CHAPTER 6

SIGNAL PROCESSING TECHNIQUES TO IMPROVE PRECISION OF SPECTRAL FIT ALGORITHM

After developing the Spectral Fit algorithm, many different signal processing techniques were investigated with the goal of improving the precision of the estimation scheme. In this chapter, each of the attempted signal processing techniques is discussed.

6.1 Growing Window Averaging

The first signal processing technique investigated was a simple extension of the Spectral Fit algorithm where the estimates for any given window size were the average of the estimates for smaller window sizes. In the algorithm, the averaging was restricted to window lengths greater than 2 mm. Furthermore, because scatterer size estimates of less than 1 μm were not physically reasonable for the range of frequencies used in the evaluation, any estimates giving a scatterer size smaller than 1 μm were also excluded from the averaging. Because estimates for window lengths have already been found for the basic Spectral Fit algorithm in Chapter 5, the same data were re-evaluated using this new algorithm. As a result, the spectra were still averaged in the log domain and the convolution effects of windowing remained uncompensated. Despite these limitations, the general performance of the algorithm could still be evaluated. Also, the choice of limiting the window length in the averaging to 2 mm was somewhat arbitrary. Hence, other limits may have slightly different performances, but the general behavior should not be drastically affected.

A plot showing the errors in the estimated total attenuation and scatterer size for the half-space with an attenuation of 0.3 dB/cm/MHz for the attempted window lengths is shown in Figure 6.1. Because scatterer estimates of less than 1 μm were automatically excluded from the averaging, not all 40 of the independent waveform groups yielded an estimate as is shown in Figure 6.1c. Also, when comparing the results shown in Figure 6.1 to the previous results given

in Figure 5.1 for window lengths greater than 4 mm, the results in Figure 6.1 have both poorer accuracy and worse precision. At smaller window sizes, the precision was slightly improved due entirely to the discarding of some of the estimates (< 1 μm). Hence, in spite of the averaging of the different windows, the overall performance of the algorithm was degraded. The degradation is even more apparent when comparing the estimates at other values of the half-space attenuation shown in Figure 6.2 to the values given in Figure 5.2.

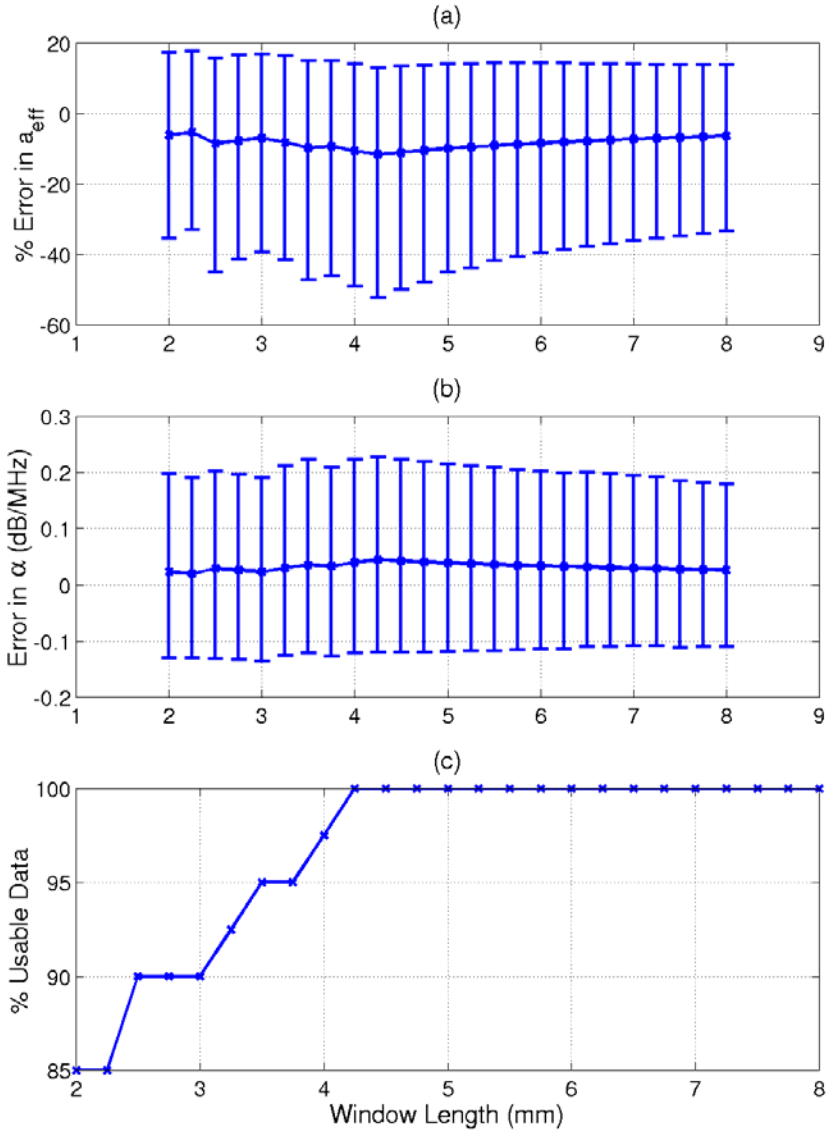

Figure 6.1: Errors in estimates of (a) scatterer size and (b) total attenuation for a half-space with an attenuation of 0.3 dB/cm/MHz using growing window averaging. (c) is the percentage of the data for which an estimate was available.

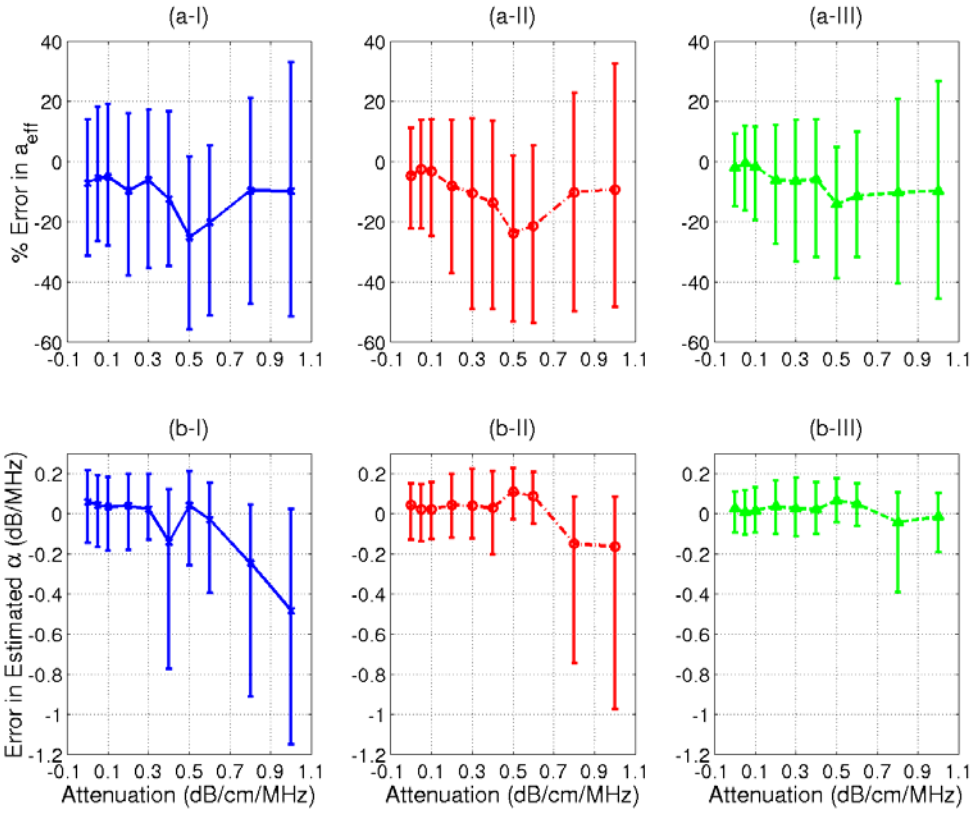

Figure 6.2: Plot of errors in estimated (a) scatterer size and (b) total attenuation for different values of attenuation with hamming window lengths of (I) 2 mm, (II) 4 mm, and (III) 8 mm in the absence of any electronic noise using growing window averaging.

The averaging of the different window lengths was also tested for varying levels of electronic noise. The results showing the noise performance for a half-space attenuation of 0.3 dB/cm/MHz are shown in Figure 6.3. Once again, the accuracy and precision of the algorithm in the presence of electronic noise are comparable to those discussed in Chapter 5. Hence, the averaging of the different window lengths does not improve the noise performance of the Spectral Fit algorithm either.

6.2 Homomorphic Deconvolution

The second signal processing technique investigated involved using homomorphic deconvolution [*Proakis and Manolakis*, 1996] to smooth the backscattered spectrum and hopefully remove the impact of the random scatterer spacing, thereby improving the precision.

Each of the backscattered waveforms was gated using a hamming window $g_{win}(t)$, converted to the frequency domain, and then filtered according to

$$|V_i|^2 = \left|\exp\left(\mathrm{Re}\left(\mathrm{IFFT}\left\{V_{cep_i}\left(\tau_{cep}\right)\cdot g_{win2}\left(\tau_{cep}\right)\right\}\right)\right)\right|^2, \tag{6.1}$$

where

$$V_{cep_i}\left(\tau_{cep}\right) = \mathrm{FFT}\left\{\ln\left(\left|\mathrm{IFFT}\left\{v_{refl_i}(t)\cdot g_{win}(t)\right\}\right|\right)\right\} \tag{6.2}$$

and

$$g_{win2} = \begin{cases} 1 & \left|\tau_{cep}\right| \le T_{cep} \\ 0 & \left|\tau_{cep}\right| > T_{cep} \end{cases}. \tag{6.3}$$

The filtered waveforms were then averaged in the normal spectral domain according to

$$P_{scat}(f) = \frac{1}{N_{lines}}\sum_{i=1}^{N_{lines}}\left(|V_i|^2\right). \tag{6.4}$$

The impact of convolution due to the hamming window was then compensated by the method explained in Section 5.2.2.

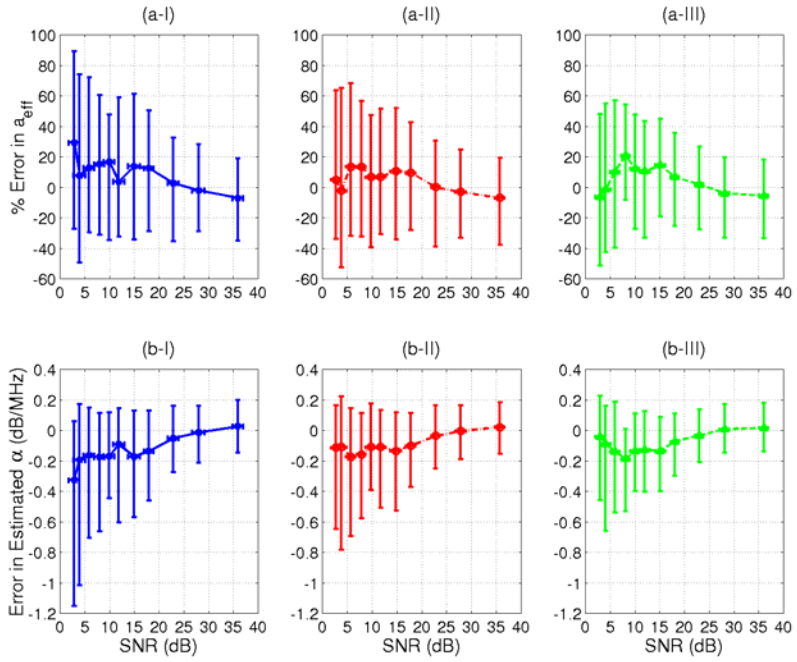

Figure 6.3: Errors in estimates of (a) scatterer size and (b) total attenuation for different levels of electronic noise for a half-space attenuation of 0.3 dB/cm/MHz and hamming window lengths of (I) 2 mm, (II) 4 mm, and (III) 8 mm using growing window averaging.

In order to test the algorithm, simulations were performed of a spherically focused *f*/4 transducer with a focal length of 5 cm sonifying an infinite half-space as was described in Chapter 5. The only difference being that this time 2000 random scatterer distributions were generated. The waveforms from each distribution were then initially grouped in sets of 40 to yield 50 independent estimates of scatterer size and total attenuation at each hamming window length (varied from 1 to 8 mm). The waveforms were initially taken in sets of 40 instead of 25 to attempt to further improve the precision. The relationship between varying the number of waveforms in each set and precision will be discussed later in this section. In addition, the code used a sampling rate of 160 MHz when "digitizing" the signal, and the attenuation of the half-space region was kept at 0.3 dB/cm/MHz. The range of frequencies used by the minimization scheme to perform the estimate was selected so that every frequency f used in the fit satisfied

$$10\log\left(\frac{P_{scat}(f)}{\max\limits_{\forall f}\left(\left\{P_{scat}(f)\right\}_{\substack{Gaussian\\ fit}}\right)}\right) > \max\left(\left[-20 \quad \text{mean}\left(10\log\left(\frac{P_{scat}(f_{N-200}:f_N)}{\max\limits_{\forall f}\left(\left\{P_{scat}(f)\right\}_{\substack{Gaussian\\ fit}}\right)}\right)\right)+6\right]\right) \quad (6.5)$$

(i.e., 6 dB above the noise floor with a maximum set to the –20 dB bandwidth).

In the simulations, the value of T_{cep} used to set the amount of filtering was varied as 0.15 μs, 0.31 μs, 0.46 μs, and 0.62 μs. The locations of all of these cutoffs relative to the main spectral peak are shown with an example signal plotted in the cepstrum domain in Figure 6.4 along with the original unfiltered spectrum. After applying the filter in the cepstrum domain, the example spectra have been significantly smoothed as is shown in Figure 6.5. The smaller the value of T_{cep}, the greater the filtering, and the smoother the resulting spectra. The simulation results for the four different filter lengths are shown in Figure 6.6. For a filter length of 0.15 μs, the accuracy of the scatterer size estimates has been seriously degraded, especially for small window lengths. Hence, this filter length will not be considered in our more detailed analysis of the accuracy and precision given in Figure 6.7. In addition to simulating each of the four filtered cases, the simulation was also run without filtering (i.e., $T_{cep} = \infty$). Hence, the performance of the filtered cases could be compared with the performance available from the straight Spectral Fit algorithm with 40 waveforms in each set averaged in the normal spectral domain while compensating for the convolution effects of windowing.

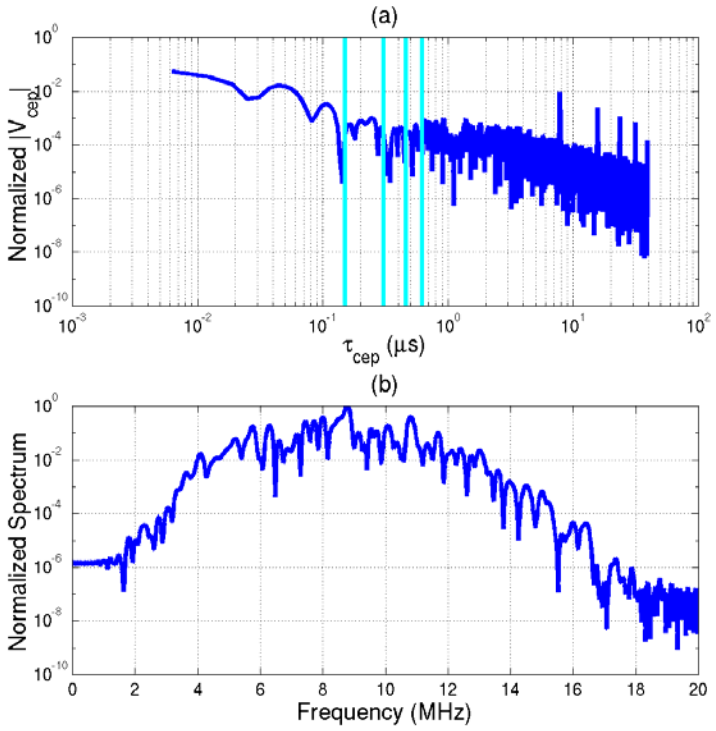

Figure 6.4: (a) Example spectrum for a single RF echo in cepstrum domain showing locations of T_{cep} cutoffs used in the homomorphic deconvolution with (b) the corresponding normalized spectra before any homomorphic processing has been applied.

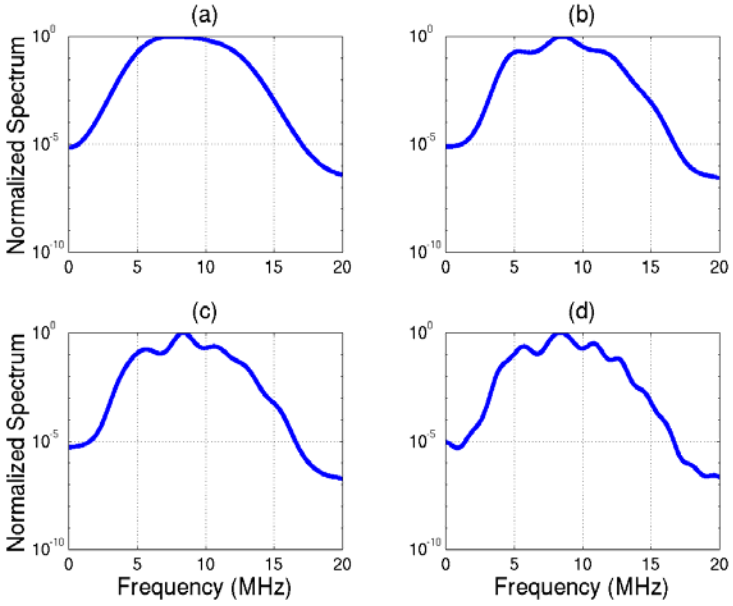

Figure 6.5: Example normalized spectrum for a single RF echo after applying homomorphic deconvolution with (a) T_{cep} = 0.15 μs, (b) T_{cep} = 0.31 μs, (c) T_{cep} = 0.46 μs, and (d) T_{cep} = 0.62 μs.

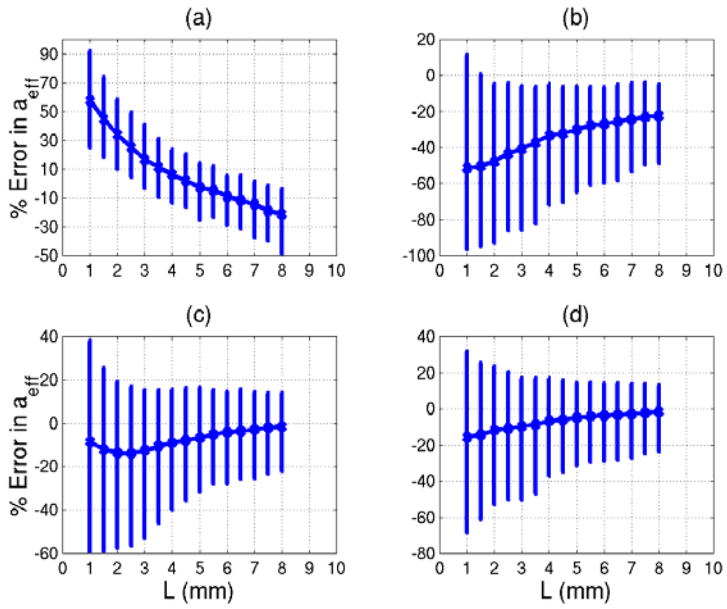

Figure 6.6: % Error in scatterer size estimate for different window lengths after homomorphic deconvolution with (a) T_{cep} = 0.15 μs, (b) T_{cep} = 0.31 μs, (c) T_{cep} = 0.46 μs, and (d) T_{cep} = 0.62 μs.

In Figure 6.7, the accuracy and precision for filter lengths of 0.31 μs, 0.46 μs, and 0.62 μs are compared to the results obtained with no filtering. The accuracy comparison is done by dividing the mean value of the percent error in the scatterer size estimate with filtering by the mean value of the percent error in the scatterer size estimate without filtering (Figure 6.7a). Likewise, the precision comparison is done by dividing the upper (Figure 6.7b) and lower (Figure 6.7c) σ-values as given by Equation (4.18) with filtering by the corresponding values without filtering. Hence an improvement in accuracy or precision would be indicated by a ratio less than one, and degradation in accuracy or precision would be indicated by a ratio greater than one.

From Figures 6.6 and 6.7, it is clear that the accuracy of the scatterer size estimates is degraded by small filter lengths (0.15 μs and 0.31 μs). At larger filter lengths, the accuracy of the scatterer size estimates is slightly improved, especially at larger window lengths (Figure 6.7a). In terms of precision, however, the filtering consistently degrades the performance. The ratio in Figure 6.7b is always greater than one while the ratio in 6.7c is greater than or approximately equal to one. Because it is the precision of the estimates that serves as the limiting factor of the Spectral Fit algorithm, not the accuracy, using homomorphic deconvolution

to smooth each spectrum before averaging provides no benefit to the algorithm when there are 40 waveforms per set.

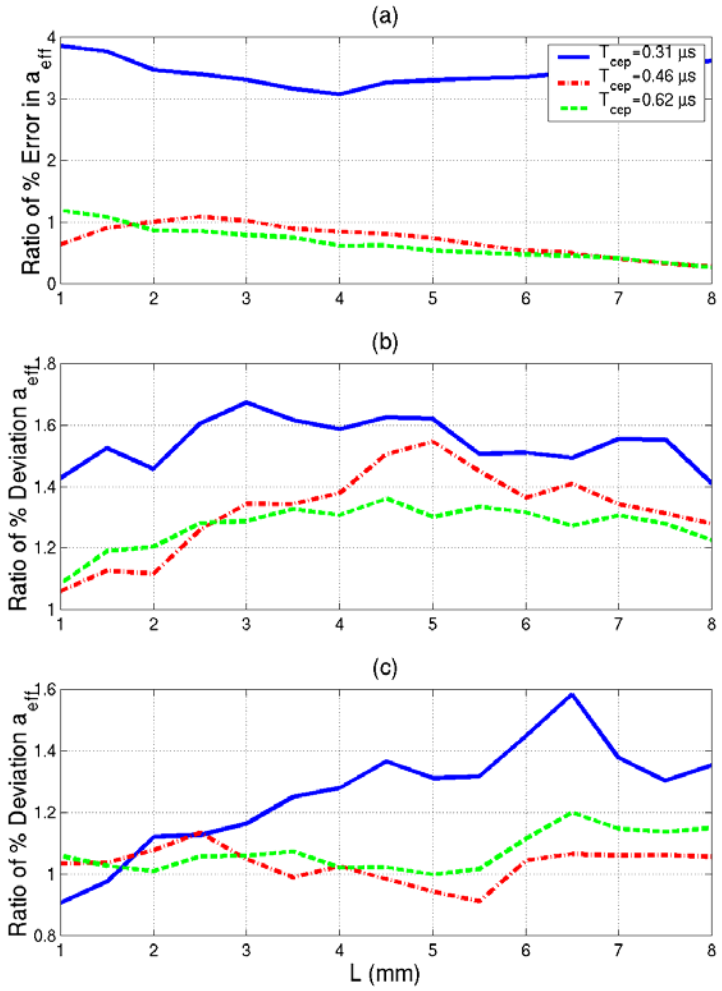

Figure 6.7: Simulation results comparing the performance of different homomorphic filter lengths to the performance of the basic Spectral Fit algorithm by plotting (a) the ratio of the average percentage error of the filtered and unfiltered results, (b) the ratio of $\sigma_{a_{upper}}$ of the filtered and unfiltered results, and (c) the ratio of $\sigma_{a_{lower}}$ of the filtered and unfiltered results.

Before concluding the use of homomorphic deconvolution, two other cases were evaluated. First, it was thought that although the filtering does not provide any benefit when there are 40 waveforms per set, it may provide some improvement when there are fewer waveforms available for averaging in the spectral domain. Hence, the simulations were re-evaluated for a hamming window length of 3 mm with the waveforms grouped into independent

sets of 2, 5, 10, and 25 waveforms. For the sake of comparison, a total of 50 estimates were made for each of the different waveform groupings. Hence, not all 2000 of the previously generated waveforms were needed. Figure 6.8 shows the precision results for each grouping of the waveforms where the upper and lower σ-values from Equation (4.18) are plotted above and below the axis respectively. The first bar in each group corresponds to the case where no homomorphic filtering was applied to the waveforms while the second, third, fourth, and fifth bars correspond to results for filter lengths of 0.15 μs, 0.31 μs, 0.46 μs, and 0.62 μs, respectively.

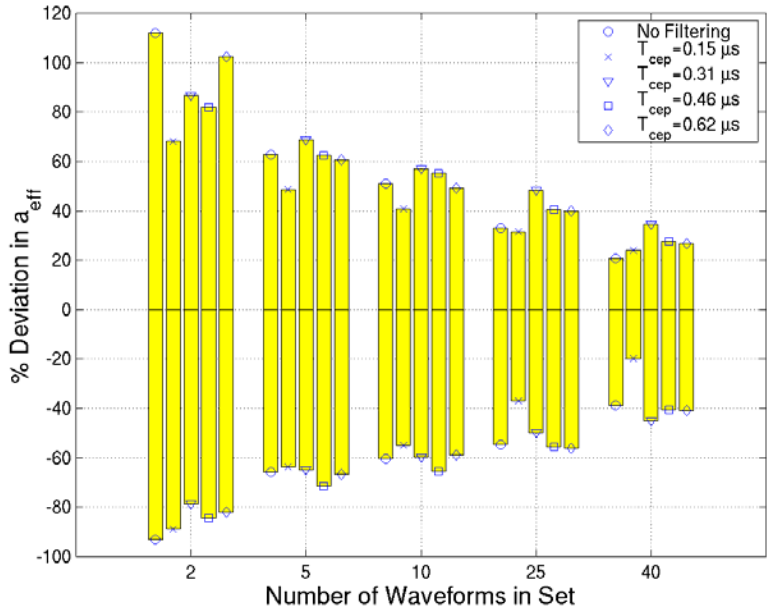

Figure 6.8: $\sigma_{a_{upper}}$ and $\sigma_{a_{lower}}$ values when a different number of waveforms are averaged in each set before perfroming the estimate for the different homomorphic deconvolution parameters.

In Figure 6.8, there is a systematic improvement in precision as the number of independent waveforms used to obtain the estimate is increased. However, for a given number of waveforms in a set, there is no improvement in precision provided by homomorphic deconvolution. The only possible exception occurs when there are only two waveforms per set, but the overall precision for this case is so poor that the amount of improvement is insignificant. Hence, homomorphic deconvolution will not allow for fewer waveforms to be used in the estimation scheme for the same level of precision.

The final case considered before concluding the consideration of homomorphic deconvolution involved applying the filtering to the spectra after averaging the 40 independent waveforms in the normal spectral domain rather than applying the filtering to each waveform

individually. Once again, the value of T_{cep} used to set the amount of filtering was varied as 0.15 μs, 0.31 μs, 0.46 μs, and 0.62 μs, and the simulation was also run without filtering (i.e., $T_{cep} = \infty$). In Figure 6.9, the accuracy and precision for the different filter lengths applied after averaging are compared to the results obtained with no filtering as was done previously. Hence, the accuracy comparison is done by dividing mean value of the percent error in the scatterer size estimate with filtering by mean value of the percent error in the scatterer size estimate without filtering (Figure 6.9a). Likewise, the precision comparison is done by dividing the upper (Figure 6.9b) and lower (Figure 6.9c) σ-values as given by Equation (4.18) with filtering by the corresponding values without filtering. The accuracy of the estimates is still slightly degraded by the filtering for small filter lengths, and the accuracy is still improved at larger filter lengths and window lengths greater than 3 mm. However, the precision of the estimates is relatively unaffected by the filtering. Hence, using homomorphic deconvolution to smooth each spectrum after averaging also provides no real benefit to the algorithm because it is the precision of the estimates that needs to be improved.

6.3 Averaging of Combinations

After considering homomorphic deconvolution, improving the precision by averaging together estimates obtained from different combinations of the available waveforms was investigated. In the basic Spectral Fit algorithm, all the RF echoes (i.e., 25) for a particular set would be averaged in the normal spectral domain and then used to obtain a single estimate for the scatterer size and total attenuation. However, these same RF echoes could also be grouped into different combinations (i.e., 300 by taking in combinations of 23 at a time). Each combination could then be used to obtain an estimate (i.e., 300 estimates) of scatterer size and total attenuation. The estimates from the different combinations could then be averaged to obtain a single final estimate for the scatterer size and total attenuation for the tissue region of interest possibly improving the precision of the estimate.

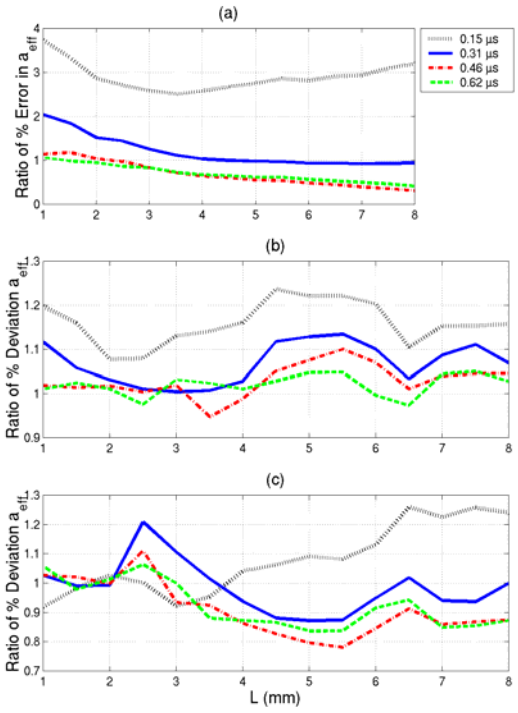

Figure 6.9: Simulation results comparing the performance of different homomorphic filter lengths applied before averaging the waveforms in the spectral domain to the performance of the basic Spectral Fit algorithm by plotting (a) the ratio of the average percentage error of the filtered and unfiltered results, (b) the ratio of $\sigma_{a_{upper}}$ of the filtered and unfiltered results, and (c) the ratio of $\sigma_{a_{lower}}$ of the filtered and unfiltered results.

In order to test whether the averaging of different waveform combinations could improve the precision, other sets of simulations were performed. Once again, the simulations were performed for a spherically focused *f*/4 transducer with a focal length of 5 cm sonifying an infinite half-space containing scatterers with Gaussian impedance distributions that were randomly positioned at a density of 35/mm^3. Initially, the scatterers had an effective radius of 25 µm. The half-space had an attenuation of 0.3 dB/cm/MHz, and the backscattered waveforms from 1000 independent scatterer distributions were generated. The waveforms from each distribution were then grouped in sets of 25 to yield 40 independent sets of waveforms. Within each of the 40 sets, the waveforms were placed in combinations of 23, 24, 25, 2, and 3 at a time yielding 300, 25, 1, 300, and 2300 different combinations, respectively. For each combination, an estimate was obtained. The estimates from the combinations were then averaged to obtain the final estimate. The process was repeated for each of the 40 sets of 25 waveforms, yielding 40

total estimates of scatterer size and total attenuation. A hamming window length of 1 mm was initially used to gate the waveforms, and the convolution effects of the window were compensated as described in Chapter 5. The code used a sampling rate of 53 MHz when "digitizing" the signal, and no electronic noise was added.

The results for estimating the scatterer size and total attenuation for each of the combinations are shown in Figure 6.10 where the error bars are once again given by Equations (4.18) and (4.19). The accuracy and precision is the same when the 25 waveforms are taken in combinations of 25, 24, and 23. This is reasonable because the sets obtained from these combinations would be highly correlated. However, there appears to be some improvement in both the accuracy and the precision of the scatterer size estimate when the 25 waveforms are combined into groups of 2 and 3, with the groups of 2 yielding slightly better performance. Hence, the combinations of 2 will be the focus for the remainder of the evaluation.

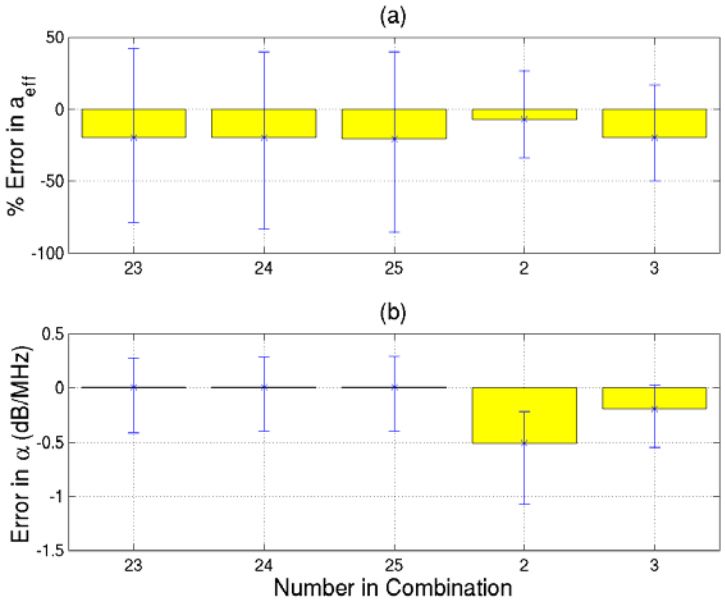

Figure 6.10: Simulation results for (a) scatterer size estimates and (b) total attenuation estimates for using different combinations of the 25 waveforms and averaging all of the estimates from the different combinations to obtain a final estimate.

While performing the simulations for the waveforms taken in combinations of 2, it was noticed that some of the combinations yielded scatterer size and attenuation measurements that were not physically reasonable. Some combinations gave negative values of attenuation while others gave unreasonably small scatterer size estimates (<< 1 μm) compared to the wavelengths used in the simulations. It was then hypothesized that if the unreasonable estimates were

discarded from the 300 combinations before averaging, the overall performance of the estimator would be improved. Hence, the analysis was redone where all estimates giving scatterer sizes less than 1 μm and negative values of attenuation were discarded from the combinations. The case where only the estimates giving negative values of attenuation were discarded, without a concern for estimated scatterer size, was also considered. Both of the results from discarding bad estimates are compared to the case where all of the estimates were kept for 25 waveforms taken 2 at a time in Figure 6.11.

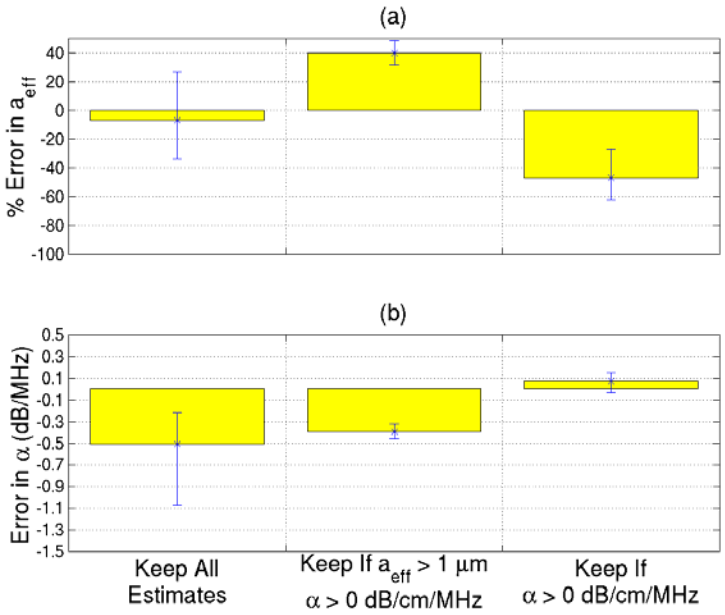

Figure 6.11: Simulation results for (a) scatterer size estimates and (b) total attenuation estimates for taking the 25 waveforms in combinations of 2 and then either keeping all 300 estimates from the combinations or discarding the unreasonable estimates before averaging.

It is clear that the precision of the attenuation and size estimates is dramatically improved when the unreasonable estimates are discarded, but there is also a loss in accuracy for the size estimate. There is a dramatic improvement in accuracy for the attenuation estimate when we only discard estimates yielding negative values of attenuation. In practice, if the total attenuation could be reliably determined, then the scatterer size could be reliably estimated using the traditional estimation scheme discussed in Chapter 2. Hence, there are two cases that warrant further investigation. First, the scatterer size and total attenuation could be found from all 300 of the estimates from the combinations of 2. Second, the total attenuation could be found while only using the estimates from the combinations of 2 that yielded positive attenuation values.

After the attenuation has been estimated, the scatterer size can be determined using a single combination of all 25 waveforms.

In the previous simulations, a hamming window length of 1 mm was used. However, the results of Chapter 5 indicated that longer window lengths tend to have better precision. Hence, the performance for each of the above mentioned two cases was evaluated for hamming window lengths from 1 to 8 mm. The results are given in Figure 6.12 along the with results from the simple Spectral Fit algorithm where the waveforms are taken in groups of 25 (i.e., only a single combination). The deviation in Figure 6.12 is once again found by summing the upper and lower σ-values from Equations (4.18) and (4.19). In Figure 6.2, the accuracy and precision are better for the cases when the 25 waveforms are grouped in combinations of 2 than they are with the basic Spectral Fit algorithm (take all 25 as a single combination), especially at small window lengths. Also, keeping only the estimates that give positive values of attenuation to solve for the total attenuation and then resolving for the scatterer size gives slightly better performance than just averaging all of the estimates from the combinations together.

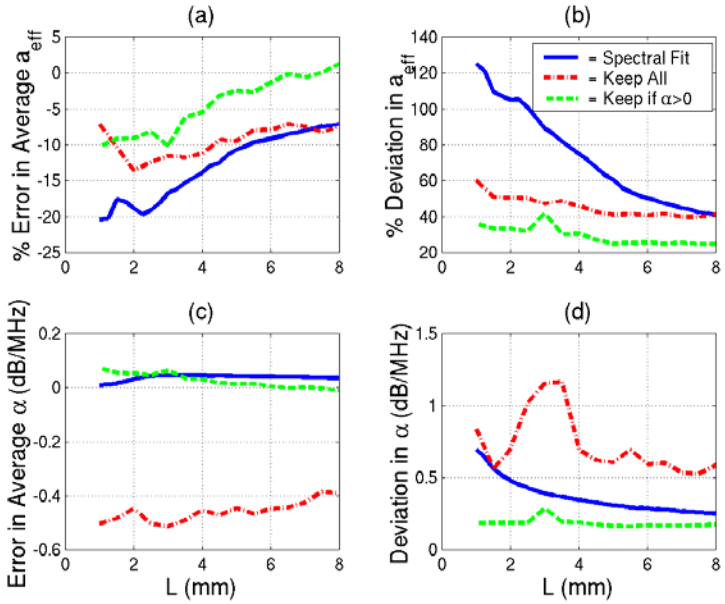

Figure 6.12: A comparison between the basic Spectral Fit algorithm (take all 25 as a single combination) and taking the 25 waveforms in combinations of 2 and then either keeping all 300 estimates from the combinations or discarding the estimates corresponding to α>0 showing (a) the percent error in the average scatterer size, (b) the percent deviation in the scatterer size, (c) the error in the average total attenuation, and (d) the deviation in the total attenuation for different window lengths for an a_{eff} of 25 μm.

Although the results for averaging the combinations have appeared very promising thus far, it is important to demonstrate the applicability of the technique to any appropriate scatterer size as well as any reasonable value of attenuation. Hence, the simulation reported in Figure 6.12 was repeated using a scatterer with an effective radius of 45 μm instead of 25 μm. Once again, the case of estimating the total attenuation by discarding all of the estimates from the combinations of 2 that yield negative attenuation values before averaging the combinations, followed by solving for scatterer size after the total attenuation had been estimated, was considered. In addition, the case where all of the estimates from the combinations of 2 were averaged to get an estimate of scatterer size and total attenuation was evaluated. Both cases are compared to the results for the basic Spectral Fit algorithm (all 25 waveforms taken as a single combination) in Figure 6.13. Both the accuracy and precision for the two different combination cases are worse than the basic Spectral Fit algorithm for a scatterer size of 45 μm. Hence, the results in Figure 6.12 that showed an improvement are a special case that cannot be generalized for all scatterers and attenuations.

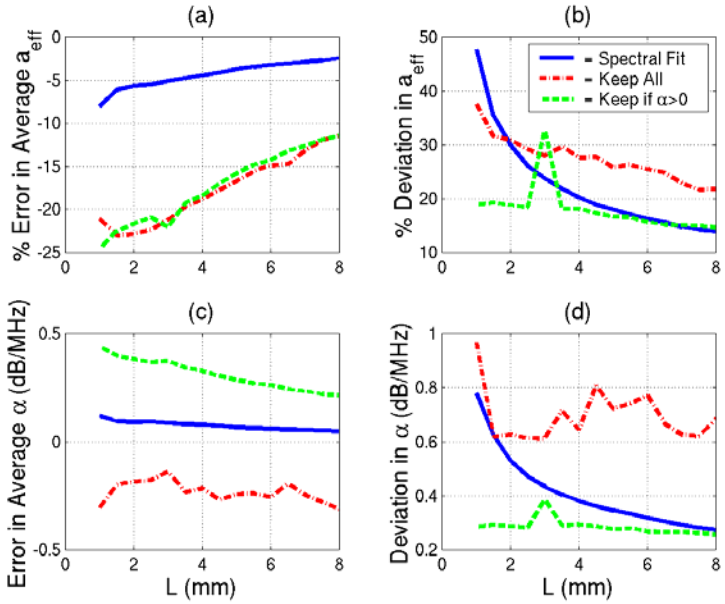

Figure 6.13: A comparison between the basic Spectral Fit algorithm (take all 25 as a single combination) and taking the 25 waveforms in combinations of 2 and then either keeping all 300 estimates from the combinations or discarding the estimates corresponding to α>0 showing (a) the percent error in the average scatterer size, (b) the percent deviation in the scatterer size, (c) the error in the average total attenuation, and (d) the deviation in the total attenuation for different window lengths for an a_{eff} of 45 μm.

In order to understand why averaging of the combinations provided improvement for the 25 μm scatterer size but not for the 45 μm scatterer size, consider the example histograms for each scatterer size of the 300 different combinations of the 25 waveforms shown in Figure 6.14. The histograms were obtained using a hamming window with a length of 3 mm and a half-space attenuation of 0.3 dB/cm/MHz. The number of estimates of scatterer size from the combinations near zero is much greater when the scatterer size is 25 μm (Figure 6.14a) as opposed to 45 μm (Figure 6.14c). Also, when focusing on the main distribution of the estimates (Figure 6.14b and 6.14d), it appears that the estimates are approximately following a Gaussian distribution for the 45 μm scatterers (Figure 6.14d) but a Rayleigh distribution for the 25 μm scatterers (Figure 6.14b). In addition, the variance of the distribution is larger for the smaller scatterer size. For the 25 μm scatterers, the larger number of estimates near zero happen to counterbalance the larger number of estimates greater than the true scatterer size. As a result, by chance the averaging of the combinations gave good size estimates for the 25 μm scatterers when in general the performance would not be enhanced by this technique.

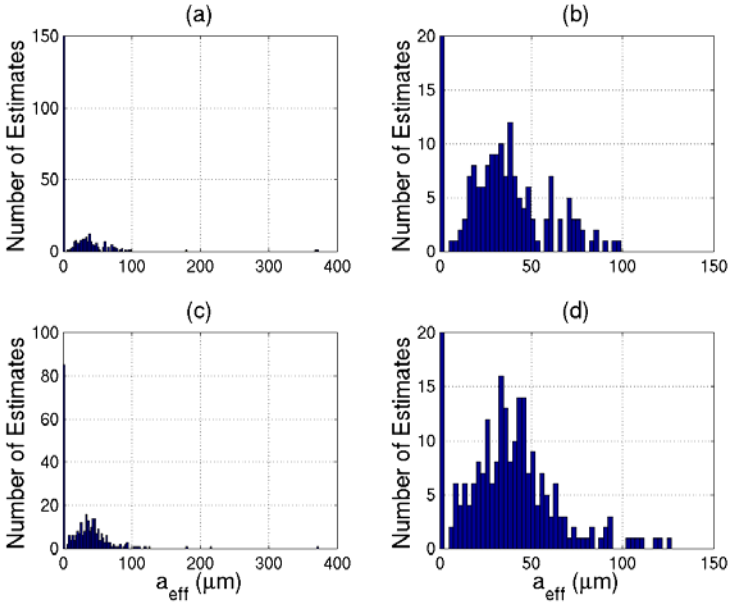

Figure 6.14: Examples of (a) complete histogram and (b) histogram focused on main distribution of scatterer sizes for scatterer size of 25 μm, and (c) complete histogram and (d) histogram focused on main distribution of scatterer sizes for scatterer size of 45 μm for the 300 different combinations of one of the sets of 25 waveforms taken in combinations of 2 at a time.

6.4 Varying Form Factor

The last signal processing technique investigated to improve the precision of the Spectral Fit algorithm involved varying the assumed frequency dependence of the form factor describing the scatterer geometry. The correct form factor when using scatterers with a Gaussian impedance distribution is given by

$$F_\gamma\left(f, a_{eff}\right) = \exp\left(-0.827 \cdot \left(ka_{eff}\right)^n\right), \tag{6.6}$$

where n=2. However, different estimates for scatterer size and total attenuation can be found by varying the value of n and solving the minimization routine of the Spectral Fit algorithm for each value. A final estimate is then obtained by fitting a line to the estimates versus n values and extrapolating the scatterer size for an n value of 2. In the analysis, the n values were weighted by 1 over the minimum value of *ASD* from Equation (5.3). Hence, the estimates corresponding to the best fits from the minimization routine would have a larger impact on the final estimate.

Some simulation results using varying form factors are presented in Figure 6.15 along with the results for the basic spectral fit algorithm. These results were obtained using a half-space with an attenuation of 0.3 dB/cm/MHz containing Gaussian scatterers with effective radii of 25 μm while varying the assumed value of n from 2 to 3.6 in steps of 0.1. Smaller values of n (i.e., less than 2) were also considered, but the results were not qualitatively different from the ones shown. Once again, 1000 waveforms were generated corresponding to 1000 independent scatterer distributions. The waveforms were grouped into sets of 25 (yielding 40 final estimates), windowed with hamming windows, and averaged in the normal spectral domain. The convolution impact of windowing was compensated. Once again, the percent deviation is found by adding the upper and lower σ-values from Equations (4.18) and (4.19). Varying the frequency dependence of the form factor from 2 to 3.6 provides an insignificant amount of improvement to the precision of the estimates (Figure 6.15b and 6.15d) while degrading the accuracy of the estimates (Figure 6.15a and 6.15c).

The lack of improvement can be understood by considering the results for each assumed form factor along with the linear fit for an example spectrum shown in Figure 6.16. A hamming window with a length of 3 mm was used to obtain these results. The curve is smooth. Hence, variations in the assumed frequency dependence of the form factor do not produce statistically independent estimates and as a result will not improve the precision of the estimation scheme.

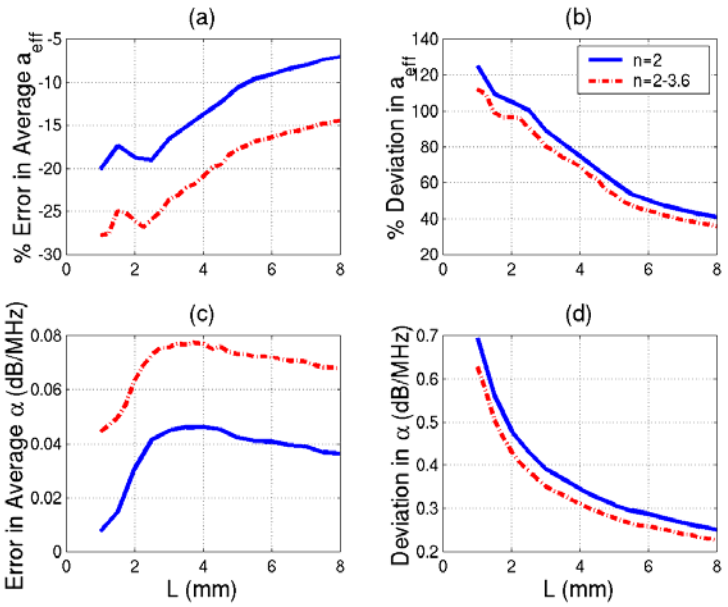

Figure 6.15: A comparison between the basic Spectral Fit algorithm (n=2) and varying the assumed form factor (n=2-3.6) and then extrapolating to find the final estimate for (a) the percent error in the average scatterer size, (b) the percent deviation in the scatterer size, (c) the error in the average total attenuation, and (d) the deviation in the total attenuation for different window lengths.

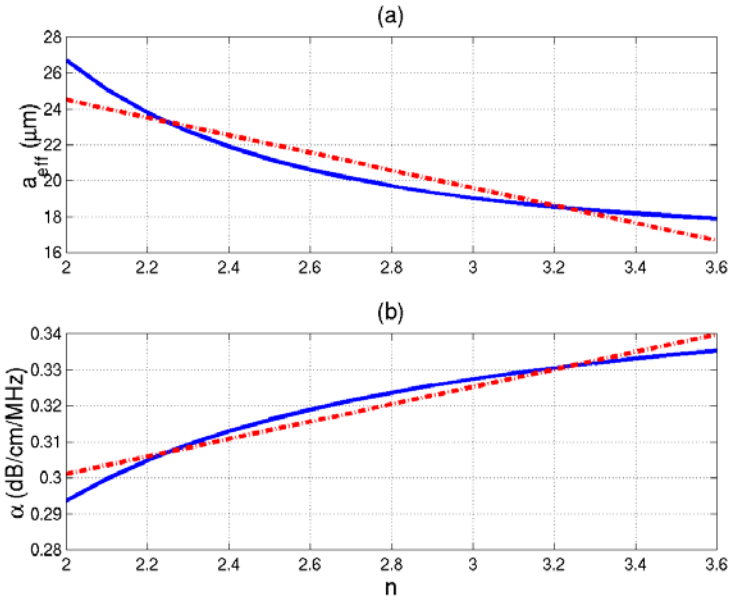

Figure 6.16: Example scatterer size and attenuation estimates at each value of the assumed form factor (solid line) along with the corresponding extrapolation line (dashed line).

The curves shown in Figure 6.16 also demonstrate that if the correct form factor (i.e., scatterer geometry) was not known, reasonable estimates for the total attenuation may still be obtainable. For example, if a frequency dependence of 2.25 is selected (similar to the frequency

dependence of a spherical shell type scatterer over some frequency bandwidths), then the estimated attenuation value will differ by only about 0.01 dB/cm/MHz from the value given for the true frequency dependence of 2. Unfortunately, the scatterer size will still vary with the form factor because the constant term (i.e., 0.827 for Gaussian scatterers) will vary depending upon the scatterer geometry. However, applications such as acoustic radiation force impulse imaging and thermal dose predicting, which only require the attenuation and not the scatterer size, would still benefit from the estimation scheme based on these results even when the correct form factor is not known. This conclusion is important because the correct form factor for biological tissue remains to be determined.

6.5 Chapter Summary

In this chapter, four different signal processing strategies were investigated with the hope of improving the precision of the Spectral Fit algorithm. The first averaged together estimates from different window lengths, but the precision was not improved. The second strategy involved using homomorphic signal processing to smooth the spectra before the fit. Homomorphic signal processing did not improve the precision, nor did it reduce the number of RF echoes required to achieve the same level of precision. The third strategy involved taking the RF echoes in different combinations and then averaging the results together. The strategy worked well for 25 μm scatterers by coincidence, but the algorithm had poor accuracy at other scatterer sizes. The last strategy involved varying the frequency dependence of the form factor, obtaining estimates at each assumed form factor, and then extrapolating to find the estimate at the true form factor. Once again, there was no improvement in the precision due to the correlation between estimates with only slightly different form factors. However, this also means that it may be possible to obtain accurate attenuation estimates when the correct form factor is not known.