Ultrasound Image Denoising by Spatially Varying Frequency Compounding*

Yael Erez[1], Yoav Y. Schechner[1], and Dan Adam[2]

[1] Dept. Electrical Engineering, Technion – Israel Inst. Tech., Haifa 32000, Israel
{yaele@tx, yoav@ee}.technion.ac.il
[2] Dept. Biomedical Engineering, Technion – Israel Inst. Tech., Haifa 32000, Israel
dan@bm.technion.ac.il

Abstract. Ultrasound images are very noisy. Along with system noise, a significant noise source is the speckle phenomenon, caused by interference in the viewed object. Most past approaches for denoising ultrasound images essentially blur the image, and they do not handle attenuation. Our approach, on the contrary, *does not* blur the image and *does* handle attenuation. Our denoising approach is based on frequency compounding, in which images of the same object are acquired in different acoustic frequencies, and then compounded. Existing frequency compounding methods have been based on simple averaging, and have achieved only limited enhancement. The reason is that the statistical and physical characteristics of the signal and noise vary with depth, and the noise is correlated. Hence, we suggest a spatially varying frequency compounding, based on understanding of these characteristics. Our method suppresses the various noise sources and recovers attenuated objects, while maintaining high resolution.

1 Introduction

Ultrasound is an imaging technique that uses high frequency acoustic waves. It is safe, suitable for many applications and is relatively cheap. It is used in sonar, medical imaging and material science work. However, there are some problems that interfere with the diagnosis. Fig. 1 illustrates some of these problems. The most prominent problem, which distinguishes ultrasound from most imaging techniques, is strong speckle noise. Speckles appear as grains of different sizes and intensities, that result from the coherent nature of the ultrasound radiation [2]. The speckle image is signal dependent. It is time invariant and thus cannot be suppressed by temporal averaging. A second problem is attenuation. The acoustic signal propagating in the medium is scattered and absorbed [2],
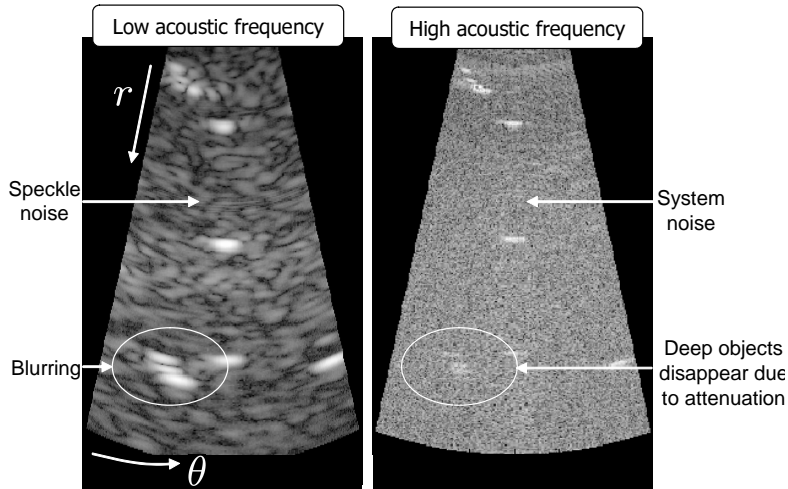

Fig. 1. Problems that disrupt diagnosis in ultrasound imaging. The depth range is $0 - 12cm$.

and hence attenuated. This phenomenon is more pronounced in high acoustic frequencies. When the attenuated signal is amplified, it is accompanied by amplification of system noise, which is signal independent. Below a certain level of signal to noise ratio (SNR), objects are overwhelmed by system noise, thus amplification in post-processing does not reconstruct these objects.

Most past approaches for denoising ultrasound images have used standard image reconstruction tools, such as weighted median filter [9], wavelet based methods [4] [5], Gaussian non-linear filters [3] and anisotropic diffusion [14]. All these methods essentially blur the image. Moreover, they do not handle spatially varying physical effects, as attenuation. Another approach is *frequency compounding*,[3] in which images of an object are acquired in different acoustic frequencies, and then compounded [10]. Existing compounding methods [1] have used simple processing methods such as pointwise arithmetic averaging, and have achieved only limited enhancement.

In this paper we present a method that does not suffer from the mentioned disadvantages. It is based on frequency compounding, and the images are analyzed in a stochastic manner. The stochastic denoising is spatially varying and it is based on statistical and physical characteristics of the signal and noise as a function of depth and acoustic frequency. The stochastic denoising shows significant speckle reduction, with no resolution loss, while deep objects are reconstructed as well.

---

[3] *Spatial compounding* is also possible. Yet, it introduces a complex registration problem, and it does not improve detection in deep regions.

2   Theoretical Background

Let us first model blur. We assume the ultrasound images to be two-dimensional (2D), given in their polar coordinates (r, θ). The r coordinate (radial axis) is the axis of wave propagation, and θ (lateral axis) represents a serial scan of the direction of the radiating ultrasound beam. The 2D signal measured by the system is the result of a natural filtering of the 2D tissue reflectivity function $a_0(r, \theta)$ with a 2D point spread function (PSF). This PSF is space variant. In particular, its lateral support changes with the depth r: the acoustic beam is focused at a certain depth, where the lateral PSF is narrowest, while at other depths this PSF gradually widens. Yet, in small regions we can assume this filter to be space invariant. There, the measured signal is

$$a^{\mathrm{RF}}(r, \theta) = a_0(r, \theta) * h(r, \theta) \ . \tag{1}$$

Following [12], it is reasonable to assume the PSF to be separable. The PSF also depends on system properties, such as acoustic frequency [2].

Image formation is also affected by attenuation of ultrasound in the medium [2]. A general simple and effective model of the amplitude of the signal is

$$a^{\mathrm{RF}}(r, \theta) = e^{-2\alpha r f_{\mathrm{acoustic}}} a_0(r, \theta) * h(r, \theta) \ , \tag{2}$$

where α is the *attenuation coefficient* of the acoustic amplitude, and f_{acoustic} is the acoustic frequency. A rule of thumb [2] is: attenuation in tissue is approximately $1 dB/(cm \cdot MHz)$, for a signal going from a probe to the object and then returning. It is clear from Eq. (2) that attenuation depends on the acoustic frequency: high acoustic frequencies suffer from stronger attenuation and thus a lower SNR, particularly at large depths. This is evident in Fig. 1

In ultrasound systems, the measured signal a^{RF} undergoes several standard conversion steps. First, attenuation is compensated for. Then, the acoustic modulation is extracted: note that a^{RF} is a high-frquency (MHz) signal, which is modulated by the tissue reflectivity function. To extract the tissue information, the envelope of the attenuation-compensated a^{RF} is detected, yielding

$$a^{\mathrm{magnitude}}(r, \theta) = \left| \mathtt{envelope} \left[e^{2\alpha r f_{\mathrm{acoustic}}} \cdot a^{\mathrm{RF}}(r, \theta) \right] \right| \ , \tag{3}$$

where $\mathtt{envelope}\,[g(r)]$ is an operator [2] that extracts the envelope of a modulated wave $g(r)$ (recall that r is the axis of wave propagation). Note that Eq. (3) derives the *modulus* of the envelope, since the envelope is complex, in general.

Speckle Noise

Speckle noise has a granular texture, as presented in Fig. 2. Speckles degrade the ability to resolve details and detect objects of size comparable to the speckle size. This noise stems from point scatterers in an homogenous tissue, that cannot be resolved by the ultrasound system. These point scatterers, which are much smaller than the ultrasound wavelength, scatter the wave. Two or more waves
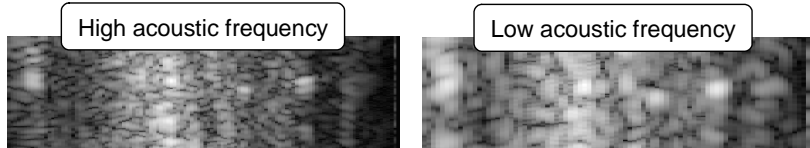
Fig. 2. Speckle appearance of the same tissue in different acoustic frequencies. High acoustic frequency speckles are smaller than the low acoustic frequency speckles.

travelling to the probe from such scatterers may interfere with each other, constructively or destructively, creating bright and dark spots, termed speckles. For interference, the backscattered signal from the scatterers should overlap in time and space. This happens when the distance between them is within the PSF (radially and laterally) support. This is an *important* point to remember: the speckle typical size is similar to the PSF support. Since the PSF changes with depth, the statistics of this noise are space (depth)-variant. Furthermore, they change when the acoustic frequency used to acquire the image changes, as shown in Fig. 2, as the PSF does. We exploit these properties in this paper.

Speckle is generally modelled as multiplicative noise [7]. The overall detected magnitude is

$$a^{\text{total}}(r, \theta) = a^{\text{magnitude}}(r, \theta) \cdot s^{\text{magnitude}}(r, \theta) + \eta(r, \theta) , \qquad (4)$$

where the real number $s^{\text{magnitude}}$ represents real nonnegative speckle noise at certain coordinates, and η represents system noise there. The system noise increases with depth, due to the attenuation compensation done in Eq. (3). Still, assume for a moment that the additive noise is sufficiently small compared to the multiplicative noise. Then, a log operation on Eq. (4) transforms speckles to additive noise

$$\underbrace{\log \left[a^{\text{total}}(r, \theta) \right]}_{a^{\text{log}}} \approx \underbrace{\log \left[a^{\text{magnitude}}(r, \theta) \right]}_{\log(a^{\text{magnitude}})} + \underbrace{\log \left[s^{\text{magnitude}}(r, \theta) \right]}_{s^{\text{log}}} . \qquad (5)$$

The logarithm operation is standard when displaying ultrasound images on a computer screen [2], since the dynamic range of a^{total} is very large [2]. Therefore, in the image used for display, the speckle noise is already additive.

3   Solution

Our solution is spatially varying frequency compounding, based on the best linear unbiased estimator (BLUE), also known as Gauss-Markov or weighted least squares (WLS) [8]. This stochastic method relies on the following principles:

– The compounding should be space (depth) variant, since the statistics of noise change with the depth r, as the PSF.

- In speckles, adjacent pixels are correlated [12]. Therefore, it is desirable that compounding would account for this spatial correlation.
- Speckles in different images of the same object, acquired with different acoustic frequencies, are correlated [13]. Therefore, simple averaging is not very efficient for speckle reduction. On the contrary, we should account for the cross-correlation between different acoustic channels.
- The method is not intended for sharpening. Therefore, it does not include deblurring. Nevertheless, we do not want to further blur existing information.
- In general, deep objects are not visible in high acoustic frequency (due to increased attenuation). However, thanks to our use of a low acoustic frequency image in the compounding, we should end up seeing even the deepest objects.
- In general, spatial resolution is low, when using a low acoustic frequency (due to a wider PSF). However, thanks to our use of a high acoustic frequency image in the compounding, we should end up with high spatial resolution, at least in close distance.

In the following we detail our solution.

3.1   Speckle Model

We refer to the signals $\mathbf{a}^{\mathrm{magnitude}}$ and $\mathbf{a}^{\mathrm{log}}$ as discrete $N \times 1$ vectors. When acquiring K images in different acoustic frequencies, then based on Eq. (5),

$$\begin{pmatrix} \mathbf{a}_1^{\mathrm{log}} \\ \mathbf{a}_2^{\mathrm{log}} \\ \vdots \\ \mathbf{a}_K^{\mathrm{log}} \end{pmatrix} = \begin{pmatrix} \log \mathbf{a}_1^{\mathrm{magnitude}} \\ \log \mathbf{a}_2^{\mathrm{magnitude}} \\ \vdots \\ \log \mathbf{a}_K^{\mathrm{magnitude}} \end{pmatrix} + \begin{pmatrix} \mathbf{s}_1^{\mathrm{log}} \\ \mathbf{s}_2^{\mathrm{log}} \\ \vdots \\ \mathbf{s}_K^{\mathrm{log}} \end{pmatrix} . \tag{6}$$

At this point we use the principle mentioned above, of not attempting to invert blur, thus we do not consider the blur h in the reconstruction. Using a δ function for h in Eq. (2) can estimate $\hat{a}_0(r,\theta) = e^{2\alpha r f_{\mathrm{acoustic}}} a^{\mathrm{RF}}(r,\theta)$. Therefore, we set

$$\mathbf{a}_k^{\mathrm{magnitude}} \approx |\mathtt{envelope}\left(\hat{\mathbf{a}}_0\right)| , \tag{7}$$

for all k. Now, the frames $\mathbf{a}_k^{\mathrm{magnitude}}$ differ in the noise, which is indeed different, especially the speckle noise. All frames include a similar object content, i.e.,

$$\mathbf{a}_1^{\mathrm{magnitude}} \approx \mathbf{a}_2^{\mathrm{magnitude}} \approx ... \approx \mathbf{a}_K^{\mathrm{magnitude}} = \mathbf{a}^{\mathrm{magnitude}} , \tag{8}$$

Hence, Eq. (6) reduces to

$$\begin{pmatrix} \mathbf{a}_1^{\mathrm{log}} \\ \mathbf{a}_2^{\mathrm{log}} \\ \vdots \\ \mathbf{a}_K^{\mathrm{log}} \end{pmatrix} = \begin{pmatrix} \mathbf{I} \\ \mathbf{I} \\ \vdots \\ \mathbf{I} \end{pmatrix} \log\left(\mathbf{a}^{\mathrm{magnitude}}\right) + \begin{pmatrix} \mathbf{s}_1^{\mathrm{log}} \\ \mathbf{s}_2^{\mathrm{log}} \\ \vdots \\ \mathbf{s}_K^{\mathrm{log}} \end{pmatrix} . \tag{9}$$

3.2   BLUE

Consider data $\mathbf{a}^{\mathrm{data}}$ in the general linear model

$$\mathbf{a}^{\mathrm{data}} = \mathbf{Ha} + \mathbf{n} \ , \tag{10}$$

where $\mathbf{H}$ is a known $KN \times N$ matrix (operator), $\mathbf{a}$ is an $N \times 1$ vector of variables to be estimated, and $\mathbf{n}$ is an $N \times 1$ noise vector with zero mean and covariance $\mathbf{C}$. The Gauss-Markov theorem [8] states that the BLUE of $\mathbf{a}$ is

$$\hat{\mathbf{a}} = (\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{C}^{-1} \mathbf{a}^{\mathrm{data}} \ . \tag{11}$$

To apply the BLUE on Eq. (9), we substitute $\mathbf{a} = \log\left(\mathbf{a}^{\mathrm{magnitude}}\right)$ as in Eqs. (10,11), while $\mathbf{a}^{\mathrm{data}}$ represents the vector on the left-hand-side of Eq. (9). Now, the noise covariance matrix $\mathbf{C}$ used in Eq. (11) has the form

$$\mathbf{C} = \begin{pmatrix} \mathbf{C}_{\mathbf{s}_1^{\log}\mathbf{s}_1^{\log}} & \mathbf{C}_{\mathbf{s}_1^{\log}\mathbf{s}_2^{\log}} & \cdots & \mathbf{C}_{\mathbf{s}_1^{\log}\mathbf{s}_K^{\log}} \\ \mathbf{C}_{\mathbf{s}_2^{\log}\mathbf{s}_1^{\log}} & \mathbf{C}_{\mathbf{s}_2^{\log}\mathbf{s}_2^{\log}} & \cdots & \mathbf{C}_{\mathbf{s}_2^{\log}\mathbf{s}_K^{\log}} \\ \vdots & & \ddots & \\ \mathbf{C}_{\mathbf{s}_K^{\log}\mathbf{s}_1^{\log}} & \mathbf{C}_{\mathbf{s}_K^{\log}\mathbf{s}_2^{\log}} & \cdots & \mathbf{C}_{\mathbf{s}_K^{\log}\mathbf{s}_K^{\log}} \end{pmatrix} \ , \tag{12}$$

where $\mathbf{C}_{\mathbf{s}_k^{\log}\mathbf{s}_i^{\log}}$ is the cross-covariance matrix between two speckle images $\mathbf{s}_k^{\log}$ and $\mathbf{s}_i^{\log}$ in different acoustic frequencies. Eq. (11) performs a linear combination of all data $\mathbf{a}^{\mathrm{data}}$ (all pixels in all images) in order to estimate the value in each pixel of $\hat{\mathbf{a}}$. Therefore, the BLUE may potentially perform deconvolution, in addition to noise averaging. Nevertheless, in our case

$$\mathbf{H} = \left(\mathbf{I}, \mathbf{I}, \ldots, \mathbf{I}\right)^T \ , \tag{13}$$

since we do not attempt deblurring. The BLUE exploits the correlation between variables. This enables denoising based on partially correlated variables, in contrary to a simple average, which implicitly assumes uncorrelated variables.

3.3   Spatially Varying BLUE

To use the BLUE we need to know the noise mean and covariance (*statistics*), in the set of images we use. When applying the method for noise reduction, we need to consider the noise statistics. We estimate the covariance functions from the data itself. We performed empirical measurements of these functions. This empirical study revealed that the noise is not stationary. This is not surprising, since according to [12], the auto and cross correlations of speckles depend on the system PSF, which (Sec. 2) changes significantly with depth.

Let us first examine a certain block in the image. We can assume stationarity within this block. However, the statistics change in different image regions. Is there a need to divide the whole image to blocks, and measure the statistics
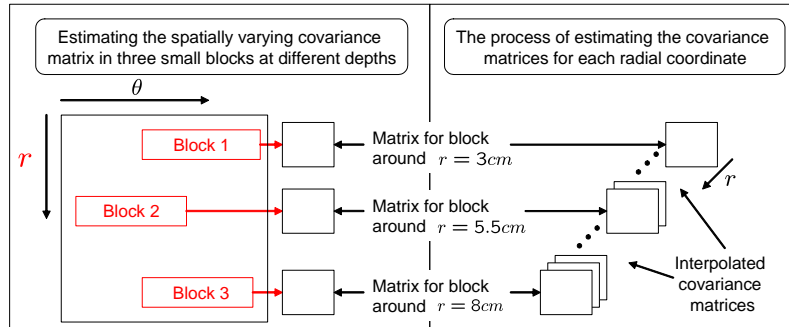
Fig. 3. Estimating the spatially varying covariance matrix.

within each of them? Practically, the answer is *No*. Since the statistics change gradually, it is possible to examine a few blocks in the field of view (FOV) as illustrated in the left side of Fig. 3, and measure the noise statistics only within them. This processing is applied in the polar coordinate space, as illustrated in Fig. 3. Then the speckle statistics around any point in the FOV can be deduced. The measurement of the statistics in these few selected blocks is described in Sec 3.4, as is the inference from these few blocks to any other region.

The BLUE requires the cross-correlation between different channels. As any cross-correlation function, it depends on the lag between pixels. When taking into account a maximum lag of $d_{\max}^{\mathrm{radial}}$ in the radial direction and a maximum lag of $d_{\max}^{\mathrm{lateral}}$ in the lateral direction, the size of the covariance matrix equals $\left(d_{\max}^{\mathrm{radial}} \cdot d_{\max}^{\mathrm{lateral}} \cdot K\right)^2$. Empirical measurements that we performed in several images showed a fast decrease in the off-diagonal elements of $\mathbf{C}_{\mathbf{s}_k^{\log}\mathbf{s}_i^{\log}}$. We conclude that the lengths of the spatial correlation are short. Hence, small lags are sufficient to reflect the statistics. We are thus allowed to use small regions, for which the radial maximum lag is ≈ 40 pixels corresponding to $\approx 1.5mm$ in our system.

We now have the statistics in a few blocks. Then, using interpolation, we infer the statistics in any region centered on any pixel in the FOV. Subsequently, we can apply the BLUE around each pixel in the image. In other words, around each pixel, we define a small region, and since the noise statistics in this region has been estimated in the previous steps, we can apply the BLUE for it, and estimate $\log\left(\mathbf{a}^{\mathrm{magnitude}}\right)$ at that location.

3.4   Measuring Statistics

We have seen in Sec. 3.3 that we use few small blocks in the image, to measure the covariance matrix, which is spatially variant. We chose blocks in which there is no meaningful object detail.[4]

---

[4] Practically, we would not expect a physician to select such blocks manually in each session. Hence, the typical covariance matrix can be learned using sets of typical speckle images of arbitrary objects. This is a matter we intend for future research.

Covariance depends on both lateral and radial lags. Furthermore, radial and lateral correlations differ. Based on the separability of the PSF [12], the covariance matrix is also separable [12]. For each matrix element $[d, q]$

$$\hat{\mathbf{C}}[d, q] = \hat{\mathbf{C}}^{\text{radial}}[d, q] \cdot \hat{\mathbf{C}}^{\text{lateral}}[d, q] \ , \tag{14}$$

where $\hat{\mathbf{C}}^{\text{radial}}$ and $\hat{\mathbf{C}}^{\text{lateral}}$ are the noise covariance matrices in the radial and lateral directions, respectively. Both matrices are measured in a similar way. For example, the cross covariance in the radial direction between two acoustic frequencies k and i, is estimated as

$$\hat{\mathbf{C}}^{\text{radial}}_{\mathbf{s}_k \mathbf{s}_i}[d, d + d^{\text{radial}}] = Z \sum_l \left\{ \mathbf{s}_k[l] - \hat{\mu}_{\mathbf{s}_k} \right\} \left\{ \mathbf{s}_i[l + d^{\text{radial}}] - \hat{\mu}_{\mathbf{s}_i} \right\} \ , \tag{15}$$

where $0 \leq d < L$, l is a pixel index, d^{radial} is the radial lag between pixels, Z is a normalization factor, L is the block length and $\hat{\mu}_{\mathbf{s}}$ is the estimated noise mean (this mean is estimated using the same data). This estimator of the covariance matrix is unbiased.

The estimated covariance functions of the selected blocks do not apply to the entire radial dimension. We still need to evaluate it in between (see the left side of Fig. 3). For this, we assume that between points in the FOV, the statistics change gradually. Hence, we can fill the missing data by interpolation. One can use interpolation methods of matrix-valued images [6], that preserve the semi-definiteness of the covariance matrix.

4   Experiment

In the experiment, we used a commercial medical ultrasonic system, the GE Vivid 3. The electronic signal generated by this system is a square burst with duration of three half periods. The probes used are phased arrays by GE, named **3s** and **5s**. The algorithm was applied on data obtained from a tissue-mimicking phantom, so that controlled and repeatable data can be generated. Fat was placed on top of the phantom to demonstrate an attenuating layer. The acquired images are presented in Fig. 1. One image was acquired with a burst frequency of $1.5MHz$ and the **3s** probe (referred to as low acoustic frequency image). The second image was acquired with a burst frequency of $2.5MHz$ and the **5s** probe (referred to as high acoustic frequency image). As illustrated in Fig. 2 speckle appearance of the same tissue changes in different acoustic frequencies. Nevertheless, in the high acoustic frequency image, system noise is very significant. We have direct access to $\mathbf{a}^{\text{RF}}$, received in MHz from the medium. We then directly apply sampling, attenuation compensation, envelope detection and log operation.

The input for the algorithm is $\mathbf{a}_k^{\text{log}}$. The BLUE was applied based on the two images, as illustrated in Fig. 4. The stochastic reconstruction significantly reduces speckle noise, along with high spatial resolution and reconstruction of deep objects. A by-product of the stochastic reconstruction is system noise reduction, due to the weighted averaging of the images. The peak signal to noise
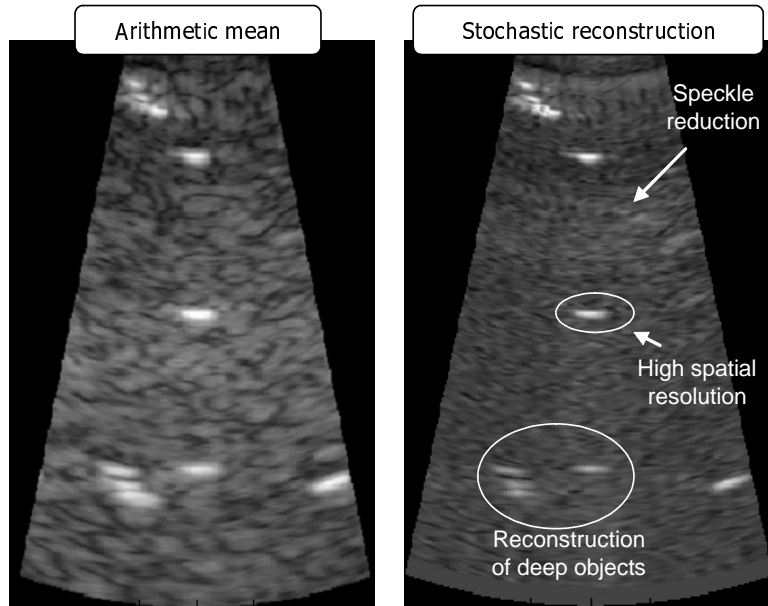
Fig. 4. Stochastic reconstruction vs. simple averaging. The stochastic reconstruction produces an image with speckle reduction, along with high spatial resolution and reconstruction of deep objects.

Table 1. In all depths, the PSNR obtained by stochastic reconstruction is higher than the PSNR obtained by arithmetic mean.

Depth (cm)	Arithmetic mean	Stochastic reconstruction
6	$66:1$ $(18dB)$	$117:1$ $(21dB)$
8	$48:1$ $(17dB)$	$73:1$ $(19dB)$
10	$78:1$ $(19dB)$	$124:1$ $(21dB)$

ratio (PSNR) was calculated. The results are presented in Table 1. The stochastic reconstruction presents a higher PSNR in all depths.

5   Discussion

The method reduces noise of ultrasound images. It also exposes deep objects while it maintains high resolution and does not blur the object to achieve denoising. Our approach requires a fast acquisition of two or more acoustic frequencies. There exists enabling technology [15] allowing that.

Future research can focus on the acquisition process as well as on the processing. In particular, it is worth studying which acoustic frequencies are optimal in this paradigm. In addition, more advanced mathematical tools can be used. For

example, diffusion methods [14] and adaptive subdivision coupled to statistical estimation [11] may be useful to this frequency compounding approach.

We wish to thank Zvi Friedman and Yonina Eldar for useful discussions.

References

1. I. Amir, N. M. Bilgutay, and V. L. Newhouse. Analysis and comparison of some frequency compounding algorithms for the reduction of ultrasonic clutter. *IEEE Trans. on Ultrasonics Ferroelectric and Frequency Control*, 4:402–411, 1986.
2. B. A. J. Angelsen. *Ultrasound Imaging Waves, Signals, and Signal Processing.* Emantec, Norway, 2000.
3. V. Aurich and J. Weule. Non-linear gaussian filters performing edge preserving diffusion. In *Proc. 17th DAGM Symposium*, pages 538–545, 1995.
4. G. Cincotti, G. Loi, , and M. Pappalardo. Frequency decomposition and compounding of ultrasound medical images with wavelet packets. *IEEE Trans. on Medical Imaging*, 20:764–771, 2001.
5. S. Gupta, R. C. Chauhan, and S. C. Sexana. Wavelet-based statistical approach for speckle reduction in medical ultrasound images. *Medical & Biological Engineering & Computing*, 42:189–192, 2004.
6. H. Hagen and J. Weickert. *Visualisation and Processing of Tensor Images.* Springer, Berlin, 2006.
7. A. K. Jain. *Fundamentals of Digital Image Processing.* Prentice-Hall, Englewood Cliffs, N.J., 1989.
8. S. M. Kay. *Fundamentals of Statistical Signal Processing. Estimation Theory.* Prentice-Hall, Englewood Cliffs, N.J., 1993.
9. T. Loupas, W. N. Mcdicken, and P. L. Allan. An adaptive weighted median filter for speckle suppression in medical ultrasonic images. *IEEE Trans. on Circuits and Systems*, 36:129–135, 1989.
10. P. A. Magnin, O. T. Von Ramm, and F. L. Thurstone. Frequency compounding for speckle contrast reduction in phased array images. *Ultrasonic Imaging*, 4:267–281, 1982.
11. J. Polzehl and V. G. Spokoiny. Adaptive weights smoothing with applications to image restoration. *J. R. Statist. Soc. B.*, 62:335–354, 2000.
12. R. F. Wagner, S. W. Smith, J. M. Sandrik, and H. Lopez. Statistics of speckle in ultrasound B-scans. *IEEE Trans. on Sonics and Ultrasonics*, 30:156–163, 1983.
13. W. F. Walker and G. E. Trahey. The application of K-space in pulse echo ultrasound. *IEEE Trans. on Ultrasonics, Ferroelectrics, and Freq. Control*, 45:541–558, 1998.
14. J. Weickert. A review of nonlinear diffusion filtering. *Lecture Notes in Computer Science*, 1252:3–28, 1997.
15. J. Xuecheng, I. Ladabaum, and T. K. Butrus. The microfabrication of capacitive ultrasonic transducers. *Journal of Microelectromechanical Systems*, 7, 1998.