annual meetings of the Radiological Society of North America from 1993 to 1996 (ftp://wuerlim.wustl.edu/pub/dicom/images/). We considered two criteria for selecting images for our test set. First, we looked for good quality images in order to be able to consider them as noise-free. Moreover, since speckle appears to some extent in any ultrasound image, we have first processed the actual images using the homomorphic Wiener filter [1] and considered the resulting images as reasonable approximations of the speckle free data. Also, we were interested in performing experiments on images of different organs and from various sources in order to be able to obtain modeling results, which we can claim to be general enough. Because of limited space, in this paper we describe the modeling of eight representative images. Each image is referred with the name of the organ that was imaged. All of them have a 256 gray-level resolution.

We proceed in two steps: First, we assess whether the data deviate from the normal distribution and if they have heavy tails. To determine that, we make use of the normal probability plots.

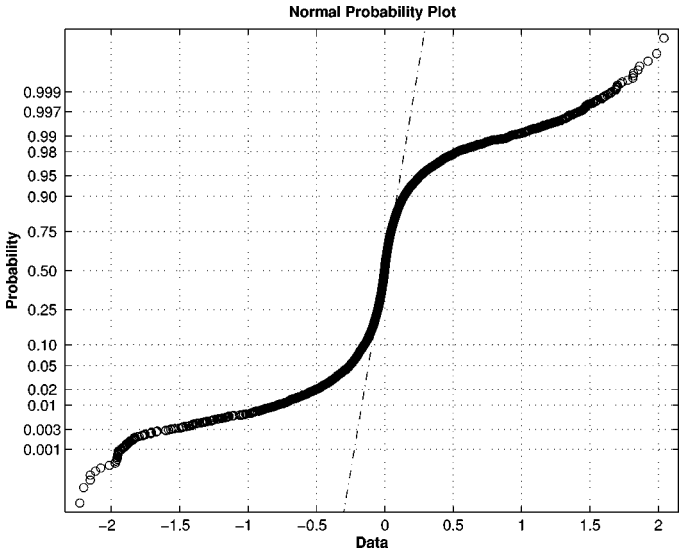

Fig. 4. Normal probability plot of the vertical subband at the first level of decomposition of the gallbladder image (GLBD_Vsbd_1lvl data for short). Characterization of data non-Gaussianity. The "∘" marks correspond to the empirical probability density versus the data value for each point in the sample. Since the circles are in a curve that does not follow the straight Gaussian line, the normality assumption is violated for this data.

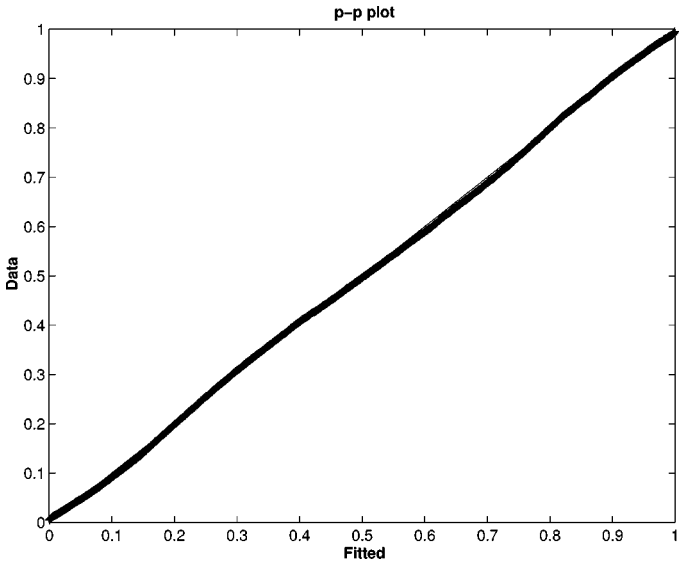

Fig. 5. Stabilized p-p plot for $S\alpha S$ fit of data set GLBD_Vsbd_1lvl. The "+" marks, denoting the empirical probability density, are in a curve that very accurately follows the straight $S\alpha S$ line corresponding to $\alpha = 1.069$.

Then, we check if the data is in the stable domain of attraction by estimating the characteristic exponent, α, directly from the data and by providing the related confidence intervals. Several methods have been proposed for estimating stable parameters. Here, we use the ML method described by Nolan in [33], which gives reliable estimates and provides the most tight confidence intervals. As further stability diagnostics, we employ probability density plots that give a good indication of whether the $S\alpha S$ fit matches the data near the mode and on the tails of the distribution.

Fig. 4 depicts the normal probability plot of the vertical subband at the first level of decomposition of the gallbladder image (data denoted by "GLBD_Vsbd_1lvl," for short). The plot provides strong evidence that the underlying distribution is not normal. The circles in the plot show the empirical probability versus the data value for each point in the sample. The circles are in a curve that does not follow the straight Gaussian line and thus, the normality assumption is violated for this data. While non-Gaussian stable densities are heavy-tailed, not all heavy-tailed distributions are stable. Hence, in Figs. 5 and 6 we assess the stability of the data. First, the characteristic exponent is estimated and the data sample is fitted with the corresponding stable distribution. For the particular case shown here, the characteristic exponent of the $S\alpha S$ distribution which best fits the data was estimated to be $\hat{\alpha} = 1.069$. The stabilized p-p $S\alpha S$ plot in Fig. 5 shows a highly accurate stable fit for this data set.

Naturally, the real question is whether the stable fit describes the data more accurately than other PDF functions proposed in the literature. Here, we compare the $S\alpha S$ fits with those provided by the generalized Laplacian density function proposed by Mallat in [18] and also used by Simoncelli in [13]

$$f_{s,p}(c) = \frac{e^{-|c/s|^p}}{Z(s,p)} \tag{9}$$

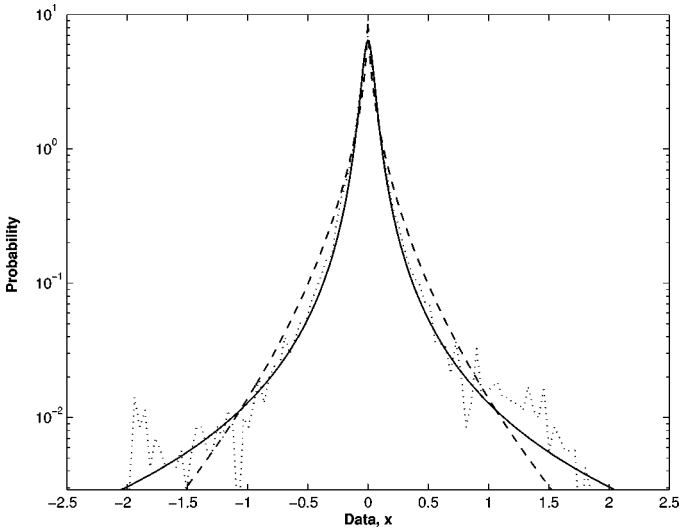

Fig. 6. Modeling of the ultrasound image wavelet coefficients GLBD_Vsbd_1lvl with the $S\alpha S$ and the generalized Laplacian density functions, depicted in solid and dashed lines, respectively. The $S\alpha S$ distribution has characteristic exponent $\alpha = 1.069$ and dispersion $\gamma = 0.051$ while the generalized Laplacian has parameters $p = 0.498$ and $s = 0.022$ [cf. (9)]. The dotted line denotes the empirical PDF. Note that the $S\alpha S$ PDF provides a better fit to both the mode and the tails of the empirical density of the actual data.

where $Z(s, p) = 2(s/p)\Gamma(1/p)$. The parameters s and p can be computed from the second and fourth moments of the data

$$\sigma^2 = \frac{s^2\Gamma\left(\frac{3}{p}\right)}{\Gamma\left(\frac{1}{p}\right)}, \qquad k = \frac{\Gamma\left(\frac{1}{p}\right)\Gamma\left(\frac{5}{p}\right)}{\Gamma^2\left(\frac{3}{p}\right)} \tag{10}$$

where σ^2 is the distribution variance, and k is the kurtosis. Fig. 6 shows that the $S\alpha S$ distribution is superior to the generalized Laplacian distribution because it provides a better fit to both the mode and the tails of the empirical density of the actual data.

TABLE I
ALPHA-STABLE MODELING OF WAVELET SUBBAND COEFFICIENTS OF ACTUAL ULTRASOUND IMAGES. ML PARAMETER ESTIMATES AND 95% CONFIDENCE INTERVALS FOR THE $S\alpha S$ CHARACTERISTIC EXPONENT, α. THE TABULATED KEY PARAMETER α DEFINES THE DEGREE OF NON-GAUSSIANITY AS DEVIATIONS FROM THE VALUE $\alpha = 2$, WHICH CORRESPONDS TO THE GAUSSIAN CONDITION. THE SIZE OF EACH IMAGE IS GIVEN IN PARENTHESES

IMAGE	Level	Image Subbands		
		Horizontal	Vertical	Diagonal
breast (361 × 361)	1	1.279 ± 0.016	0.965 ± 0.013	1.128 ± 0.015
	2	1.380 ± 0.032	1.178 ± 0.028	1.248 ± 0.031
	3	1.349 ± 0.059	1.073 ± 0.053	1.303 ± 0.058
gallbladder (256 × 256)	1	1.349 ± 0.022	1.069 ± 0.020	0.974 ± 0.019
	2	1.516 ± 0.044	1.265 ± 0.043	1.352 ± 0.040
	3	1.508 ± 0.081	1.295 ± 0.077	1.181 ± 0.075
kidney (293 × 293)	1	1.382 ± 0.019	1.121 ± 0.018	1.124 ± 0.017
	2	1.494 ± 0.038	1.308 ± 0.038	1.427 ± 0.039
	3	1.348 ± 0.071	1.126 ± 0.067	1.495 ± 0.074
liver (256 × 256)	1	1.469 ± 0.020	1.197 ± 0.021	1.269 ± 0.022
	2	1.482 ± 0.045	1.254 ± 0.042	1.546 ± 0.045
	3	1.112 ± 0.072	1.008 ± 0.068	1.391 ± 0.080
pancreas (230 × 230)	1	1.443 ± 0.023	1.159 ± 0.023	1.264 ± 0.023
	2	1.491 ± 0.050	1.230 ± 0.045	1.523 ± 0.050
	3	1.237 ± 0.085	0.936 ± 0.071	1.501 ± 0.091
spinal cord (256 × 256)	1	1.253 ± 0.022	1.110 ± 0.020	1.196 ± 0.021
	2	1.416 ± 0.044	1.298 ± 0.043	1.528 ± 0.045
	3	1.424 ± 0.081	1.116 ± 0.072	1.487 ± 0.082
urinary bladder (190 × 190)	1	1.167 ± 0.027	1.008 ± 0.026	1.002 ± 0.027
	2	1.279 ± 0.056	1.167 ± 0.054	1.141 ± 0.053
	3	1.408 ± 0.101	1.006 ± 0.088	0.958 ± 0.085
spline (256 × 256)	1	1.417 ± 0.020	1.104 ± 0.020	1.178 ± 0.021
	2	1.470 ± 0.041	1.046 ± 0.038	1.306 ± 0.043
	3	1.154 ± 0.073	0.922 ± 0.064	1.197 ± 0.075

For every image we iterated three times the separable wavelet decomposition described in Section II and we modeled the coefficients of each subband by using the $S\alpha S$ family. The wavelet decomposition was done using Daubechies' Symmlet 4 basis wavelet because we found this basis to be the most effective in decorrelating the data. The results are summarized in Table I, which shows the ML estimates of the characteristic exponent α together with the corresponding 95% confidence intervals. It can be observed that the confidence interval depends on the size of the images and on the particular level of decomposition. The confidence interval becomes larger as the size decreases and as the level increases since the number of samples used for estimating α gets smaller. The table demonstrates that the coefficients of different subbands and decomposition levels exhibit various degrees of non-Gaussianity. The important observation is that all subbands exhibit distinctly non-Gaussian characteristics, with values of α varying between 0.9 and 1.6, away from the Gaussian point of $\alpha = 2$. Our modeling results clearly point to the need for the design of Bayesian processors that take into consideration the non-Gaussian heavy-tailed character of the data to achieve close to optimal speckle mitigation performance.

IV. A BAYESIAN PROCESSOR FOR ULTRASOUND SPECKLE REMOVAL

Current state-of-the-art wavelet-based denoising and image enhancement techniques employ a combination of wavelet shrinkage by soft and hard thresholding together with a generalized adaptive gain (GAG) for feature emphasis (see [9] and references therein). In particular, Zong *et al.* apply soft thresholding at fine scales (levels 1 and/or 2) and hard thresholding within middle levels 3 and/or 4 to eliminate noise, followed by nonlinear processing of feature energy to enhance contrast. The regularized threshold parameter used in [9] is related to the noise level, orientation, and scale through a judiciously chosen but at the same time *ad hoc* linearly decreasing function. Moreover, the five parameters which determine the empirical GAG function in [9] are tuned experimentally to achieve the appropriate nonlinear stretching of wavelet coefficients that accomplishes the desired contrast enhancement.

In this section, our goal is the design of a formal Bayesian estimator that recovers the signal component of the wavelet coefficients in ultrasound images by using an alpha-stable signal prior distribution. The proposed processor is motivated by the modeling studies in Section III, it is based on solid statistical theory, and it does not depend on the use of *ad hoc* thresholding and stretching parameters.

In a Bayesian framework, referring to (5), $d_{j,k}$, $s_{j,k}$, and $\xi_{j,k}$ are considered as samples of the random variables d, s, and ξ, respectively. The signal component s is modeled according to a $S\alpha S$ distribution with zero location parameter, while the noise component ξ is modeled as a zero-mean Gaussian random variable. Our goal is to find the Bayes risk estimator $\hat{s}$ that minimizes the conditional risk, which is the loss averaged over the

conditional distribution of s, given the set of wavelet coefficients, d

$$\hat{s}(d) = \arg\min_{\hat{s}} \int L[s,\, \hat{s}(d)] P_{s|d}(s|d)\, ds. \tag{11}$$

The Bayes risk estimator under a quadratic cost function minimizes the mean-square error (mse) and is given by the conditional mean of s, given d

$$\hat{s}(d) = \int s P_{s|d}(s|d) \cdot ds. \tag{12}$$

Of course, the mse metric is defined for random variables that possess finite second-order moments. In this work, the signal component of the wavelet coefficients is modeled as an alpha-stable random variable that does not have finite second-order statistics. Hence, we use the absolute error $|s - \hat{s}(d)|$ as the loss function in expression (11). Under this loss function, expression (11) is well defined for all $S\alpha S$ random variables with characteristic exponent α grater than one. The Bayesian estimator that we consider minimizes the mean absolute error and can be shown to be the conditional median of s, given d [34]. But, since the conditional density $P_{s|d}(s|d)$ is symmetric around zero, the conditional median coincides with the conditional mean. Hence, the Bayesian estimator for the absolute error cost function is again given by (12).

Bayes' theorem gives the *a posteriori* probability density function of s based on the measured set of wavelet coefficients

$$P_{s|d}(s|d) = \frac{P_{d|s}(d|s) P_s(s)}{\int P_{d|s}(d|s) P_s(s) \cdot ds} \tag{13}$$

where $P_s(s)$ is the prior PDF of the signal component of the wavelet coefficients of the ultrasound image and $P_{d|s}(d|s)$ is the *likelihood* function. Substituting (13) into (12), we get

$$\hat{s}(d) = \frac{\int P_{d|s}(d|s) P_s(s) s \cdot ds}{\int P_{d|s}(d|s) P_s(s) \cdot ds} = \frac{\int P_\xi(d-s) P_s(s) s \cdot ds}{\int P_\xi(d-s) P_s(s) \cdot ds} = \frac{\int P_\xi(\xi) P_s(s) s \cdot ds}{\int P_\xi(\xi) P_s(s) \cdot ds} \tag{14}$$

where $P_\xi(\xi)$ is the PDF of the wavelet coefficients corresponding to the noise.

In order to be able to construct the Bayesian processor in (14), first we estimate the parameters of the prior distributions of the signal (s) and noise (ξ) components of the wavelet coefficients (d). Then, we use the parameters to "build" the two prior PDFs $P_\xi(\xi)$ and $P_s(s)$ and the nonlinear (in general) I/O relationship $\hat{s}(d)$. Observing (14), we note that the denominator, referred to as the "evidence" is the PDF of the noisy observation, d, computed as the convolution between the noise and signal PDFs. Hence, we need a parameterized model for the two PDFs that provides a good fit to the statistics of the ultrasound images. Furthermore, the distribution parameters should be estimated from the noisy observations in an efficient manner.

Motivated by our modeling results in Section III-B, we use a two-parameter $S\alpha S$ model for the signal component while we use a zero-mean Gaussian model for the noise component. In other words, the observed signal is a mixture of $S\alpha S$ signal and Gaussian noise. Moreover, we consider the signal and noise components to be independent. Because of the lack of closed-form expressions for the general $S\alpha S$ PDF, we propose a method that is based on characteristic functions. In particular, since the PDF of the measured coefficients is the convolution between the PDFs of the signal and noise components, the associated characteristic function of the measurements is given by the product of the characteristic functions of the signal and noise

$$\Phi_d(\omega) = \Phi_s(\omega) \cdot \Phi_\xi(\omega) \tag{15}$$

where

$$\Phi_s(\omega) = \exp(-\gamma_s |\omega|^{\alpha_s}), \qquad 1 < \alpha \leq 2$$

and

$$\Phi_\xi(\omega) = \exp\left(-\frac{\sigma^2}{2} |\omega|^2\right).$$

Using expression (15), we estimate the parameters α_s, γ_s, and σ by fitting the Fourier transform of the empirical PDF of the measured coefficients with function $\Phi_d(\omega)$ in the least-squares (LS) sense

$$\{\hat{\alpha}_s,\, \hat{\gamma}_s,\, \hat{\sigma}\} = \arg \min_{\hat{\alpha_s}, \hat{\gamma}_s, \hat{\sigma}} \sum_i^n [\Phi_d(\omega_i) - \Phi_{d_\epsilon}(\omega_i)]^2 \tag{16}$$

where $\Phi_{d_\epsilon}(\omega)$ denotes the empirical characteristic function. In practice, we first estimate the level of noise σ, and we optimize (16) only with respect to the $S\alpha S$ parameters α_s and γ_s. As proposed in [35], a robust estimate of the noise standard deviation, σ, is obtained in the finest decomposition scale by the measured wavelet coefficients as

$$\hat{\sigma} = \frac{1}{0.6745} \operatorname{MAD}\left(\{d_{J,k},\, 0 \leq k < 2^J\}\right) \tag{17}$$

where the operator MAD signifies the *median absolute deviation* and J denotes the highest level of wavelet decomposition. We found that this method for estimating the $S\alpha S$ parameters gives reliable estimates, it is not computationally expensive and, more importantly, it allows us to estimate the parameters from the noisy measurements. We should note here that Paulson *et al.* used a similar approach to estimate the parameters of alpha-stable densities and observed that LS fitting in the characteristic function domain produces estimates within two standard errors of the actual values and with a bias that is inversely proportional to the sample size [36].

As expected, in the general case the Bayesian processor described in (14) does not have a closed-form expression. Only for the case of Gaussian signal in Gaussian noise a well-known closed-form solution exists

$$\hat{s}(d) = \frac{\sigma_s^2}{\sigma_s^2 + \sigma^2}\, d \tag{18}$$

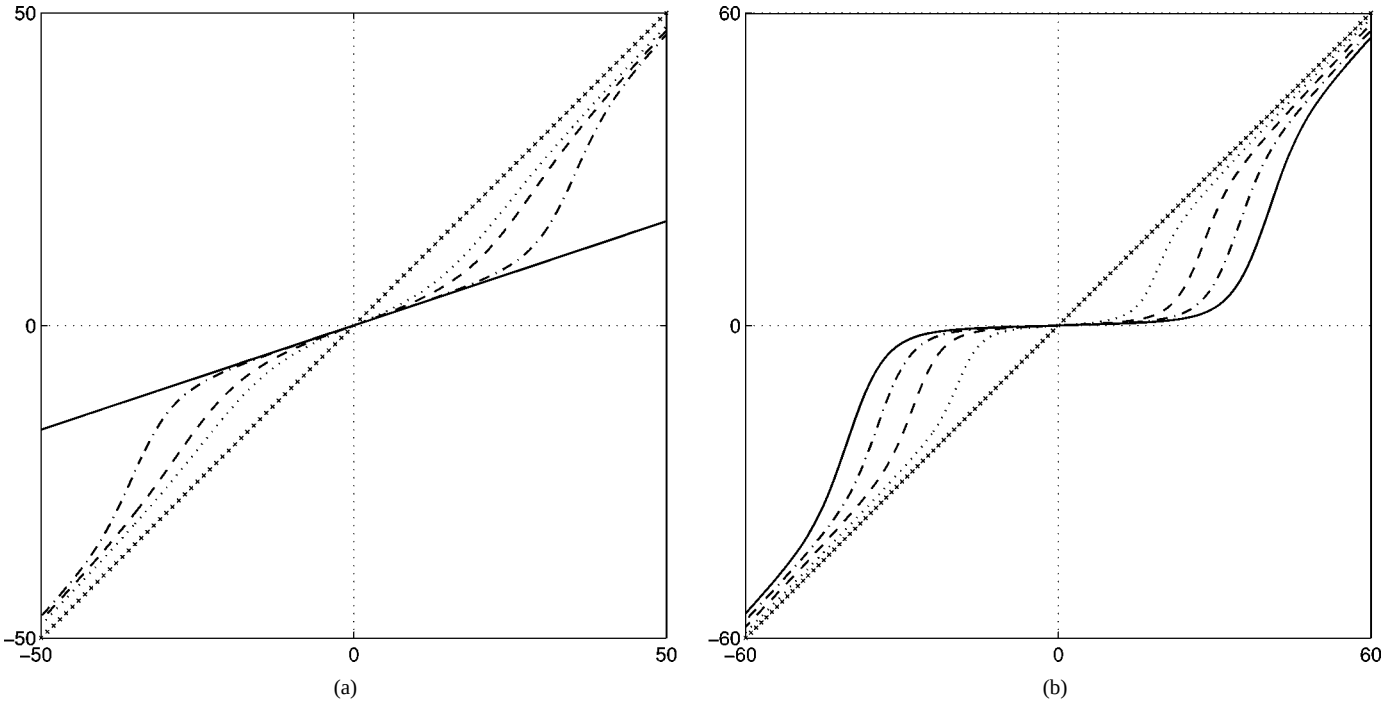

Fig. 7. Bayesian processor input–output (I/O) curves for alpha-stable signal ($1 < \alpha \leq 2$) and Gaussian noise prior distributions. The straight line with $\times$s indicates the identity function. (a) Bayesian curves for constant γ_s/σ ratio and four different signal statistics corresponding to $\alpha = 2$ (Gaussian signal, solid line), $\alpha = 1.95$ (slightly non-Gaussian signal, dash-dotted line), $\alpha = 1.5$, and $\alpha = 1.05$ (considerably heavy-tailed signal, dashed and dotted lines, respectively). (b) Bayesian curves for $S\alpha S$ signal $\alpha = 1.5$ and four different γ_s/σ ratios: 0.1 (solid), 0.12 (dash-dotted), 0.14 (dashed), and 0.2 (dotted). Note the processor nonlinear "coring" operation, which preserves large-amplitude observations and suppresses small-amplitude values in a statistically optimal fashion.

where σ_s^2 is the Gaussian signal variance. In other words, the processing is a simple linear rescaling of the measurement. For the general non-Gaussian $S\alpha S$ signal case, we computed numerically the Bayesian processor function in (14). Fig. 7(a) depicts the Bayesian I/O curves for four different values of the signal characteristic exponent, α, namely, $\alpha = 2$ (Gaussian data), $\alpha = 1.95$ (slightly non-Gaussian data), $\alpha = 1.5$, and $\alpha = 1.05$ (considerably heavy-tailed data). All curves except the case $\alpha = 2$, correspond to a nonlinear "coring" operation, i.e., large-amplitude observations are essentially preserved while small-amplitude values are suppressed. This is expected since small measurement values are assumed to come from signal values close to zero. Fig. 7(a) also illustrates the processor dependency on the parameter α of the signal prior PDF. Specifically, for a given ratio γ_s/σ, the amount of shrinkage decreases as α decreases. The intuitive explanation for this behavior is that the smaller the value of α, the heavier the tails of the signal PDF and the greater the probability that the measured value is due to the signal.

On the other hand, Fig. 7(b) shows how the processor nonlinearity is varied for certain signal statistics ($\alpha = 1.5$) and various noise levels. It is evident from the curves that as the noise level increases, the amount of shrinkage also increases. We should note at this point that curves similar to the ones in Fig. 7 are chosen adaptively by our processor since the PDF parameters are estimated by means of (16) and (17) at each level of wavelet decomposition and for each orientation. In actual ultrasound images, at the first decomposition levels where the wavelet coefficients arising from noise are predominant, the Bayesian shrinkage function would resemble to that corresponding to low signal-to-noise ratio (SNR). As the resolution decreases, in general the noise level decreases and the nonlinearity applied to the wavelet coefficients corresponds to the high SNR curves in Fig. 7(b) gradually approaching the identity function as the SNR becomes very high.

V. Experimental Results

We tested our proposed multiscale Bayesian speckle suppressing algorithm on the ultrasound images modeled in Section III-B. In order to obtain speckle images, we degraded the original test images by multiplying them with unit-mean random fields, as shown in expression (2). We generated spatially correlated speckle noise $\eta_m(x, y)$ by lowpass filtering a complex Gaussian random field and taking the magnitude of the filtered output. We controlled the correlation length of the speckle by appropriately setting the size of the kernel used to introduce correlation to the underlying Gaussian noise. A kernel size of one corresponds to white noise. On the other hand, in order to allow the noise correlation to taper gradually to zero we could not set the kernel size arbitrarily large. Thus, a short-term correlation obtained with a kernel of size three was sufficient to model reality. In our experiments, we considered three different levels of simulated speckle noise.

We compared the results of our approach with other speckle reduction techniques including median filtering, homomorphic Wiener filtering, and wavelet shrinkage denoising using soft and hard thresholding. All the parameters involved in these methods were selected by trial-and-error in order to get an optimal result from each method. Specifically, for the median filter we used a 3×3 mask for the lowest level of noise and 5×5 masks for

TABLE II
IMAGE ENHANCEMENT MEASURES OBTAINED BY THE FIVE DENOISING METHODS TESTED ON THE KIDNEY-ULTRASOUND IMAGE. THE S/mse IS GIVEN IN dB. VALUES OF THE CORRELATION MEASURE, β, CLOSE TO UNITY DENOTE OPTIMAL EDGE PRESERVATION PERFORMANCE

Method	MSE	S/MSE	β	MSE	S/MSE	β	MSE	S/MSE	β
Without Filtering	26.0528	5.61	0.2872	16.2945	9.69	0.4336	7.2111	16.77	0.7357
Median Filtering	13.7002	11.19	0.2138	9.7630	14.13	0.3449	6.6770	17.43	0.5708
Homomorphic Wiener	13.8381	11.10	0.1776	8.8103	15.03	0.4988	6.3944	17.81	0.6287
Soft Thresholding	13.6370	11.23	0.3364	8.9242	14.91	0.5880	6.1133	18.20	0.8062
Hard Thresholding	13.5001	11.32	0.3160	8.6400	15.20	0.5576	5.5608	19.02	0.7569
Bayesian Denoising	12.7398	11.82	0.4559	8.2037	15.65	0.6253	4.8869	20.15	0.8249

the other two levels. The homomorphic Wiener filter was implemented using a window of size 5×5 pixels for the highest level of noise and 3×3 pixels in all the other cases. For soft thresholding we used a threshold $t = 1.5\sigma_d$, while for hard thresholding we have chosen $t = 3\sigma_d$, σ_d being the standard deviation of the wavelet coefficients. Both wavelet shrinkage soft and hard thresholding schemes were developed using Daubechies' Symmlet 8 mother wavelet. Denoising results using this basis wavelet have been found to be less affected by pseudo-Gibbs phenomena [9]. Moreover, in order to minimize such side effects, we have embedded all wavelet-based methods (including the Bayesian approach) into the cycle spinning algorithm [37]. This consists in averaging the result of the wavelet shrinkage method over all circulant shifts of the input image. In practice we found that a number of 8 translations is sufficient. The maximum number of wavelet decompositions we used was 5.

In order to quantify the achieved performance improvement, three different measures were computed based on the original and the denoised data. For quantitative evaluation, an extensively used measure is the mse defined as

$$\text{mse} = \frac{1}{K} \sum_{i=1}^{K} \left(\hat{S}_i - S_i \right)^2 \tag{19}$$

where

S original image;
$\hat{S}$ denoised image;
K image size.

The standard signal to noise ratio (SNR) is not adequate to evaluate the noise suppression in case of multiplicative noise. Instead, a common way to achieve this in coherent imaging is to calculate the signal-to-mse (S/mse) ratio, defined as [10]

$$\text{S/mse} = 10 \log_{10} \left(\sum_{i=1}^{K} S_i^2 \Big/ \sum_{i=1}^{K} \left(\hat{S}_i - S_i \right)^2 \right). \tag{20}$$

This measure corresponds to the classical SNR in the case of additive noise.

Remember that in ultrasound imaging, we are interested in suppressing speckle noise while at the same time preserving the edges of the original image that often constitute features of interest for diagnosis. Thus, in addition to the above quantitative performance measures, we also consider a qualitative measure for edge preservation. More specifically, we used a parameter β originally defined in [19] and [23]

$$\beta = \frac{\Gamma\left(\Delta S - \overline{\Delta S}, \widehat{\Delta S} - \overline{\widehat{\Delta S}}\right)}{\sqrt{\Gamma\left(\Delta S - \overline{\Delta S}, \Delta S - \overline{\Delta S}\right) \cdot \Gamma\left(\widehat{\Delta S} - \overline{\widehat{\Delta S}}, \widehat{\Delta S} - \overline{\widehat{\Delta S}}\right)}} \tag{21}$$

where ΔS and $\widehat{\Delta S}$ are the high-pass filtered versions of S and $\hat{S}$ respectively, obtained with a 3×3 pixel standard approximation of the Laplacian operator, and

$$\Gamma(S_1, S_2) = \sum_{i=1}^{K} S_{1_i} \cdot S_{2_i}. \tag{22}$$

The correlation measure, β should be close to unity for an optimal effect of edge preservation.

The obtained values of mse, S/mse, and β for all methods applied to the kidney image are given in Table II. It is evident from the table that the three wavelet-based methods (i.e., soft and hard thresholding and our proposed Bayesian denoising technique) are more successful in speckle noise suppression than median and homomorphic Wiener filtering. In terms of mse and S/mse, the soft thresholding scheme achieves comparable performance with the homomorphic Wiener filter, but the visual quality of the soft threshold processed image seems to be better (cf. Fig. 8). This is due to the fact that the soft thresholding approach is not intended to minimize the mse, the result being an estimator which achieves a low variance at the expense of bias [35]. Observing the β metric values, we see that the multiresolution techniques exhibit a clearly better performance in terms of edge preservation, as expected. Among them, our proposed Bayesian approach exhibits the best performance according to all three metrics.

Fig. 8 shows a representative result from the processing of the noisy kidney image. The simulated speckle image shown in Fig. 8(b) corresponds to a S/mse value of 9.69 dB. For this noise level, all the methods that we tested achieved a good speckle suppression performance. However, the median and homomorphic Wiener filters loose many of the signal details and the resulting images are blurred [cf. Fig. 8(c) and (d)]. On the other hand, the images processed by soft and hard thresholding are oversmoothed [cf. Fig. 8(e) and (f)]. Clearly, as it can be seen in Fig. 8(g), our proposed Bayesian processor effectively reduces

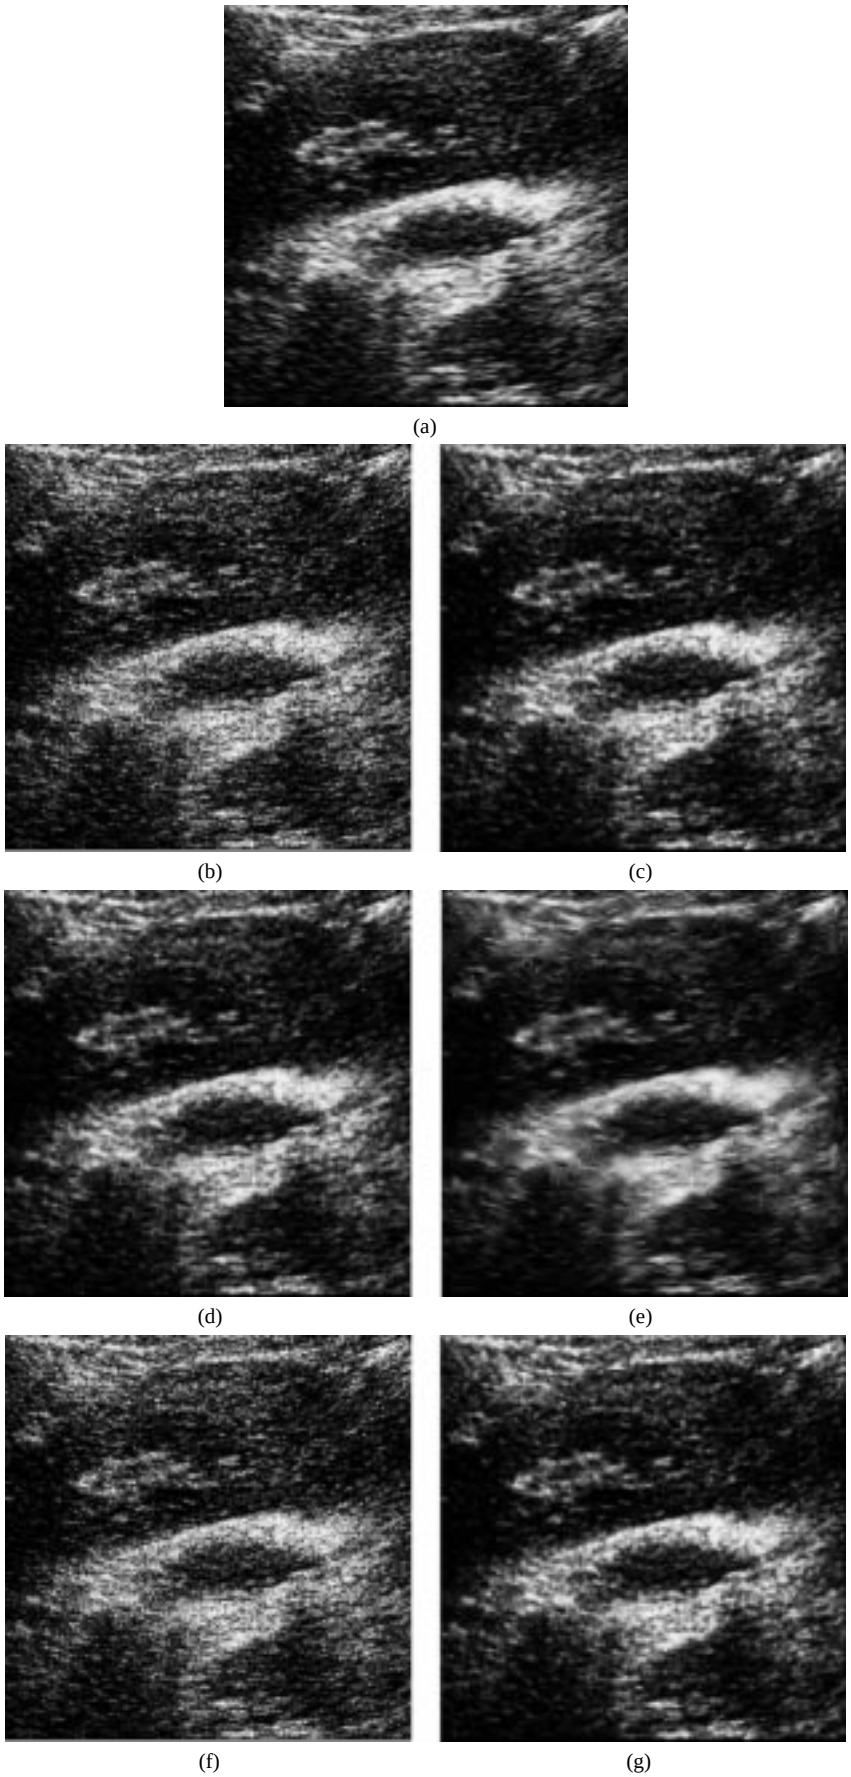

(a) (b) (c) (d) (e) (f) (g)

Fig. 8. Results of various speckle suppressing methods. (a) Original kidney ultrasound image. (b) Image degraded with simulated speckle noise S/mse 9.69 dB. (c) Median filtering. Results of various speckle suppressing methods. (d) Homomorphic Wiener filtering. (e) Soft thresholding. (f) Hard thresholding. (g) Proposed Bayesian denoising.

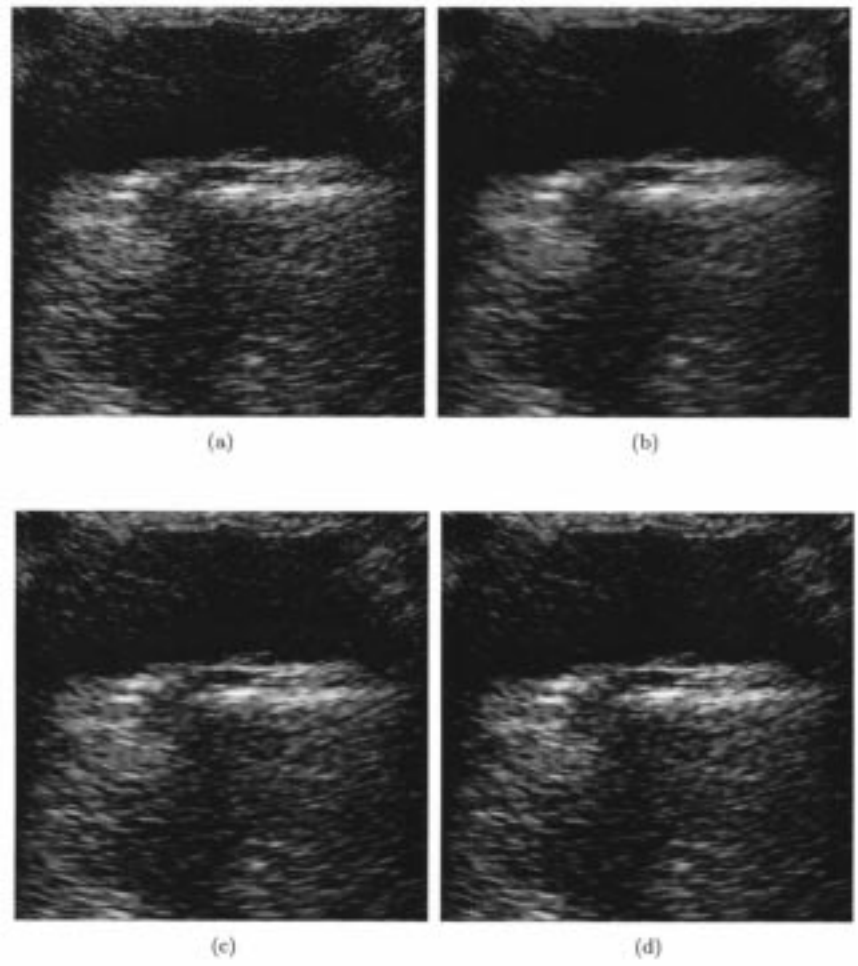

Fig. 9. (a) Noisy ultrasound image of the bladder. (b) Image denoised using translation-invariant soft thresholding. (c) Image denoised using translation-invariant hard thresholding. (d) Image enhanced using our Bayesian algorithm.

speckle, it preserves step edges, and it enhances fine signal details, better than the other methods.

Indeed, the problem with the mse, S/mse, and β measures, or with any other metric, is their connection to the visual interpretation of an human observer. A radiologist, in analyzing ultrasound images, does not compute any of the above measures. Hence, in order to visually study the merit of the proposed $S\alpha S$ subband coefficient modeling and the Bayesian processor, we chose a noisy ultrasound image, applied the algorithm without artificially adding noise, and visually evaluated the denoised image. The results of this experiment are shown in Fig. 9. The figure only shows results obtained using the wavelet based schemes, which were proved to give better results for simulated speckle noise. Although qualitative evaluation in this case is highly subjective, the results of this experiment seem to be consistent with the simulation results. It appears that the Bayesian processor performs like a feature detector, retaining the features that are clearly distinguishable in the speckled data.

VI. Conclusion

In this paper, we introduced a novel multiscale nonlinear homomorphic method for speckle suppression in ultrasound images. Three are the main processing stages of our approach. First, similarly to existing multiresolution techniques [9], [19], the logarithm of the image is decomposed into several scales through a multiresolution analysis employing the 2-D wavelet transform. This step guarantees that the speckle is transformed from multiplicative into additive and its characteristics are differentiated from the signal characteristics in each decomposition level. The second and third steps differentiate our technique from existing ones. After decomposing the original image, the signal and noise components in various scales are modeled as $S\alpha S$ and Gaussian processes, respectively. The parameters of the distributions are estimated from the measurements by means of a LS fitting in the characteristic function domain. We showed that the class of $S\alpha S$ distributions is more effective in modeling detail image histograms than other exponentially tailed densities.

In the third step, a Bayesian processor based on a $S\alpha S$ signal prior is built at each scale for statistically optimal signal feature extraction and speckle suppression. The main advantage of our method is that the obtained I/O shrinkage functions are optimal in the Bayesian sense. As a result, a more accurate signal reconstruction is achieved in each scale. Our processor was tested and found to be more effective than thresholding methods, which are *ad hoc* in the sense that they do not allow for an exact matching of the signal and noise distributions at different scales and orientations. The method could be easily adapted for the purpose

of denoising other types of biomedical images where the noise can be (eventually after an appropriate transformation) modeled as additive Gaussian and signal-independent. Naturally, our approach is more computationally expensive due to the fact that the prior distribution parameters need to be estimated at each scale of interest. However, this is not a serious problem for off-line processing.

Currently, we are addressing several issues related to the work we presented in this paper. One major issue is the choice of a statistical model for the speckle noise component of the wavelet coefficients that is more appropriate than the currently used Gaussian model. It is to be tested whether the $S\alpha S$ family is a good model also for the noise component. In this case, our problem will be formulated as Bayesian signal detection from measurements that are mixtures of $S\alpha S$ signal in $S\alpha S$ noise with different characteristic exponents, in general. Research in this direction is under way and it will be presented in a future correspondence.

ACKNOWLEDGMENT

The authors would like to thank Dr. C. F. Starmer and the IT Lab at the Medical University of South Carolina for providing the ultrasound images used in this paper.

REFERENCES

[1] A. K. Jain, *Fundamental of Digital Image Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[2] J. W. Goodman, "Some fundamental properties of speckle," *J. Opt. Soc. Amer.*, vol. 66, pp. 1145–1150, Nov. 1976.

[3] J. G. Abbott and F. L. Thurstone, "Acoustic speckle: Theory and experimental analysis," *Ultrason. Imag.*, vol. 1, pp. 303–324, 1979.

[4] A. Ioannidis, D. Kazakos, and D. D. Watson, "Application of median filtering on nuclear medicine scintigram images," in *Proc. 7th Int. Conf. Pattern Recognition*, 1984, pp. 33–36.

[5] E. R. Ritenour, T. R. Nelson, and U. Raff, "Application of the median filter to digital radiographic images," in *Proc. IEEE Int. Conf. Acoust. Speech, Signal Processing*, 1984, pp. 23.1.1–23.1.4.

[6] T. Loupas, W. N. Mcdicken, and P. L. Allan, "An adaptive weighted median filter for speckle suppression in medical ultrasonic images," *IEEE Trans. Circuits Syst.*, vol. 36, pp. 129–135, Jan. 1989.

[7] D. L. Donoho, "Denoising by soft-thresholding," *IEEE Trans. Inform. Theory*, vol. 41, pp. 613–627, May 1995.

[8] E. P. Simoncelli and E. H. Adelson, "Noise removal via Bayesian wavelet coring," in *Third Int Conf. Image Processing*, vol. 1, September 1996, pp. 379–382.

[9] X. Zong, A. F. Laine, and E. A. Geiser, "Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing," *IEEE Trans. Med. Imag.*, vol. 17, pp. 532–540, Aug. 1998.

[10] L. Gagnon and A. Jouan, "Speckle filtering of SAR images—A comparative study between complex-wavelet based and standard filters," in *SPIE Proc.*, vol. 3169, 1997, pp. 80–91.

[11] M. Popescu, P. Cristea, and A. Bezerianos, "Multiresolution distributed filtering: A novel technique that reduces the amount of data required in high-resolution electrocardiography," *Future Gen. Comput. Syst.*, no. 15, pp. 195–209, 1999.

[12] S. Papadimitriou and A. Bezerianos, "Multiresolution analysis and denoising of computer performance evaluation data with the wavelet transform," *J. Syst. Architect.*, vol. 42, pp. 55–65, 1996.

[13] E. P. Simoncelli, "Bayesian denoising of visual images in the wavelet domain," in *Bayesian Inference in Wavelet Based Models*, P. Muller and B. Vidakovic, Eds. New York: Springer-Verlag, June 1999, ch. 18, pp. 291–308.

[14] P. Tsakalides, P. Reveliotis, and C. L. Nikias, "Scalar quantization of heavy-tailed signals," *IEE Proc.—Vision, Imag. Signal processing*, vol. 147, pp. 475–484, Oct. 2000.

[15] Y. Meyer, "Principe d'incertitude, bases hilbertiennes et algèbre d'opérateurs," in *Proc. Bourbaki Seminar*, 1985–1986, Paper 662.

[16] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Commun. Pure Appl. Math.*, vol. 41, pp. 909–996, 1988.

[17] S. Mallat, *A Wavelet Tour of Signal Processing*. New York: Academic Press, 1998.

[18] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, pp. 674–692, July 1989.

[19] X. Hao, S. Gao, and X. Gao, "A novel multiscale nonlinear thresholding method for ultrasonic speckle suppressing," *IEEE Trans. Med. Imag.*, vol. 18, pp. 787–794, Sept. 1999.

[20] H. A. Chipman, E. D. Kolaczyk, and R. E. McCulloch, "Adaptive Bayesian wavelet shrinkage," *J. Amer. Statist. Assoc.*, vol. 92, pp. 1413–1421, 1997.

[21] B. Vidakovic, "Nonlinear wavelet shrinkage with Bayes rules and Bayes factors," *J. Amer. Statist. Assoc.*, vol. 93, pp. 173–179, 1998.

[22] R. N. Czerwinski, D. L. Jones, and W. D. O'Brien, Jr., "Line and boundary detection in speckle images," *IEEE Trans. Image Processing*, vol. 7, pp. 1700–1714, Dec. 1998.

[23] F. Sattar, L. Floreby, G. Salomonsson, and B. Lovstrom, "Image enhancement based on a nonlinear multiscale method," *IEEE Trans. Image Processing*, vol. 6, pp. 888–895, June 1997.

[24] M. A. Kutay and A. P. Petropulu, "Power-law shot noise model for the ultrasound RF echo," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Processing*, June 2000, pp. 3787–3790.

[25] M. A. Kutay, A. P. Petropulu, and C. W. Piccoli, "On modeling biomedical ultrasound RF echoes using a power-law shot noise model," *IEEE Trans. Ultrason., Ferroelect. Freq. Contr.*, July 2001.

[26] S. Cambanis, G. Samorodnitsky, and M. S. Taqqu, Eds., *Stable Processes and Related Topics*. Boston, MA: Birkhauser, 1991.

[27] G. Samorodnitsky and M. S. Taqqu, *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. New York: Chapman and Hall, 1994.

[28] C. L. Nikias and M. Shao, *Signal Processing with Alpha-Stable Distributions and Applications*. New York: Wiley, 1995.

[29] R. Adler, R. Feldman, and M. S. Taqqu, *A Guide to Heavy Tails: Statistical Techniques and Applications*. Boston, MA: Birkhauser, 1998.

[30] E. F. Fama and R. Roll, "Some properties of symmetric stable distributions," *J. Amer. Statist. Assoc.*, vol. 63, pp. 817–836, 1968.

[31] W. H. DuMouchel, "Stable distributions in statistical inference," Ph.D. dissertation, Dept. Statistics, Yale Univ., New Haven, CT, 1971.

[32] B. W. Brorsen and S. R. Yang, "Maximum likelihood estimates of symmetric stable distribution parameters," *Commun. Statist.-Simulat.*, vol. 19, pp. 1459–1464, 1990.

[33] J. P. Nolan, "Maximum likelihood estimation and diagnostics for stable distributions," Tech. Rep., Dept. Math. Statist., Amer. Univ., Washington, DC, June 1999.

[34] L. L. Scharf, *Statistical Signal Processing: Detection, Estimation and Time Series Analysis*. Menlo Park, CA: Addison Wesley, 1991.

[35] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, pp. 425–455, 1994.

[36] A. S. Paulson, E. W. Holcomb, and R. A. Leitch, "The estimation of the parameters of the stable laws," *Biometrika*, vol. 62, pp. 163–170, 1975.

[37] R. R. Coifman and D. L. Donoho, "Translation-invariant de-noising," in *Wavelets and Statistics, Lecture Notes in Statistics 103*, A. Antoniadis and G. Oppenheim, Eds. Berlin, Germany: Springer-Verlag, 1995.