

PARAMETER INITIALISATION FOR
THREE-DIMENSIONAL ULTRASOUND
DECONVOLUTION
H. Gomersall, N. G. Kingsbury, R. W. Prager,
H. Shin,
G. M. Treece and A. H. Gee
CUED/F-INFENG/TR 640

Abstract

Conventional beamforming used in pulse-echo ultrasound images results in an inherent blurring. This blurring limits the resolution, and hence the utility, of pulse-echo ultrasound in clinical applications. Ng *et al.* developed an expectation-maximisation algorithm to enhance blurred, two-dimensional RF ultrasound images. It is important that the algorithm is initialised well in order that convergence is rapid. We present here a methodology, with rationale, for estimating the initialisation parameters for the deconvolution algorithm described by Ng *et al.*

1 Introduction

Pulse-echo ultrasound, as a method of imaging anatomical features, has several advantages over other imaging techniques; it is cheap, safe and can capture three-dimensional volumes in real-time. Despite these advantages, its utility is limited because of its poor image resolution, which is caused by limitations of conventional beam-forming, resulting in blurring.

With knowledge of the beam-forming that has been used, it is possible to formulate exactly how the blurring occurred (subject to a few assumptions). Techniques that could recover the ideal field of scatterers given such knowledge of the blurring would go some way to addressing the problems of ultrasound imaging.

Unfortunately, applying a blur to an image results in a loss of information. This means the image can never be exactly reconstructed. It is, however, possible to make an estimate of the original image by having a good statistical model of the system.

Ng *et al.* described a deconvolution algorithm based on a linear model of the blurring inherent in an ultrasound image[1]. The final formulation of the algorithm requires a pair of parameters that should be initialised with some value. Although the parameters are updated as an integral part of the algorithm, by starting with parameters that are close to the correct value, the time to converge is improved over initialising with a poor estimate of the parameters. Such initialisation requires some insight into the nature of the parameters and the physics which dictates their value.

The contribution of this technical report is to attempt to present a methodology with rationale for establishing initialisation parameters for the deconvolution algorithm described by Ng *et al.*

2 The Expectation-Maximisation deconvolution algorithm

Ng *et al.* described how ultrasound can be modelled as a linear shift-variant system [1]. Following this work, an algorithm was developed for wavelet restoration of ultrasound datasets in an expectation-maximisation (EM) framework [2]. This section considers some aspects of the papers by Ng *et al.* The aim is to give the reader context and insight without repeating much of the content of [1] and [2]. For a full exposition of the algorithm and its background, the reader should refer to those papers.

2.1 The linear model

We wish to consider the analytic signal \hat{q}_{RF} of the RF ultrasound time trace q_{RF} (an A-line), which gives us the relatively slowly varying complex envelope. Using the definition of an analytic signal, we have that \hat{q}_{RF} is:

$$\hat{q}_{RF}(\mathbf{r}_0, t) = q_{RF}(\mathbf{r}_0, t) + j\mathcal{H}_t\{q_{RF}(\mathbf{r}_0, t)\} \quad (1)$$

where $\mathcal{H}_t\{\cdot\}$ denotes the Hilbert transform and \mathbf{r}_0 is a vector to a point on the aperture. Effectively, \mathbf{r}_0 encompasses the fact that each time trace can be considered as a positional shift in the active elements of the ultrasound transducer. \mathbf{r} is a vector to some point in space and t is time.

Additionally and separately, we introduce a point spread blurring function, which we call $h_{RF}(\mathbf{r}, t)$, that maps a field of scatterers, $f_m(\mathbf{r})$, onto an RF ultrasound time trace, $q_{RF}(\mathbf{r}_0, t)$:

$$q_{RF}(\mathbf{r}_0, t) = h_{RF}(\mathbf{r}, t) *_{\mathbf{r}} f_m(\mathbf{r}) \Big|_{\mathbf{r}=\mathbf{r}_0} \quad (2)$$

The definition of $h_{RF}(\mathbf{r}, t)$ is:

$$h_{RF}(\mathbf{r}, t) = v_{pe}(t) *_{t} h_{pe}(\mathbf{r}, t) \quad (3)$$

The term $v_{pe}(t)$ encapsulates the impulse response of the ultrasound transducer, and the term $h_{pe}(\mathbf{r}, t)$ describes how a point scatterer contributes to the ultrasound field (the pe subscript denotes that we are considering the *pulse-echo* responses, that is the effect of the full transmit and receive path). The symbol $*$ denotes convolution; in (2) and (3) it is over \mathbf{r} and t respectively.

Drawing on the analysis and assumptions from [1], it can be shown that:

$$h_{RF}(\mathbf{r}, t) + j\mathcal{H}_t\{h_{RF}(\mathbf{r}, t)\} = \tilde{h}(\mathbf{r}, t)e^{j(\omega_0 t - 2k_0 r_3)} \quad (4)$$

The term on the right hand side of equation (4) shows how the Hilbert transform of the signal represents the signal as a slowly varying complex envelope, $\tilde{h}(\mathbf{r}, t)$, and a modulation term, $e^{j(\omega_0 t - 2k_0 r_3)}$, where ω_0 and k_0 are the centre frequency and centre wave number respectively of the axially propagating pulse. Note that equation (4) encapsulates the implicit assumption that the wave is propagating only in the axial direction, which we define as the r_3 axis. This is reasonable when we are considering a conventional electronically and acoustically focussed beam.

The above equations give us an expression for the demodulated analytic signal in terms of the scatterer field and the blurring function [1].

$$\hat{q}_{RF}(\mathbf{r}_0, t) e^{-j\omega_0 t} = \left[\tilde{h}(\mathbf{r}, t) e^{-2jk_0 r_3} \right] *_{\mathbf{r}} f_m(\mathbf{r}) \Big|_{\mathbf{r}=\mathbf{r}_0} \quad (5)$$

To simplify the analysis, it is useful to align the coordinate system with the aperture. As already described, we set the r_3 -axis to be aligned with the axial direction and the r_1 - and r_2 -directions to be aligned with the lateral and elevational directions respectively. This gives us that, for a planar aperture, \mathbf{r}_0 always has a zero r_3 -component. Given these constraints of the system, we can say $\mathbf{r}_0 = (r_1, r_2, 0)$, for some r_1 and r_2 .

The final point to note is that for some arbitrary functions $u(z)$ and $v(z)$,

$$u(-z) * v(z) \Big|_{z=0} = \int u(z) v(z) dz$$

We can now write down a final expression for the time trace at lateral position r_1 and elevational position r_2 (remembering $\mathbf{r} = (r_1, r_2, r_3)$):

$$q(r_1, r_2, t) = \int h(\mathbf{r}, t) \underset{r_1}{*} \underset{r_2}{*} f_m(\mathbf{r}) dr_3 \quad (6)$$

where $q(r_1, r_2, t) = \hat{q}_{RF}(r_1, r_2, 0, t) e^{-j\omega_0 t}$ and $h = \tilde{h}(\mathbf{r}, t) e^{-2jk_0 r_3}$.

Since this is a linear equation, for discretised, real-world signals, this can be formulated in matrix vector notation as

$$\mathbf{y} = H\mathbf{x} + \mathbf{n} \quad (7)$$

where H encompasses the $\tilde{h}(\mathbf{r}, t) e^{-2jk_0 r_3}$ term and its operation on the vector of scatterers, \mathbf{x} , which is the vector form of a discretised $f_m(\mathbf{r})$. \mathbf{y} is the vector form of a discretised $q(r_1, r_2, t)$. Effectively, this is a mapping from the vector of scatterers, \mathbf{x} , onto a vector of time samples \mathbf{y} . Measurement noise is allowed for with the vector \mathbf{n} . We assume \mathbf{n} to be normally distributed with zero mean and covariance $\Sigma_{\mathbf{n}}$.

It is important to realise that, when working on demodulated complex data, it is necessary to take into account the spatial phase de-rotation, $e^{-2jk_0 r_3}$, when defining the blurring matrix H .

2.1.1 Statistical reconstruction

Treating \mathbf{x} and \mathbf{y} as random vectors, and casting the problem into a Bayesian framework, the maximum a posteriori (MAP) estimate may be written as

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} [\ln p(\mathbf{y} | \mathbf{x}) + \ln p(\mathbf{x})] \quad (8)$$

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

$$p(\mathbf{y}|\mathbf{x}) \propto \exp\left(-\frac{1}{2}(\mathbf{y} - H\mathbf{x})^H \Sigma_{\mathbf{n}}^{-1} (\mathbf{y} - H\mathbf{x})\right). \quad (9)$$

If we model \mathbf{x} as a zero-mean Gaussian random vector with diagonal covariance matrix, $C_{\mathbf{x}}$, which has been proposed as a good prior [1], we have:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \left(\frac{1}{2}(\mathbf{y} - H\mathbf{x})^H \Sigma_{\mathbf{n}}^{-1} (\mathbf{y} - H\mathbf{x}) + \frac{1}{2}\mathbf{x}^T C_{\mathbf{x}}^{-1} \mathbf{x} \right) \quad (10)$$

to which the closed form solution is the Wiener filter,

$$\hat{\mathbf{x}} = (H^H \Sigma_{\mathbf{n}}^{-1} H + C_{\mathbf{x}})^{-1} H^H \Sigma_{\mathbf{n}}^{-1} \mathbf{y}. \quad (11)$$

This well-defined and well-understood estimator for \mathbf{x} forms the core of the EM algorithm. It is apparent that if we knew the inverse of the signal covariance matrix $C_{\mathbf{x}}$ and the inverse of the additive noise covariance matrix $\Sigma_{\mathbf{n}}^{-1}$, we would have essentially solved the problem we have set (subject to computational considerations). In loose terms, the formal EM algorithm introduced in [2] solves this by iterating between:

1. Finding $\hat{\mathbf{x}}$, given our best estimate of $C_{\mathbf{x}}$ and $\Sigma_{\mathbf{n}}$

2. Updating the estimates of $C_{\mathbf{x}}$ and $\Sigma_{\mathbf{n}}$.

In this technical report, we are interested in gaining insight into the nature of $\Sigma_{\mathbf{n}}$ and how it might be initialised. Producing a good initial estimate of $\Sigma_{\mathbf{n}}$ is important in the context of producing a rapidly converging EM algorithm. In addition, if the estimates of both parameters, $C_{\mathbf{x}}$ and $\Sigma_{\mathbf{n}}$, are good enough not to require updating, the algorithm becomes significantly simplified and reduces to solving equation (11) just once for these initialised parameters.

3 The Noise Parameter, $\Sigma_{\mathbf{n}}$

In this section, we propose a technique for investigating the noise and methods for initialising the noise parameter, $\Sigma_{\mathbf{n}}$ in a practicable algorithm.

3.1 Some Initial Thoughts

In thinking about how to estimate $\Sigma_{\mathbf{n}}$, it is useful to consider the nature of regularisation in our problem. If we consider the case where the additive noise is stationary, then the covariance matrix reduces $\Sigma_{\mathbf{n}}^2 I$, i.e. just a single scalar noise variance. In this case, equation (11) reduces to the following form:

$$\hat{\mathbf{x}} = (H^H H + \Sigma_{\mathbf{n}}^2 C_{\mathbf{x}})^{-1} H^H \mathbf{y}. \quad (12)$$

The term on the right inside the brackets is a diagonal matrix of the squared noise-to-signal ratio (NSR) at each point in the image. It can be considered as “filling-in” any eigenvalues that are missing from the $H^H H$ matrix, making the term inside the brackets well conditioned and invertible.

The implication of this line of thinking is that it is important to consider the quasi null-space of the H matrix in estimating $\Sigma_{\mathbf{n}}$, given that our model enforces white noise. That is, we care mostly about the noise variance in the space where the signal is small. We say the quasi null-space to imply that space for which the eigenvalues of H are small, rather than necessarily zero.

As we will argue, although the noise variance may not be stationary over the image, there is still a reasonable case to be made that it can be treated as stationary for initialisation.

3.2 The Nature of Noise in an Ultrasound Imaging System

In many data acquisition systems, the dominant noise source is the input amplifier. In such a case, the noise tends to be white with a variance that is a function of the amplifier gain. This is true noise, in that taking a pair of acquisitions of the same source will yield noise that is uncorrelated.

In ultrasound, we find that the dominant source of noise is due some signal dependent effect. It should be noted that we use the term “noise” in a fairly loose manner, to mean any portion of the signal that is not adequately modelled. This may mean that the noise is not true noise in the sense that repeated scans of the same specimen may well yield “noise” that is invariant between scans. This same model is presented by Angelsen[3, p. 1.44][4, p. 11.4], in which he discusses *acoustic noise* as being that noise due to reverberation and sidelobes.

In the remainder of this technical report, we refer to noise to mean any signal that has not been modelled adequately, and so needs to be dealt with as unwanted signal¹.

In the context of ultrasound, we propose the following potential sources of noise, which include electrical noise as well as the signal dependent noise described in [4]:

- An incorrect point spread function, H .
- Non-linear effects, for example caused by second order scatterers. That is, signal that has been reflected off more than one scatterer in the medium of interest.
- Probe interface reverberation. This is in essence a non-linear effect, in which the ultrasound pulse reverberates across the probe head and the specimen interface. It is often present and visible in ultrasound images and so merits an explicit mention distinct from other non-linear effects.
- Reflections from scatterers in the sidelobes of the ultrasound beam. In this case, the problem is that the point spread function is assumed to be truncated to some finite width, and never to radiate outside the lateral and elevational limits of the data set.
- Amplifier noise. This is thermal noise introduced at the amplifier stage and is generally well understood as additive, white, Gaussian noise, with variance set by the amplifier gain.
- Cross channel coupling in which signal from one piezo-transducer couples into the signal wire of another piezo-transducer.

By considering *in vitro* datasets of some known phantom, it is possible to get an understanding of the dominant source of the noise in an ultrasound image of the kind described by equation (7).

3.2.1 A Technique to Estimate the Noise

The blurring matrix H , can be considered to be a convolution in the lateral and elevational directions along with a spatially varying blur in the depth direction, as described in the continuous sense by equation (6). This means that, in the lateral and elevational directions, we can present the action of H as a simple multiplication in the Fourier domain. Consequently, in considering the regions of the Fourier transform in the lateral and elevational directions of our image \mathbf{y} for which the Fourier coefficients in the lateral and elevational directions representing our blurring are zero, or close to zero, we will be considering only the noise in the signal. In other words, the portion of the signal that exists in those Fourier coefficients of our blurring operator that are close to zero must contain primarily noise.

Since we assume that the acquisition procedure is invariant in the lateral and elevational directions, it makes sense to say that by scanning a phantom that is invariant in the lateral and elevational directions, the noise must be invariant in these directions also. We cannot make this same assertion for the depth direction, as the signal is attenuated as it propagates through the medium.

¹It should be noted that this may mean we are throwing away useful information. Such aspects are beyond the scope of this report and could be the subject of future research.

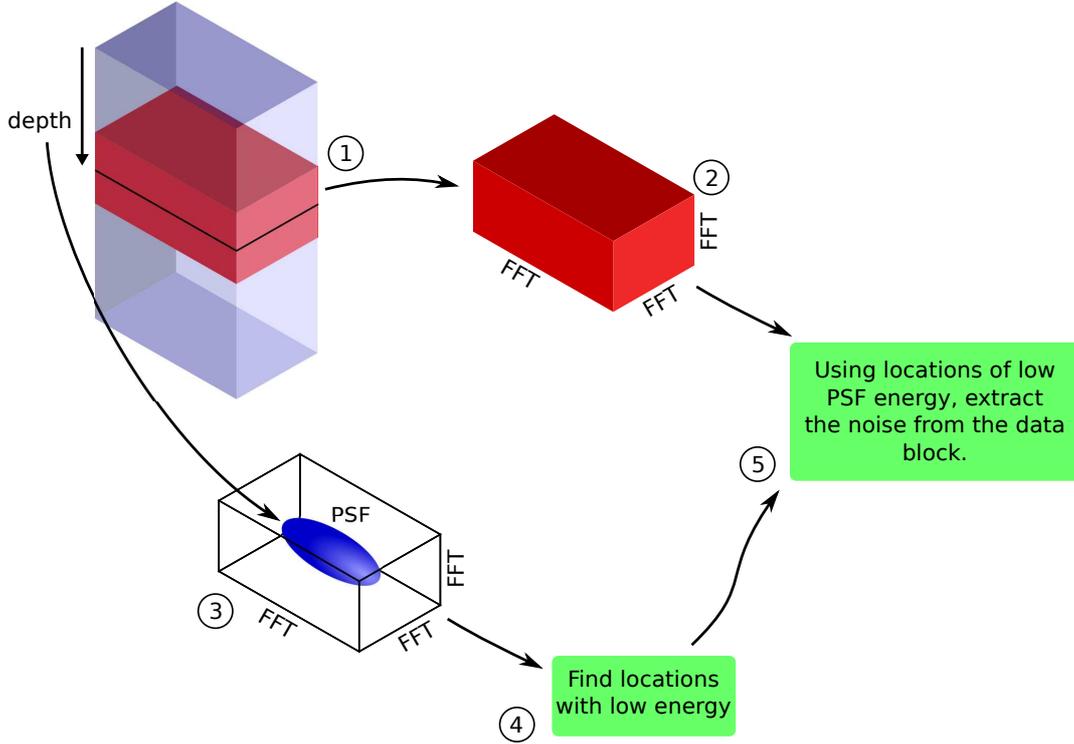


Figure 1: A graphical representation of the noise estimation technique described in section 3.2.1

Following from this, we propose the following method for estimating the noise at some depth:

1. Consider a subblock of the input data \mathbf{y} about the depth of interest. This is a block which spans the full lateral and elevational extent of the data, but is limited in depth. It is assumed that within this small block the blurring in the depth direction is invariant.
2. Take the discrete Fourier transform in all three dimensions of the data subblock.
3. Take the discrete Fourier transform of the point spread function that corresponds to the middle depth of the extracted data subblock.
4. Threshold the discrete Fourier transformed point spread function at some small² value and store the indices corresponding to the zero values.
5. Find the mean square value of the fourier coefficients of the data subblock corresponding to the indices extracted in the previous step³.

This procedure is demonstrated graphically in figure 1.

²The threshold is largely arbitrary. The important point is that the threshold is significantly below the resultant estimated noise value. In practice, it was found to be a simple parameter to set heuristically.

³Finding the mean square value of the discrete Fourier coefficients is equivalent to finding the mean square value of the signal component corresponding to those coefficients, assuming that the Fourier transform has a $\frac{1}{\sqrt{N}}$ factor included, where N is the total number of coefficients.

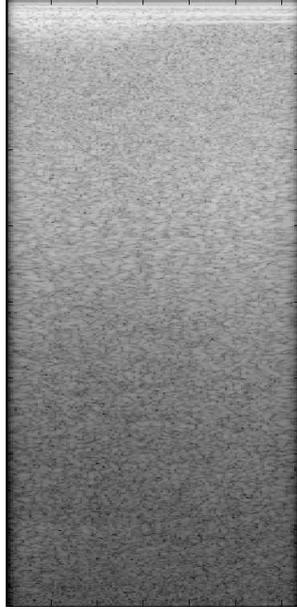


Figure 2: A log-amplitude b-scan of a field of homogeneous scatterers. There is no time-gain compensation applied to this data. Notable features are the bright reverberation artifacts at the top of the image and the brightening of the image about a third of the way down corresponding to the focal depth. The edges of the image are tapered to zero amplitude in an attempt to alleviate edge effects. It also should be noted that this is a single slice through a three-dimensional dataset.

3.2.2 Consideration of the Noise for Various *in vitro* datasets

The first dataset to study is that of a phantom of uniform and homogeneous scatterers. This phantom is an attenuating phantom that attenuates propagating ultrasound at a rate which varies linearly with frequency. It should be noted that in all the datasets considered in this report, the receive amplifier gain is constant as a function of time. That is, there is no time-gain compensation applied. This is so that the receive circuitry noise has constant variance across the whole image. A B-scan showing a field of homogeneous scatterers of the kind of which this first dataset is comprised is shown in figure 2.

If we assume a hypothesis that signal-dependent noise is a simple scaling of the signal intensity, then the expected curve of the noise as a function of depth would be similar in shape to the curve of signal as a function of depth, tending to the noise floor introduced by the receive circuitry. Although this makes a few implicit assumptions about the precise method or methods of noise production (as described in section 3.2), it is found that this hypothesis is reasonable.

Figure 3 shows a plot of noise versus depth and signal versus depth for this homogeneous case. It is apparent that both plots decay with depth, confirming the hypothesis that the noise is signal dependent. Less apparent is why the log noise versus depth plot does not have the same flattening at shallower depths that is apparent in the signal depth. We suggest that an influential factor at these depths is the probe interface reverberation that is clearly visible at the top of the b-scan, figure 2. This reverberation occurs between the probe head and the phantom surface. If the reverberation manifests as itself as wideband noise, with its energy spread across the whole spectrum (in three dimensions), then given our technique for estimating the noise by considering the energy in subbands, we would expect that the reverberations become insignificant with respect to the signal quicker

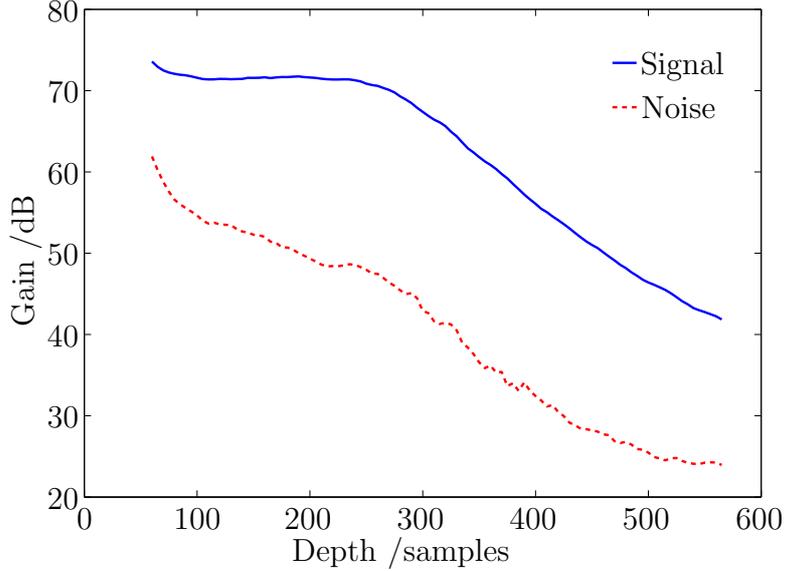


Figure 3: Log signal versus depth and log noise versus depth for a scan of a homogeneous field of scatterers. A constant offset would imply a simple scaling relationship, which we propose is a reasonable model for the data shown. The shape of the noise curve at shallower depths is not the same shape as the signal curve. We tentatively hypothesise that this is due to reverberation from the probe interface and is discussed in section 3.2.2.

than they become insignificant with respect to the noise, since the noise has lower energy than the signal.

The second dataset of interest is that showing the mean noise versus the mean transmit signal. The same setup is used as in the first dataset, but this time the transmit signal amplitude is varied. Figure 4 shows how the mean RMS noise varies with the mean RMS signal. In each of these cases, we take the mean over the depth. It is apparent that the relationship is approximately linear, with a vertical offset. The offset is to be expected and is due to the amplifier noise being invariant to the transmit signal.

The third dataset presents a crude attempt to show that the noise is dependent as much on signal intensity due to changes in the back scatter from the specimen being scanned as on the insonifying amplitude. In this case, we scan two pairs of spheres which are inhomogeneous with respect to the background. B-scans from the two resulting datasets are shown in figure 5. The log RMS noise versus depth and log RMS signal versus depth plots for each of the bright sphere and dark sphere datasets are shown in figure 6. The hypothesis that the noise varies with scattered signal intensity appears to be supported by these plots. A caveat should be attached to this data in that the spectral technique used to estimate the noise implicitly assumes that the data is homogeneous in the lateral and elevational directions (and slowly varying in the depth direction). Since the spheres are distinctly limited in the lateral and elevational directions, some errors should be expected.

The final dataset is used to show that the sidelobes seem not to be a dominant effect in the noise. In this experiment, we take scans of the the phantom with the bright spheres of the previous dataset adjacent to, but not in, the volume being scanned. The volume being scanned consists of the homogeneous scatterers similar to those used in the first dataset. We have one scan with the sphere laterally adjacent, one scan with the sphere

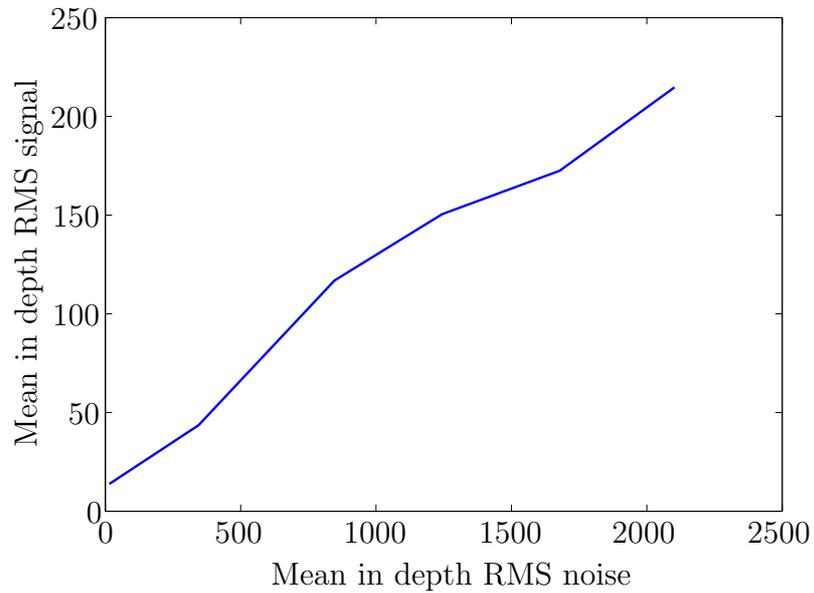


Figure 4: Mean RMS noise versus mean RMS signal, with the means taken in the depth direction. The curve is approximately linear with a y-intercept at the noise floor set by the receive amplifier.

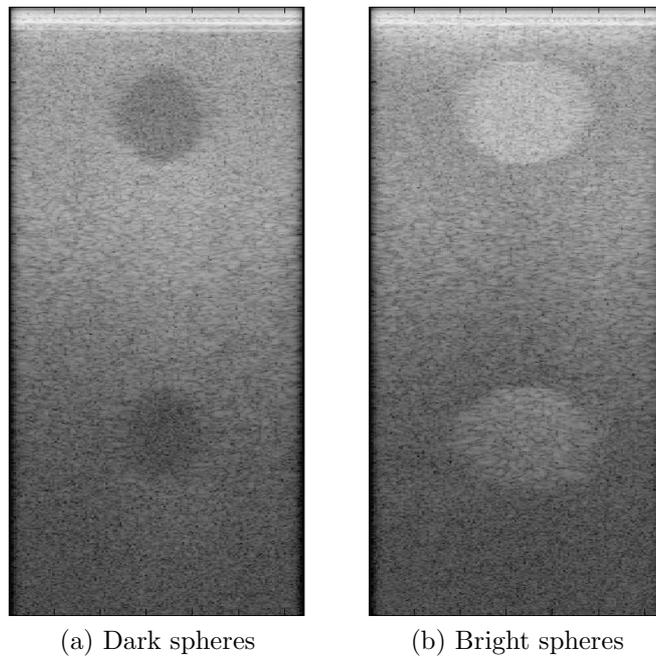
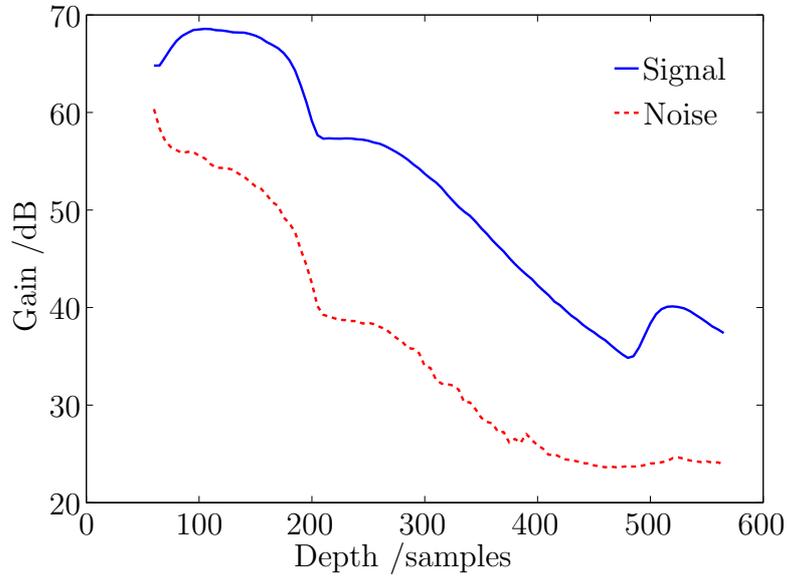
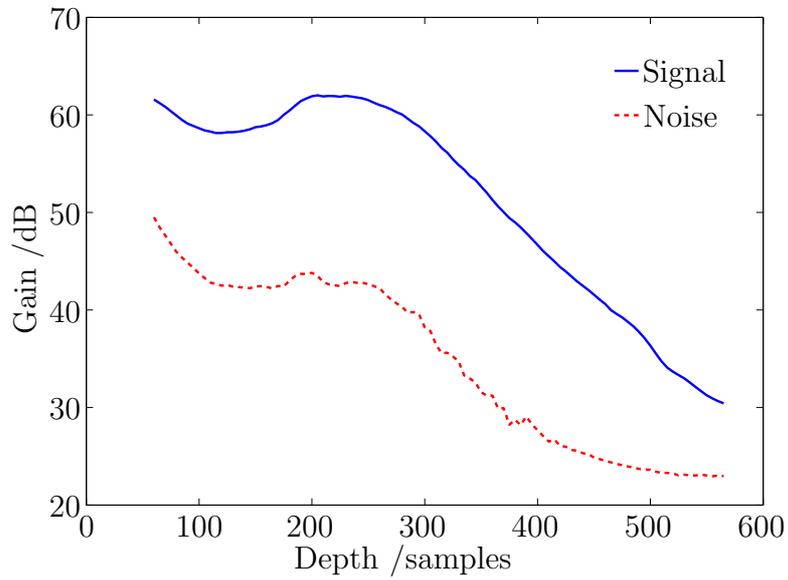


Figure 5: B-scans extracted from the scans of dark spheres and bright spheres.



(a) Bright spheres



(b) Dark spheres

Figure 6: Log signal versus depth and log noise versus depth for two scans of inhomogeneous spheres. It is apparent that the noise varies with the signal backscatterer.

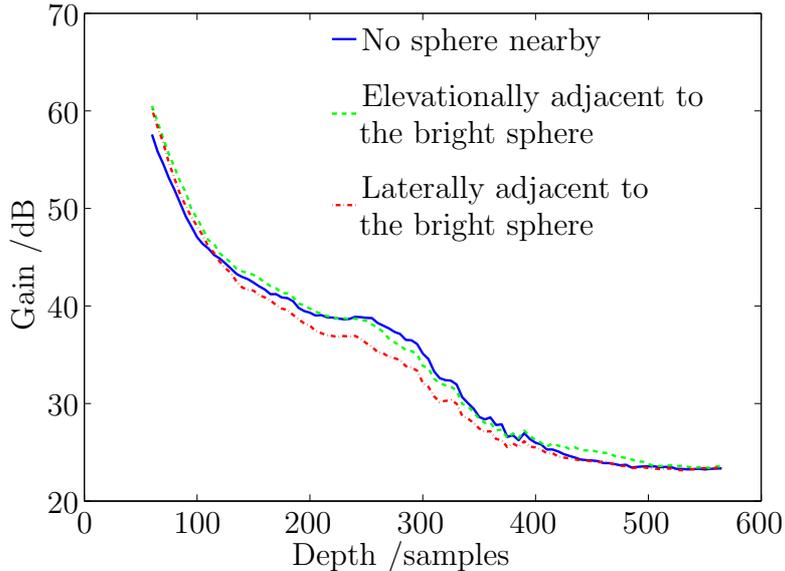


Figure 7: The log noise versus depth plots showing little variation between the case when our scan volume is some distance from any changes in the scatterers compared to when there is a distinct bright sphere adjacent to the scan volume.

elevationally adjacent and one scan for control purposes some distance from any spheres. The noise versus depth plots for all three of these datasets are shown in figure 7. There are no distinct differences between all three and we infer from this that the sidelobes are, at most, only weakly contributing to the noise.

3.3 An Estimate for the Noise Parameter

Based on the studies described in section 3.2.2, we suggest that a reasonable initialisation for the noise can be derived by a simple scaling of the signal covariance matrix, $C_{\mathbf{x}}$. The rationale being that this covariance matrix describes, to a reasonable approximation, the signal amplitude as a function of space. A point to note is that we cannot simply use a scaled version of the signal, as might be expected, because the signal is inherently noise-like, varying significantly from pixel to pixel, such is the nature of speckle. The investigatory techniques used in section 3.2.2 make the assumption that the noise variance is locally stationary and so we require a locally stationary estimate for the signal amplitude, which is exactly what the signal covariance matrix describes. However, in section 4, we will describe how having an explicit estimation for $\Sigma_{\mathbf{n}}$ is unnecessary.

4 Initialising the Noise-to-Signal Ratio in Place of the Covariance Parameter, $C_{\mathbf{x}}$, and the Noise Parameter, $\Sigma_{\mathbf{n}}$

In section 3.3 we suggested that a reasonable initialisation for the noise parameter, $\Sigma_{\mathbf{n}}$, would be a simple scaling of the covariance matrix $C_{\mathbf{x}}$. This gives us:

$$\Sigma_{\mathbf{n}} = \eta C_{\mathbf{x}} \quad (13)$$

where η is the scaling coefficient.

The importance of this is that it gives us a method to initialise our algorithm that relies only on a single parameter that is no longer either the signal or the noise directly, but simply the ratio between them, which is, we assume, unchanging.

Considering equation (11), if we assume that the noise varies slowly, such that it can reasonably be assumed to be constant across the support of the point spread function, then we can make the approximation:

$$\Sigma_{\mathbf{n}}^{-1}H \approx H\Sigma_{\mathbf{n}}^{-1} \quad (14)$$

Substituting (13) and (14) into (11), we simplify our initial estimator for \mathbf{x} to:

$$\hat{\mathbf{x}} \approx (H^H H + \eta)^{-1} H^H \mathbf{y}. \quad (15)$$

η can be found by considering the ratio of the mean square noise to the mean square signal found using the techniques of section 3.2.1, and this becomes the only initialisation value we require.

It is worth noting that equation (15) is simply the zero-order Tikhonov regularisation approximation to the Wiener filter.

References

- [1] J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee, “Modeling ultrasound imaging as a linear shift-variant system,” *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 53, no. 3, pp. 549–563, 2006.
- [2] —, “Wavelet restoration of medical pulse-echo ultrasound images in an EM framework,” *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 54, no. 3, pp. 550–568, 2007.
- [3] B. A. J. Angelsen, *Ultrasound Imaging*. Emantec, Norway, 2000, vol. 1.
- [4] —, *Ultrasound Imaging*. Ementec, Norway, 2000, vol. 2.