

# An Iterative, Wavelet-Based Deconvolution Algorithm for the Restoration of Ultrasound Images in an EM Framework

J. K. H. Ng, R. W. Prager, N. G. Kingsbury, G. M. Treece and A. H. Gee

Department of Engineering, University of Cambridge,  
Trumpington Street, Cambridge CB2 1PZ, UK.

## ABSTRACT

The quality of medical ultrasound images is limited by inherent poor resolution due to the finite temporal bandwidth of the acoustic pulse and the non-negligible width of the system point-spread function. One of the major difficulties in designing a practical and effective restoration algorithm is to develop a model for the tissue reflectivity that can adequately capture significant image features without being computationally prohibitive. The reflectivities of biological tissues do not exhibit the piecewise smooth characteristics of natural images considered in the standard image processing literature; while the macroscopic variations in echogenicity are indeed piecewise smooth, the presence of sub-wavelength scatterers adds a pseudo-random component at the microscopic level. This observation leads us to propose modelling the tissue reflectivity as the product of a piecewise smooth echogenicity map and a unit-variance random field. The chief advantage of such an explicit representation is that it allows us to exploit representations for piecewise smooth functions (such as wavelet bases) in modelling variations in echogenicity without neglecting the microscopic pseudo-random detail. As an example of how this multiplicative model may be exploited, we propose an expectation-maximisation (EM) restoration algorithm that alternates between inverse filtering (to estimate the tissue reflectivity) and logarithmic wavelet denoising (to estimate the echogenicity map). We provide simulation and *in vitro* results to demonstrate that our proposed algorithm yields solutions that enjoy higher resolution, better contrast and greater fidelity to the tissue reflectivity compared with the current state-of-the-art in ultrasound image restoration.

## 1. INTRODUCTION

The use of ultrasound in the non-invasive imaging of soft tissue is now well established in medical diagnostics. In standard clinical practice, ultrasound images are acquired with a transceiving probe which is placed in contact with the subject to transmit radio-frequency (RF) pressure pulses into the soft tissue and to record echoes backscattered by acoustic inhomogeneities. Since the backscattered signal strength is strongly correlated with structural features, an image of the anatomy can be constructed by lining the RF echoes up in image space and displaying their amplitudes as grey-scale levels. Because these amplitude images usually have very large dynamic range, it is common to display them logarithmically compressed as so-called *B-scan* images.

Despite the popularity of ultrasound imaging as a diagnostic tool, the diagnostic utility of ultrasound images is limited by inherent blurring due to the finite temporal bandwidth of the transceiving probe and the non-negligible width of the pulse-echo acoustic beam. There is thus significant scope for the development of image restoration methods to improve diagnostic quality by correcting this blurring. Essentially, the RF image recorded by the transceiving probe is a blurred map of the *tissue reflectivity* and, in this context, image restoration is identical to the inverse problem of recovering the true tissue reflectivity given the RF image. Typical tissue reflectivities are characterised on a macroscopic scale by piecewise smooth variations in *echogenicity*, but the presence of sub-wavelength scatterers adds a pseudo-random component on a microscopic scale. The incoherent interference between pressure waves scattered by a large number of these sub-wavelength scatterers within the volume of the acoustic pulse (i.e. the *resolution cell*) gives rise to a granular texture known as *speckle*. Although speckle obscures significant image features and degrades the resolvability of structures, it also contains important textural information which is useful to clinicians for tissue identification. Hence, restoration methods for ultrasound

---

Email: [jkhn2/rwp/ngk/gmt11/ahg@eng.cam.ac.uk](mailto:jkhn2/rwp/ngk/gmt11/ahg@eng.cam.ac.uk)

images are required not only to enhance the resolvability of macroscopic structural features but also to preserve the textural information present in speckle.

Because blurring constitutes a loss of information, exact recovery of the tissue reflectivity from the RF image is impossible, but with the imposition of appropriate regularising constraints, a physically feasible approximation to the tissue reflectivity can be constructed. In a Bayesian context, these regularising constraints correspond to prior probability density functions on the tissue reflectivity and reflect our prior belief in what the tissue reflectivity should be. One of the major difficulties in designing an effective and efficient restoration algorithm is to obtain a model for tissue reflectivity which can adequately capture both macroscopic structural features and microscopic pseudo-random detail. The simplest and commonest model for tissue reflectivity is to assume Gaussian statistics which leads to the well-known *Wiener filter*.<sup>1-3</sup> Michailovich and Adam<sup>1</sup> proposed a Laplacian probability density function in place of the Gaussian on the grounds that its heavier tails admit better recovery of stronger reflections at structural boundaries.

Mainstream research in image processing has produced a variety of models for describing so-called *natural images* (such as photographs) which are predominantly piecewise smooth functions with localised discontinuities.<sup>4</sup> Models based on wavelets have been particularly successful as the wavelet transform provides a domain in which natural images can be sparsely represented and the statistics of natural images are substantially simplified (this, in turn, leads to simple but effective processing rules in the wavelet domain). Unfortunately, such models for natural images cannot adequately account for the texturally important pseudo-random behaviour that is characteristic of the acoustic response of soft tissue and so are not, by themselves, good models for tissue reflectivities.

We observe, however, that while tissue reflectivities are not strictly piecewise smooth, their echogenicities are piecewise smooth and would be well described by mainstream models of natural images. This observation leads us to suggest modelling the echogenicity and the pseudo-random component separately, the former as a piecewise smooth function and the latter as a field of independent and identically distributed random variables. In this way, we are able to assign separate priors to the echogenicity and the pseudo-random component and solve for an estimate of the tissue reflectivity in which structural features are enhanced and textural detail preserved.

The idea of modelling the echogenicity and pseudo-random component separately is not new, but was previously proposed by Husby *et al.*<sup>5</sup> who assumed Gaussian statistics for the pseudo-random component. We make the same assumption of Gaussianity, but our work differs from Husby *et al.*'s in our approach to the rest of the problem: they chose to model the echogenicity in the image domain and solved for it with Markov chain Monte Carlo (MCMC) methods, whereas we have opted to model the echogenicity in the wavelet domain and to solve for the complete tissue reflectivity using *expectation-maximisation* (EM). As we shall show, this approach reduces to alternating between Wiener filtering and wavelet-based denoising which we believe it to be more computationally efficient than MCMC.

Our hybrid Wiener/denoising scheme was inspired by the work of Figueiredo and Nowak<sup>6</sup> who also used EM to develop a similar algorithm for the deconvolution of natural images. Their work, in turn, is an extension of Neelamani *et al.*'s<sup>7,8</sup> Fourier-wavelet regularised deconvolution (ForWaRD) procedure which applies an under-regularised Wiener-style filter to the blurred image followed by wavelet-based denoising (the application of ForWaRD to ultrasound image restoration was proposed by Wan *et al.*<sup>8</sup>). We emphasise that, although these techniques share a procedural similarity with our proposed algorithm, they were designed for natural images and, unlike our algorithm, they do not assume any pseudo-random behaviour on a microscopic level.

In the sections that follow, we shall first introduce a linear description for the blurring operator and our model for the tissue reflectivity. We shall then present the EM algorithm as a framework for statistical inference in problems with unobservable hidden parameters and show that, for our problem, the EM algorithm reduces to alternating between Wiener filtering (to update the reflectivity estimate) and wavelet-based denoising (to update the echogenicity estimate). We shall present simulation and *in vitro* results to compare the performance of our algorithm with the current state-of-the-art in ultrasound image restoration.

## 2. IMAGE MODELLING

Since we deal entirely with digitised quantities, we adopt matrix-vector notation for our mathematical description of the blurring operator. Let  $\mathbf{x}$  be an  $N \times 1$  vector representing the sampled tissue reflectivity, let  $\mathbf{y}$  be an  $N \times 1$

vector representing the sampled RF image and let  $\mathbf{n}$  be an  $N \times 1$  noise vector to account for measurement error. If we assume linear blurring, then we may write

$$\mathbf{y} = H\mathbf{x} + \mathbf{n} \quad (1)$$

where  $H$  is the  $N \times N$  blurring matrix. It is usual to assume  $\mathbf{n}$  to be normally distributed with zero mean and covariance matrix  $C_n$ . In this context, the problem of restoration is to estimate  $\mathbf{x}$  given  $\mathbf{y}$ ,  $H$  and  $C_n$ .

The conventional way to cast the problem into a Bayesian framework is to treat  $\mathbf{y}$  and  $\mathbf{x}$  as random vectors and to solve for the maximum *a posteriori* (MAP) estimate of  $\mathbf{x}$  given  $\mathbf{y}$ , i.e. we seek the realisation  $\hat{\mathbf{x}}$  of  $\mathbf{x}$  which maximises its posterior probability  $p(\mathbf{x} | \mathbf{y})$ . Bayes's rule states that  $p(\mathbf{x} | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x})p(\mathbf{x})$ ; taking logarithms, the MAP estimation may be written as

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} [\ln p(\mathbf{y} | \mathbf{x}) + \ln p(\mathbf{x})] \quad (2)$$

where the log-likelihood  $\ln p(\mathbf{y} | \mathbf{x})$  enforces fidelity to the observed data and the log-prior  $\ln p(\mathbf{x})$  is a regularising constraint that reflects our prior belief. In our case, where we have assumed additive Gaussian noise,

$$p(\mathbf{y} | \mathbf{x}) \propto \exp \left[ -\frac{1}{2} (\mathbf{y} - H\mathbf{x})^T C_n^{-1} (\mathbf{y} - H\mathbf{x}) \right]. \quad (3)$$

If we model  $\mathbf{x}$  as a zero-mean Gaussian random vector with covariance matrix  $C_x$ ,

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \left[ \frac{1}{2} (\mathbf{y} - H\mathbf{x})^T C_n^{-1} (\mathbf{y} - H\mathbf{x}) + \frac{1}{2} \mathbf{x}^T C_x^{-1} \mathbf{x} \right] \quad (4)$$

to which the closed-form solution is the Wiener filter,

$$\hat{\mathbf{x}} = (H^T C_n^{-1} H + C_x^{-1})^{-1} H^T C_n^{-1} \mathbf{y}. \quad (5)$$

The special case  $C_x = \sigma_x^2 I_N$ , where  $I_N$  is the  $N \times N$  identity matrix, corresponds to a regularising constraint on the  $l_2$  norm of  $\mathbf{x}$ , which is referred to as *zero-order Tikhonov regularisation*.

If we model  $\mathbf{x}$  as a vector of independent, identically distributed Laplacian random variables with variance  $\sigma_x^2$ , each  $x_i = [\mathbf{x}]_i$  has probability density function,

$$p(x_i) \propto \exp \left( -\frac{\sqrt{2}}{\sigma_x} |x_i| \right) \quad (6)$$

and we end up constraining the  $l_1$  norm instead of the  $l_2$  norm:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \left\{ \frac{1}{2} (\mathbf{y} - H\mathbf{x})^T C_n^{-1} (\mathbf{y} - H\mathbf{x}) + \frac{\sqrt{2}}{\sigma_x} \sum_{i=1}^N |x_i| \right\}. \quad (7)$$

### 3. IMAGE RESTORATION

Instead of following the conventional approach of assigning a single prior to  $\mathbf{x}$ , we propose modelling  $\mathbf{x}$  as a product of its echogenicity and a random component. In matrix-vector notation, we write

$$\mathbf{x} = S\mathbf{w} \quad (8)$$

where  $S$  is an  $N \times N$  diagonal matrix of non-negative echogenicity values and  $\mathbf{w}$  is an  $N \times 1$  random vector. We propose treating the elements of  $\mathbf{w}$  as independent and identically distributed Gaussian random variables with zero mean and unit variance, which is consistent with the models of Michailovich and Adam<sup>1</sup> and Husby *et al.*<sup>5</sup> in which the tissue reflectivity in a uniformly echogenic region is considered to behave as white noise. Under this Gaussian assumption, we may write

$$p(\mathbf{x} | S) \propto \frac{1}{|S|} \exp \left( -\frac{1}{2} \mathbf{x}^T S^{-2} \mathbf{x} \right). \quad (9)$$

At this point, we note the following two important observations:

- If  $S$  were known exactly, then  $\mathbf{x}$  is just a Gaussian random vector with covariance matrix  $S^2$  and can be estimated from the RF image  $\mathbf{y}$  by applying a Wiener filter as given in (5).
- Conversely, if  $\mathbf{x}$  were known exactly, then we can estimate  $S$  by treating  $\mathbf{w}$  as multiplicative noise and applying a suitable denoising procedure to get rid of it.

These observations suggest an iterative restoration procedure in which we alternate between estimating  $\mathbf{x}$  and estimating  $S$  to get successively better estimates of both. This approach forms the basic structure of our restoration algorithm; to derive the exact update rules for  $\mathbf{x}$  and  $S$ , we turn to the EM algorithm which is ideally suited to problems where there are unobservable hidden parameters.

### 3.1. Expectation-Maximisation

The EM algorithm is an iterative framework for *maximum likelihood* (ML) or MAP estimation in problems where there is an unobservable hidden parameter in which we are not interested (we shall refer to this hidden parameter as a *nuisance variable*). Each iteration of the EM algorithm yields an estimate which increases the likelihood/posterior distribution of the parameter of interest and convergence to a local maximum of the likelihood/posterior distribution is guaranteed.

We present, without proof, the mechanics of the EM algorithm for MAP estimation as described by Dellaert.<sup>9</sup> Let  $\Theta$  be the parameter we wish to estimate, let  $\mathbf{U}$  be the observed data from which we wish to estimate  $\Theta$  and let  $\mathbf{J}$  be the nuisance variable. At the  $k$ th iteration, we execute:

- The *E-step*: Calculate the expected joint log-likelihood of  $\mathbf{U}$  and  $\mathbf{J}$  given  $\Theta$ ,

$$Q(\Theta | \hat{\Theta}_k) = \mathbb{E} \left[ \ln p(\mathbf{U}, \mathbf{J} | \Theta) | \mathbf{U}, \hat{\Theta}_k \right] = \int p(\mathbf{J} | \mathbf{U}, \hat{\Theta}_k) \ln p(\mathbf{U}, \mathbf{J} | \Theta) d\mathbf{J}. \quad (10)$$

- The *M-step*: Calculate the MAP estimate of  $\Theta$ ,

$$\hat{\Theta}_{k+1} = \arg \max_{\Theta} \left[ Q(\Theta | \hat{\Theta}_k) + \ln p(\Theta) \right]. \quad (11)$$

For our image restoration problem, we assign  $S$  to be the parameter for which we wish to find the MAP estimate and  $\mathbf{x}$  to be the nuisance variable. Given that we are ultimately interested in estimating  $\mathbf{x}$  and not  $S$ , this allocation of parameters may seem counter-intuitive, but as we shall demonstrate, it leads to a very elegant alternation between Wiener filtering and logarithmic denoising.

### 3.2. The E-step: Wiener Filtering

Substituting  $\mathbf{U} = \mathbf{y}$ ,  $\mathbf{J} = \mathbf{x}$  and  $\Theta = S$  into (10), we obtain

$$Q(S | \hat{S}_k) = \int p(\mathbf{x} | \mathbf{y}, \hat{S}_k) \ln p(\mathbf{y}, \mathbf{x} | S) d\mathbf{x}. \quad (12)$$

To obtain a closed-form expression for this integral, we first turn our attention to finding an expression for the ‘nuisance posterior’  $p(\mathbf{x} | \mathbf{y}, \hat{S}_k)$ . Applying Bayes’s rule and recognising that  $p(\mathbf{y} | \mathbf{x}, \hat{S}_k) = p(\mathbf{y} | \mathbf{x})$ , we have

$$p(\mathbf{x} | \mathbf{y}, \hat{S}_k) = \frac{p(\mathbf{y} | \mathbf{x}, \hat{S}_k) p(\mathbf{x} | \hat{S}_k)}{p(\mathbf{y} | \hat{S}_k)} = \frac{p(\mathbf{y} | \mathbf{x}) p(\mathbf{x} | \hat{S}_k)}{p(\mathbf{y} | \hat{S}_k)}. \quad (13)$$

The denominator in the rightmost term of (13) does not depend on  $\mathbf{x}$  and may be regarded as just a normalisation constant. Substituting (3) and (9) into the numerator and simplifying, we obtain

$$p(\mathbf{x} | \mathbf{y}, \hat{S}_k) \propto \exp \left\{ -\frac{1}{2} \left[ \mathbf{x}^T \left( H^T C_n^{-1} H + \hat{S}_k^{-2} \right) \mathbf{x} - 2 \mathbf{x}^T H^T C_n^{-1} \mathbf{y} \right] \right\} \quad (14)$$



which is a Gaussian distribution with mean  $\mathbf{m}_k$  and covariance matrix  $C_x$  given by

$$C_k = \left( H^T C_n^{-1} H + \widehat{S}_k^{-2} \right)^{-1}, \quad \mathbf{m}_k = C_x H^T C_n^{-1} \mathbf{y} = \left( H^T C_n^{-1} H + \widehat{S}_k^{-2} \right)^{-1} H^T C_n^{-1} \mathbf{y}. \quad (15)$$

We note that the mean  $\mathbf{m}_k$  is, by definition, the *minimum mean squared error* (MMSE) estimate of  $\mathbf{x}$  given the echogenicity estimate  $\widehat{S}_k$  and is therefore the best estimate of the tissue reflectivity at the  $k$ th iteration. We also point out that, comparing the expression for  $\mathbf{m}_k$  in (15) to the expression for the Wiener filter in (5), the calculation of  $\mathbf{m}_k$  is the same as applying a Wiener filter derived from  $\widehat{S}_k$  to the RF image  $\mathbf{y}$ .

Returning now to the calculation of the integral in (12), we note that  $p(\mathbf{y}, \mathbf{x} | S) = p(\mathbf{y} | \mathbf{x}) p(\mathbf{x} | S) \Rightarrow \ln p(\mathbf{y}, \mathbf{x} | S) = \ln p(\mathbf{y} | \mathbf{x}) + \ln p(\mathbf{x} | S)$ . The term  $\ln p(\mathbf{y} | \mathbf{x})$  does not depend on  $S$  and can therefore be dropped from the definition of  $Q(S | \widehat{S}_k)$ . Substituting (9) into (12) and discarding constants, we are left with

$$Q(S | \widehat{S}_k) = -\ln |S| - \frac{1}{2} \int \mathbf{x}^T S^{-2} \mathbf{x} p(\mathbf{x} | \mathbf{y}, \widehat{S}_k) d\mathbf{x}. \quad (16)$$

Finally, we recognise that the remaining integral in (16) is just the conditional expectation  $E(\mathbf{x}^T S^{-2} \mathbf{x} | \mathbf{y}, \widehat{S}_k)$ . We apply result (7.2.2) from Ref. 10 which states that, for a Gaussian random vector  $\mathbf{z}$  with mean  $\mathbf{m}$  and covariance matrix  $C$ ,  $E(\mathbf{z}^T A \mathbf{z}) = \text{Tr}(AC) + \mathbf{m}^T A \mathbf{m}$ , and we obtain

$$Q(S | \widehat{S}_k) = -\ln |S| - \frac{1}{2} [\text{Tr}(S^{-2} C_k) + \mathbf{m}_k^T S^{-2} \mathbf{m}_k]. \quad (17)$$

### 3.3. The M-step: Logarithmic Denoising

Before we launch into the derivation of the update rule for  $\widehat{S}$  in the M-step, we detour briefly to discuss the problem of estimating  $S$  given  $\mathbf{x}$  by treating  $\mathbf{w}$  in (8) as multiplicative noise to be removed. We define  $x_i = [\mathbf{x}]_i$ ,  $s_i = [S]_{ii}$  and  $w_i = [\mathbf{w}]_i$  and we rewrite (8) component-wise as  $x_i = s_i w_i$ . To turn  $\{w_i\}$  into additive noise, we take the logarithms of the moduli of both sides; defining  $\tilde{x}_i = \ln |x_i|$ ,  $\tilde{s}_i = \ln |s_i|$  and  $\tilde{w}_i = \ln |w_i|$ , we obtain

$$\tilde{x}_i = \tilde{s}_i + \tilde{w}_i, \quad i = 1, \dots, N. \quad (18)$$

The logarithmic noise term  $\{\tilde{w}_i\}$  has probability density function, mean and variance given by

$$p(\tilde{w}_i) = \sqrt{\frac{2}{\pi}} \exp \left[ \tilde{w}_i - \frac{1}{2} \exp(2\tilde{w}_i) \right], \quad E(\tilde{w}_i) = -\frac{1}{2} (\gamma + \ln 2), \quad \text{Var}(\tilde{w}_i) = \frac{\pi^2}{8} \quad (19)$$

where  $\gamma$  is the Euler-Mascheroni constant and has an approximate value of 0.5772. The derivation of this probability density function and the calculation of its mean and variance are detailed in Appendix A.

Since the echogenicity is piecewise smooth, we expect the log-echogenicity  $\{\tilde{s}_i\}$  to also be piecewise smooth and to have a sparse representation in the wavelet domain. We can therefore denoise by applying wavelet shrinkage: since most of the energy of  $\{\tilde{s}_i\}$  will be concentrated into just a few wavelet coefficients, we can modify the wavelet coefficients of  $\{\tilde{x}_i\}$  according to some shrinkage rule that attenuates the logarithmic noise term  $\{\tilde{w}_i\}$ . Most wavelet shrinkage rules are based on the assumption that the wavelet coefficients of the additive noise are Gaussian, which at first sight seems to be violated by the non-Gaussianity of the probability density function in (19). In practice, however, we have found that, because of the band-limiting of each wavelet sub-band, the central limit theorem keeps the wavelet coefficients of  $\{\tilde{w}_i\}$  approximately Gaussian. The additive noise is also usually assumed to have zero mean, so we need to add  $\frac{1}{2} (\gamma + \ln 2)$  to  $\{\tilde{x}_i\}$  before applying wavelet shrinkage.

It is well known that wavelet shrinkage corresponds to MAP estimation with a wavelet-domain prior, the exact form of which depends on the specific shrinkage rule used.<sup>11</sup> Hence, we may regard the logarithmic denoising of  $\{\tilde{x}_i\}$  to recover  $\{\tilde{s}_i\}$  (from which we estimate  $\{s_i\}$ ) as being equivalent to the MAP estimation problem

$$\begin{aligned} \widehat{S} &= \arg \max_S p(S | \mathbf{x}) = \arg \max_S p(\mathbf{x} | S) p(S) = \arg \max_S [\ln p(\mathbf{x} | S) + \ln p(S)] \\ &= \arg \max_S \left[ -\ln |S| - \frac{1}{2} \mathbf{x}^T S^{-2} \mathbf{x} + \ln p(S) \right] = \arg \max_S \left[ -\sum_{i=1}^N \left( \ln s_i + \frac{x_i^2}{2s_i^2} \right) + \ln p(S) \right] \end{aligned} \quad (20)$$

when the prior  $p(S)$  is defined in terms of the wavelet coefficients of the log-echogenicity  $\{\tilde{s}_i\}$ .

Returning now to the derivation of the update rule for  $\hat{S}$ , we substitute  $\Theta = S$  and (17) into (11),

$$\hat{S}_{k+1} = \arg \max_S \left\{ -\ln |S| - \frac{1}{2} [\text{Tr}(S^{-2}C_k) + \mathbf{m}_k^T S^{-2} \mathbf{m}_k] + \ln p(S) \right\}. \quad (21)$$

If we define  $\sigma_{k,i}^2 = [C_k]_{ii}$  and  $m_{k,i} = [\mathbf{m}_k]_i$ , we may rewrite the update rule as

$$\hat{S}_{k+1} = \arg \max_S \left[ -\sum_{i=1}^N \left( \ln s_i + \frac{m_{k,i}^2 + \sigma_{k,i}^2}{2s_i^2} \right) + \ln p(S) \right]. \quad (22)$$

If we now compare this form of the update rule with (20), we see that the two expressions are identical if we let  $x_i^2 = m_{k,i}^2 + \sigma_{k,i}^2$ , and in light of our previous discussion on logarithmic denoising, we conclude that we can calculate our next estimate of  $S$  simply by applying wavelet shrinkage to  $\left\{ \ln \left( \sqrt{m_{k,i}^2 + \sigma_{k,i}^2} \right) \right\}$ .

### 3.4. Summary of the Algorithm

We summarise our algorithm for ultrasound image restoration as follows (for notational convenience, we have dropped the subscript  $k$ ):

```

Initialise  $\hat{S}$ .
While termination condition is not satisfied
  E-step:
  Calculate  $\mathbf{m} = \left( H^T C_n^{-1} H + \hat{S}^{-2} \right)^{-1} H^T C_n^{-1} \mathbf{y}$ .
  Calculate  $\sigma_i = \left[ \left( H^T C_n^{-1} H + \hat{S}^{-2} \right)^{-1} \right]_{ii}$  for  $i = 1, \dots, N$ .
  M-step:
  Calculate  $\tilde{x}_i = \ln \left( \sqrt{m_i^2 + \sigma_i^2} \right) = \frac{1}{2} \ln (m_i^2 + \sigma_i^2)$  for  $i = 1, \dots, N$ .
  Estimate  $\{\tilde{s}_i\}$  by applying wavelet shrinkage to  $\{\tilde{x}_i + \frac{1}{2}(\gamma + \ln 2)\}$ .
  Calculate  $\hat{s}_i = \exp(\tilde{s}_i)$  for  $i = 1, \dots, N$ .
end

```

Once the algorithm has terminated, we take  $\mathbf{m}$  as our Bayesian estimate of the tissue reflectivity. We also get the MAP estimate  $\{\hat{s}_i\}$  of the echogenicity, which is essentially a despeckled copy of the envelope-detected RF image. We note that the echogenicity image, although devoid of any textural information, can be useful for examining and segmenting gross anatomical features.

### 3.5. Computational Issues

So far, we have not discussed the implementation of the blurring matrix  $H$ . Because of the typically large size of  $N$ , it is usually not feasible to compute  $H$  explicitly and multiplication by  $H$  or  $H^T$  is performed indirectly. In problems where the blurring operator is shift-invariant and can be characterised by a single *impulse response* or *point-spread function* (PSF), the blurring matrix  $H$  is block-circulant with circulant blocks (BCCB). The product  $H\mathbf{x}$  is then just the same as the convolution of the image represented by  $\mathbf{x}$  with the PSF and can be efficiently computed by multiplying the discrete Fourier transform (DFT) coefficients of  $\mathbf{x}$  with the DFT coefficients of the PSF and taking the inverse DFT of the product. The matrix  $H^T$  denotes convolution with a spatially reversed

copy of the PSF which corresponds to multiplication by the complex conjugate of the DFT coefficients of the PSF in the DFT domain.

Unfortunately, the width of the acoustic beam in ultrasound imaging is not spatially uniform but varies with axial distance, giving rise to a shift-variant blurring operator and a blurring matrix  $H$  that is not BCCB. To handle a shift-variant blurring operator, Nagy and O’Leary<sup>12,13</sup> suggested partitioning image space into a number of smaller non-overlapping regions within which the blurring operator is considered to be shift-invariant. Each of these regions can be characterised by its own PSF, and the blurring operator can be approximated by convolving each region with its PSF, weighting the results and summing them. Mathematically, we write  $H \approx \sum_n D_n H_n$  where  $H_n$  is the BCCB matrix corresponding to the  $n$ th region and  $D_n$  is a non-negative diagonal weighting matrix that satisfies  $\sum_n D_n = I_N$ . Correspondingly, we have  $H^T \approx \sum_n H_n^T D_n$  which is equivalent to weighting each region, convolving each weighted region with a spatially reversed copy of its PSF and summing the results. Since  $H$  and  $H^T$  can neither be computed explicitly nor computed indirectly via the DFT, we are forced to solve the linear system for the conditional mean  $\mathbf{m}_k$  in (15) iteratively. We chose to use the method of preconditioned conjugate gradients which solves any symmetric, positive definite linear system by minimising an equivalent multidimensional quadratic function (for a more in-depth discussion of conjugate gradients for solving linear systems, see sections 2.7 and 10.6 of Ref. 14). We also note that we cannot explicitly compute the conditional covariance matrix  $C_k$  in (15) and consequently, the diagonal elements  $\{\sigma_{k,i}^2\}$  of  $C_k$  are inaccessible to us. To overcome this problem, we recognise that each  $\sigma_{k,i}^2$  is just the variance of  $x_i$  given  $\mathbf{y}$  and  $\hat{S}_k$ , and since  $\mathbf{m}_k$  is our best estimate of  $\mathbf{x}$  given  $\mathbf{y}$  and  $\hat{S}_k$ , we assume local ergodicity and estimate each  $\sigma_{k,i}^2$  by computing the sample variance over a local neighbourhood in the image  $\mathbf{m}_k$ .

#### 4. EXPERIMENTAL RESULTS

To assess the effectiveness of our proposed restoration algorithm, we tested it on simulated and *in vitro* datasets and compared its performance to zero-order Tikhonov regularisation, Laplacian regularisation and ForWaRD (the latter two are generally considered to be state-of-the-art for ultrasound image restoration).

We assumed the additive noise  $\mathbf{n}$  in (1) to be white with variance  $\sigma_n^2$  (i.e.  $C_n = \sigma_n^2 I_N$ ) and we estimated  $\sigma_n^2$  by applying the wavelet-based median absolute deviation estimator described in section 7.3 of Ref. 15. We scaled the blurring matrix  $H$  to satisfy  $\|H\mathbf{o}\|^2 = N$ , where  $\mathbf{o}$  is an  $N \times 1$  vector of ones and, whenever an estimate of the global image variance  $\sigma_x^2$  was required, we took  $\sigma_x^2 \approx N^{-1} \|\mathbf{y}\|^2 - \sigma_n^2$ .

For wavelet-based processing, we used Kingsbury’s *dual-tree complex wavelet transform* (DTCWT)<sup>16,17</sup> which, unlike the conventional discrete wavelet transform, is shift-invariant and directionally selective with minimal redundancy and only slightly heavier computational requirements. To denoise, we used Lendur and Selesnick’s bivariate shrinkage rule<sup>11,18</sup> which, when used in conjunction with the DTCWT, gives state-of-the-art denoising.

To compute the zero-order Tikhonov and ForWaRD solutions, we substituted  $C_x = \sigma_x^2 I_N$  into (5) and solved the resulting linear system with the method of preconditioned conjugate gradients as described in Section 3.5. In the case of ForWaRD, we replaced  $\sigma_x^2$  with  $\alpha\sigma_x^2$ , where  $\alpha \in [0, 1]$  is an under-regularisation parameter and we selected  $\alpha = 0.2$  as recommended in Ref. 7. This Wiener-style inverse filtering was followed by bivariate shrinkage in the DTCWT domain to remove additive noise amplified by the under-regularisation.

To compute the Laplacian-regularised solution, we removed the singularity of the modulus operator at the origin by replacing it with its convex smooth approximation,  $|x| \approx \sqrt{x^2 + \epsilon}$ , where  $0 < \epsilon \ll 1$ , as recommended by Michailovich and Adam.<sup>1</sup> We minimised the cost function in (7) with the Polak-Ribiere variant of the preconditioned conjugate gradients algorithm as described in Section 10.6 of Ref. 14 and Chapter 5 of Ref. 19.

For the EM algorithm, we initialised  $\hat{S}_0$  to  $\sigma_x^2 I_N$ , making the E-step in the first iteration identical to zero-order Tikhonov regularisation. In each E-step, we used the conjugate gradients algorithm to perform the Wiener filtering, terminating when  $\|C_k \mathbf{m}_k - H^T C_n^{-1} \mathbf{y}\| < 0.01 \|H^T C_n^{-1} \mathbf{y}\|$ ; this is equivalent to allowing a relative error of 1% in our calculation of the conditional mean  $\mathbf{m}_k$ . The conditional mean  $\mathbf{m}_k$  from the current iteration was then used to initialise the conjugate gradients algorithm in the next E-step. In each M-step, we used bivariate wavelet shrinkage in the DTCWT domain to update the echogenicity estimate.

## 4.1. Simulation Results I

Our first set of results was produced in simulation. We created a tissue reflectivity by multiplying a drawing of a cross-section of a human kidney by an array of independent realisations of a zero-mean, unit-variance Gaussian distribution. We partitioned image space axially into 12 non-overlapping regions and calculated a PSF for each region by simulating the response to appropriately positioned point scatterers in Field II<sup>20</sup> (we used the parameters for the transceiving probe described in Section 4.3). We applied a shift-variant blur (in the manner described in Section 3.5) to this tissue reflectivity and added white Gaussian noise to give the resulting image a blurred signal-to-noise ratio of 20dB.

In Figure 1, we display B-scan images of the original drawing of the kidney’s cross-section, the tissue reflectivity (original and corrupted) and the various restoration results. To quantify the quality of each restoration, we calculated the improvement in signal-to-noise ratio (ISNR) in dB according to the formula

$$\text{ISNR} = 10 \log_{10} \left( \frac{\|\widehat{\mathbf{x}} - \mathbf{x}\|^2}{\|\mathbf{y} - \mathbf{x}\|^2} \right). \quad (23)$$

Our results in Figure 1 show that the EM algorithm outperforms the other restoration methods in both visual quality and in ISNR; the improvements in contrast and resolution are clearly visible in the EM image.

We ran our restoration algorithms in Matlab 7 on a personal computer with a 3.2GHz processor and 1GB of memory. In Figure 3(a), we have plotted the evolution of ISNR against program execution time for the zero-order Tikhonov, Laplacian and EM restorations (ForWaRD was omitted because its ISNR can only be sensibly computed at termination after wavelet shrinkage). We point out that the EM algorithm converged to its final ISNR in two iterations, which we find very encouraging, and since we have designed the first iteration of EM to be identical to zero-order Tikhonov regularisation, we consider the second iteration to be a small computational penalty to pay for significant improvement in image quality.

## 4.2. Simulation Results II

Our second set of results was also produced in simulation. This time, we used a photograph of a cross-section of a human heart’s left ventricle as our echogenicity image. We created a tissue reflectivity and blurred it in the same way as we did for the image of the kidney, but this time, we partitioned image space axially into 11 non-overlapping regions instead of 12. This set of simulation results is displayed in Figure 2 and, as before, we have quantified the quality of each restoration by its ISNR. We see that the EM algorithm has again outperformed the other restoration methods in ISNR and in visual quality: the EM solution has a 0.31dB advantage over the next best solution and, visually, the EM image exhibits better contrast, particularly between the dark region in the middle and the surrounding cardiac tissue. We have also plotted the evolution of ISNR against execution time in Figure 3(b), and we see once again that the EM algorithm achieves most of its ISNR in just two iterations.

## 4.3. *In Vitro* Results

We also applied our restoration algorithm to an *in vitro* image of a phantom containing spherical regions with echogenicities different to the background. The image was acquired with a 127-element probe with a nominal passband from 5 to 10MHz and an active (lateral) length of 40mm. The probe was operated with a single lateral focus on transmission and reception; elevational (out-of-plane) focussing was accomplished with an acoustic lens. The beam-formed RF traces were sampled at 66.6MHz. To approximate the blurring operator, we partitioned image space axially into 12 non-overlapping regions and calculated the PSF for each region by simulating the response to a point scatterer in Field II.

We present B-scan images of the phantom before and after restoration in Figure 4. We note, first of all, that all of the restoration methods have successfully corrected the diffraction effects at the axial extremes and that the circular regions there no longer look distorted. Secondly, we point out that, in this particular experiment, the EM algorithm converged to its final solution within one iteration and thus the EM image and the Tikhonov image are identical. We can see that the quality of this image is better than the ForWaRD and Laplacian images; the circular region in the top-left corner of the ForWaRD image has severe artifacts, while the circular region at the bottom of the Tikhonov image has been obscured by increased granularity in texture. Finally, we note that the EM algorithm has, within one iteration, reconstructed an image of the echogenicity which is free of speckle and in which the boundaries of the circular regions are well defined.

## 5. CONCLUSIONS

In this paper, we addressed the problem of image restoration for improving the diagnostic quality of medical ultrasound images. This is essentially the inverse problem of estimating the reflectivity of an insonified region of tissue from echoes backscattered by acoustic inhomogeneities. We highlighted the importance of having a model of tissue reflectivity which can adequately capture both macroscopic structural features and the textural information present in speckle. To this end, we proposed modelling tissue reflectivity as the product of its echogenicity and a pseudo-random component. As the echogenicity is typically piecewise smooth, it would very likely lend itself to an efficient representation in the wavelet domain. We suggested treating the pseudo-random component as a field of independent and identically distributed Gaussian random variables.

Casting the restoration problem into a Bayesian framework, we showed that adopting EM to solve for the MAP estimate of the echogenicity led to an elegant iterative algorithm that alternates between Wiener filtering to estimate the tissue reflectivity and wavelet-based denoising to estimate the echogenicity. Although we are primarily interested in the estimate of the complete tissue reflectivity for medical diagnostics, we note that the estimate of the echogenicity image is essentially a despeckled copy of the envelope-detected scan and can be useful for examination and segmentation of gross anatomical features.

We conducted a number of experiments in simulation and *in vitro* to compare our EM restoration algorithm to zero-order Tikhonov regularisation, Laplacian regularisation and ForWaRD. Our simulation results show that our algorithm produces restorations which have higher ISNR and better visual quality than these other techniques. This finding is confirmed by our *in vitro* results which showed the EM image to have better visual quality than its Laplacian and ForWaRD counterparts. Our *in vitro* results also confirmed the ability of our restoration algorithm to produce a good estimate of the echogenicity in which macroscopic structural features are well defined.

We conclude that our proposed algorithm is competitive with the current state-of-the-art in ultrasound image restoration and can, with minor computational penalty, produce results that have superior image quality in terms of resolution, contrast and fidelity to the tissue's acoustic response. Our algorithm is not restricted to two-dimensional images, and work is currently underway to apply it to three-dimensional datasets where elevational (out-of-plane) blurring can also be accounted for.

Software and data for the simulations and the *in vitro* experiment described in this paper can be downloaded from [www-sigproc.eng.cam.ac.uk/~jkh2/SPIEMedImag/](http://www-sigproc.eng.cam.ac.uk/~jkh2/SPIEMedImag/).

### APPENDIX A. DERIVATION OF THE STATISTICS OF LOGARITHMIC NOISE

In this appendix we derive the probability density function and the statistics quoted in (19) for the logarithmic noise term  $\{w_i\}$  in (18). We first state the following three useful results,

$$\int_0^{\infty} \exp(-x^2) dx = \frac{\sqrt{\pi}}{2} \quad (24)$$

$$\int_0^{\infty} \ln x \exp(-\mu x^2) dx = -\frac{1}{4} \sqrt{\frac{\pi}{\mu}} (\gamma + 2\mu \ln 2) \quad (25)$$

$$\int_0^{\infty} (\ln x)^2 \exp(-x^2) dx = \frac{\sqrt{\pi}}{8} \left[ (\gamma + 2 \ln 2)^2 + \frac{\pi^2}{2} \right] \quad (26)$$

where  $\gamma$  is the Euler-Mascheroni constant. (24) is the well-known Gaussian integral and is easily derived simply by requiring the probability density function of the Gaussian distribution to integrate to 1 over the real number line. (25) and (26) correspond respectively to results (4.333) and (4.335-2) in Ref. 21.

To derive the probability density function in (18), we define a Gaussian random variable  $W$  with zero mean and unit variance, and we define  $Z = |W|$  and  $\widetilde{W} = \ln Z$ . We have  $p_Z(z) = 2 p_W(z) = \sqrt{2\pi}^{-1} \exp(-\frac{1}{2}z^2)$  for  $z \geq 0$  and  $p_Z(z) = 0$  otherwise, and

$$p_{\widetilde{W}}(\tilde{w}) = p_Z(z) \left. \frac{dz}{d\tilde{w}} \right|_{z=e^{\tilde{w}}} = \sqrt{\frac{2}{\pi}} \exp\left(-\frac{e^{2\tilde{w}}}{2}\right) \exp(\tilde{w}) = \sqrt{\frac{2}{\pi}} \exp\left(-\frac{e^{2\tilde{w}}}{2} + \tilde{w}\right).$$

The mean of  $\widetilde{W}$  is trivial to calculate; we just apply (25) with  $\mu = \frac{1}{2}$ .

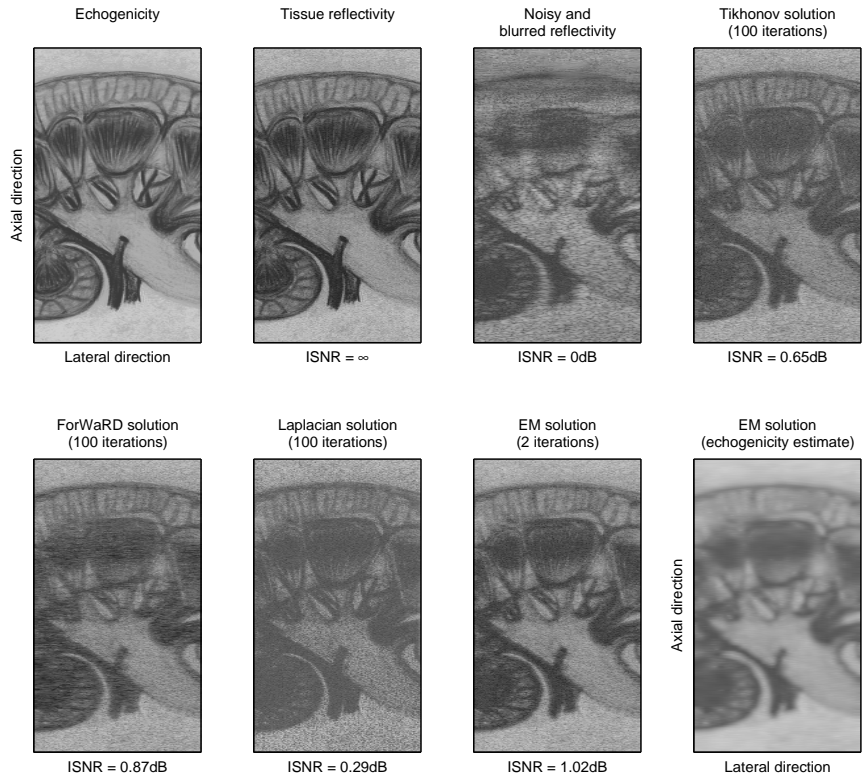
To calculate the variance of  $\widetilde{W}$ , we first derive an expression for its second moment:

$$\begin{aligned} E(\widetilde{W}^2) &= \int_{-\infty}^{\infty} (\ln z)^2 p_Z(z) dz = \sqrt{\frac{2}{\pi}} \int_0^{\infty} (\ln z)^2 e^{-\frac{z^2}{2}} dz = \sqrt{\frac{2}{\pi}} \int_0^{\infty} [\ln(\sqrt{2}z)]^2 e^{-z^2} dz \\ &= \frac{2}{\sqrt{\pi}} \left[ \frac{(\ln 2)^2}{4} \int_0^{\infty} e^{-z^2} dz + \ln 2 \int_0^{\infty} \ln z e^{-z^2} dz + \int_0^{\infty} (\ln z)^2 e^{-z^2} dz \right]. \end{aligned}$$


From (24), (25) (with  $\mu = 1$ ) and (26), it is easy to show that  $E(\widetilde{W}^2)$  reduces to  $\frac{1}{4}(\gamma + \ln 2)^2 + \frac{\pi^2}{8}$ . Subtracting the square of the mean from this gives the variance as  $\frac{\pi^2}{8}$ .

## REFERENCES

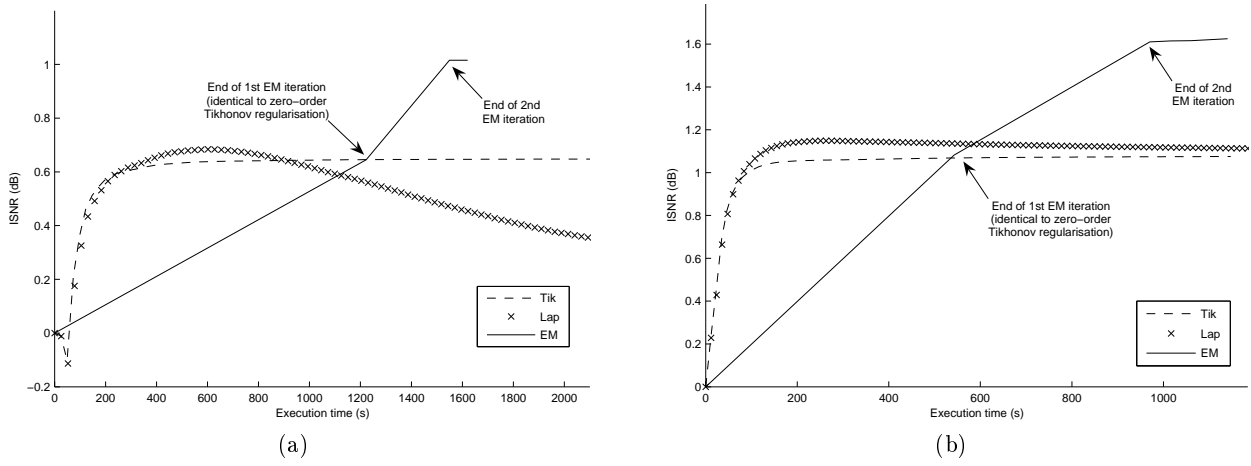
1. O. Michailovich and D. Adam, "A novel approach to the 2-D blind deconvolution problem in medical ultrasound," *IEEE Trans. Med. Imag.* **24**, pp. 86–104, Jan. 2005.
2. T. Taxt, "Three-dimensional blind deconvolution of ultrasound images," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* **48**, pp. 867–871, July 2001.
3. T. Taxt and J. Strand, "Two-dimensional noise-robust blind deconvolution of ultrasound images," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* **48**, pp. 861–866, July 2001.
4. A. Srivastava, A. B. Lee, E. P. Simoncelli, and S. C. Zhu, "On advances in statistical modeling of natural images," *Journal of Mathematical Imaging and Vision* **18**, pp. 17–33, 2003.
5. O. Husby, T. Lie, T. Lango, J. Hokland, and H. Rue, "Bayesian 2-D deconvolution: A model for diffuse ultrasound scattering," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* **48**, pp. 121–130, Jan. 2001.
6. M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Processing* **12**, pp. 906–916, Aug. 2003.
7. R. Neelamani, H. Choi, and R. Baraniuk, "ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems," *IEEE Trans. Signal Processing* **52**, pp. 418–433, Feb. 2004.
8. S. Wan, B. I. Raju, and M. A. Srinivasan, "Robust deconvolution of high-frequency ultrasound images using higher-order spectral analysis and wavelets," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* **50**, pp. 1286–1295, Oct. 2003.
9. F. Dellaert, "The expectation maximization algorithm," tech. rep., Georgia Institute of Technology, Feb. 2002.
10. K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook*. Technical University of Denmark, 2005. <http://www2.imm.dtu.dk/pubdb/p.php?3274>.
11. L. Sendur and I. W. Selesnick, "Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency," *IEEE Trans. Signal Processing* **50**, pp. 2744–2756, Nov. 2002.
12. J. G. Nagy and D. P. O’Leary, "Restoring images degraded by spatially-variant blur," *SIAM J. Sci. Comput.* **19**, pp. 1063–1082, 1998.
13. J. G. Nagy and D. P. O’Leary, "Fast iterative image restoration with a spatially-varying PSF," in *Advanced Signal Processing Algorithms, Architectures, and Implementations IV*, F. T. Luk, ed., **3162**, pp. 388–399, 1997.
14. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C++: The Art of Scientific Computing*, Cambridge University Press, second ed., 2002.
15. S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1999.
16. N. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Applied and Computational Harmonic Analysis* **10**, pp. 234–253, 2001.
17. N. Kingsbury, "Image processing with complex wavelets," *Phil. Trans. R. Soc. Lond. A* **357**, p. 2543, 1999.
18. L. Sendur and I. W. Selesnick, "Bivariate shrinkage with local variance estimation," *IEEE Signal Processing Lett.* **9**, pp. 438–441, Dec. 2002.
19. J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, 1999.
20. J. A. Jensen, "Speed-accuracy trade-offs in computing spatial impulse responses for simulating medical ultrasound imaging," *Journal of Computational Acoustics* **9**(3), pp. 731–744, 2001.
21. I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series and Products*, Academic Press, fifth ed., 1994.




**Figure 1.** Simulated B-scan images of the human kidney before and after restoration. Each B-scan image has physical dimensions of 70mm (axial) by 40mm (lateral) and, collectively, the images have a dynamic range of 80dB. The lateral focus is at an axial depth of 42mm from the top of each image. The echogenicity image was taken with permission from [www.aic.cuhk.edu.hk/web8/Hi%20res/Kidney%20cross%20section.jpg](http://www.aic.cuhk.edu.hk/web8/Hi%20res/Kidney%20cross%20section.jpg).



**Figure 2.** Simulated B-scan images of the human heart's left ventricle before and after restoration. Each B-scan image has physical dimensions of 30mm (axial) by 40mm (lateral) and, collectively, the images have a dynamic range of 40dB. The lateral focus is at an axial depth of 11mm from the top of each image. The echogenicity image was taken from [www.umdj.edu/pathnweb/syspath/syslab\\_2/Slides\\_24/Slide\\_24\\_A/slide\\_24\\_a.htm](http://www.umdj.edu/pathnweb/syspath/syslab_2/Slides_24/Slide_24_A/slide_24_a.htm).



**Figure 3.** Evolution of ISNR versus program execution time for the various restoration methods. (a) Kidney simulation. (b) Heart simulation.



**Figure 4.** *In vitro* B-scans of a phantom before and after restoration. Each B-scan has physical dimensions of 64mm (axial) by 40mm (lateral) and, collectively, the images have a dynamic range of 80dB. The lateral focus is at an axial depth of 30mm from the top of each image.