Signal Processing and Modelling of Cortical Evoked Potentials for Feature Extraction

by

Alan Ian Gale

Submitted to the Department of Electrical Engineering
and Computer Science in partial fulfillment of the
requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 1996
[February 1997]

Author .....
Electrical Engineering and Computer Science
August 30, 1996

Certified by .....
John L. Wyatt, Jr.
Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by .....
Frederic R. Morgenthaler
Chairman, Departmental Committee on Graduate Students
Electrical Engineering and Computer Science

Signal Processing and Modelling of Cortical Evoked Potentials for Feature Extraction

by
Alan Ian Gale

Submitted to the Department of Electrical Engineering and Computer Science
on August 30, 1996 in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

The research conducted in this thesis is part of an overall effort to develop a retinal prosthetic for restoring vision to subjects with diseases of the photoreceptors. Here we develop signal processing and modelling schemes verifying that the prosthetic device is working as expected, producing vision-type signals in the visual cortex in the form of evoked potentials. Efforts concentrate primarily in two areas: (1) reducing the number of cortical recordings necessary for obtaining accurate estimates of the evoked potential, and (2) modelling the recorded waveforms to further aid our understanding of the evoked response.

Three methods are used to reduce the number of recordings necessary for accurate estimation of the evoked response. These include ideal bandpass filtering designed to remove those frequency components of the noise which do not overlap frequency components of the signal, Weiner-type filtering designed to produce the best mean-square fit between the filter output and the evoked response, and active noise cancellation designed to adaptively subtract a filtered reference noise signal from the cortical recording such that an evoked response estimate is produced. The success of these methods is based on statistical characteristics of mean waveform features, and in the case of active noise cancellation it is also measured in terms of normalized mean-squared error of waveforms. The mean waveforms are formed from randomly chosen individual waveforms, and the features chosen are dependent on characteristics of waveform "size". Ideal bandpass filtering produces a moderate improvement over waveform averaging, although no consistent set of filter parameters is best over all experimental conditions and data sets. Weiner-type filtering consistently produces improvements over averaging and ideal bandpass filtering for all but the smallest amplitude data set. Active noise cancellation produces improvements over the other methods for the largest stimulus levels, but produces mixed results for the other data sets. Active noise cancellation performs consistenly well in terms of normalized mean-squared error between individual waveforms and the evoked potential estimate for all data sets and experimental conditions.

An alternate model is put forth and verified which hypothesizes that the recorded waveforms consist of the sum of a random variable ampltiude scale factor multiplying a deterministic evoked potential and additive noise. This model is verified over a time window of 20-40 ms for all data sets.

Acknowledgments

I am grateful for the opportunity to work on an extremely interesting project which potentially offers so much benefit to a large population of individuals. I would like to thank those agencies who funded this research, Defense Advanced Research Projects Agency (DARPA) and the National Science Foundation (NSF). I would also like to thank my advisor John Wyatt for continued guidance and insight, in addition to providing an environment necessary for the success of such an interdisciplinary project. Additionally I would like to thank the many people whom I work with on a daily basis for their technical expertise and friendship. Finally I would like to thank my parents for their continued support and love without which this thesis would not be possible.

Contents

1. Introduction

1.1 Project Overview

1.1.1 Visual System Description

The mammalian visual system consists of many biological subsystems each performing different biological signal processing functions. After incident light passes through the lens and vitreous of the eye, it then travels through several retinal cellular layers essentially unchanged until it reaches the photoreceptors (rods and cones) in the outermost part of the retina where it is converted to an electrical/chemical signal. From the photoreceptors the signal then is processed biologically in the intermediary retinal cellular layers until it finally reaches the innermost layer of the retina where the ganglion cells are located. From the ganglion cells the signal then travels along the optic nerve, through the optic chiasm and the lateral geniculate nucleus until eventually arriving at the visual cortex of the brain where the signal is measured using surface cortical electrodes.

Numerous causes of blindness such as Macular Degeneration and Retinitis Pigmentosa in mammals are attributed to a loss of photoreceptor cells only, where the other cellular layers remain essentially intact and functional [22]. The retinal implant project uses a silicon based microelectrode implanted on the surface of the retina to directly stimulate the ganglion cells with a current source, circumventing the damaged photoreceptors and the intermediary retinal cellular layers. Since we desire a cortical response which is equivalent to the response due to normal vision, we require the recording of cortical responses to verify that the implant is working as expected. The cortical evoked potentials are very small in amplitude with respect to background biological noise [30]. The most commonly encountered biological noise comes from myogenic electrical activity, EKG, breathing, and background EEG [2]. This unwanted noise must be reduced in order to improve the measurement of the evoked response. This biological noise, and any other noise which is picked up by recording electrodes is generally assumed in the literature to be a stochastic process which is additive to the desired response signal, and the response signal is assumed to be either a deterministic or stochastic process [5, 10, 14, 15, 24]. The importance of which model is used to characterize the signal and noise is discussed below.

1.1.2 Retinal Implant Chip Description

The current version of the retinal implant electrode is a flexible multi-electrode array which is fabricated using integrated circuit technology. It consists of 24 gold electrodes for current stimulation and a single open circuit electrode, all of which are in a square array. The electrodes are connected to contact pads via gold leads. The electrodes, gold leads and contact pads are sandwiched between two layers of polyimide which passivates the device and makes the device flexible so that it can conform to the shape of the retina [19]. Each electrode is driven by a biphasic current source which uses the stimulation waveform described below in Figure 1.1. Note that the positive and negative amplitudes A_a and A_b respectively are related to the positive and negative durations T_a and T_b respectively through the relation $|A_a T_a| = |A_b T_b|$. All stimulation currents are hereafter defined by the magnitude of a single number A_b passing through a single electrode on the array. Recording was done at 1 Hz rate, with the electrode array in monopolar configuration with the rabbit ear as the return electrode. The sampling rate used was 2 KHz.

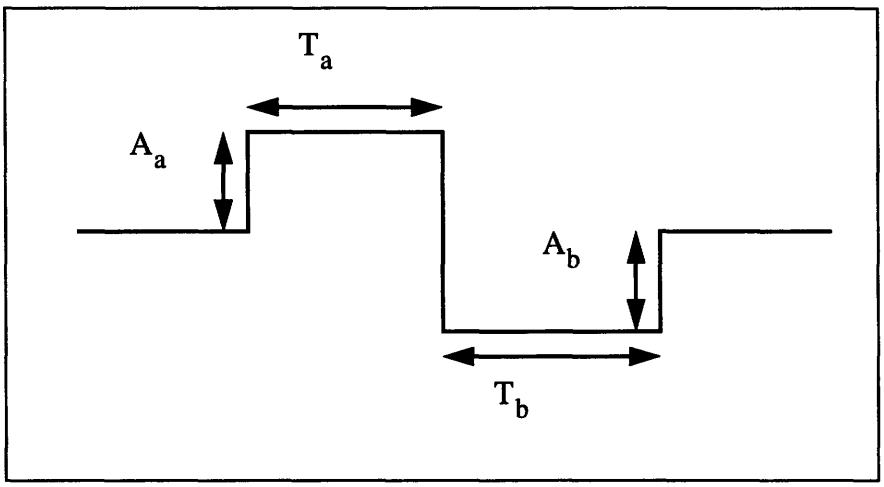

Figure 1.1 Schematic Diagram of Biphasic Stimulation Current

1.2 Evoked Potentials Background

1.2.1 Types of Evoked Potentials

Numerous types of evoked responses are routinely recorded by medical professionals for diagnosing clinical disorders, including electrical evoked potentials (EEP's) which use voltage or current as the stimulus and visual evoked potentials (VEP's) which use light as the stimulus [1, 2, 8]. Although the waveform shape of EEP's and VEP's are very similar, EEP's typically are much smaller in amplitude, and thus the recorded waveforms contain substantially larger amounts of undesirable background noise which must be removed to measure the response. This thesis emphasizes noise removal from EEP recordings only.

1.2.2 Characteristic Waveform Set

The characteristic waveform set used in this thesis consists of four sets of 500 individual cortical recordings each corresponding to different biphasic stimuli, and one set of 500 individual cortical recordings when no stimulus was presented (the noise recordings). The biphasic stimuli per electrode are +8.7/-35, +3.8/-15, +2.5/-10, +2/-7.5 μA, which correspond to A_a/A_b in Figure 1.1 above, and are hereafter referred to as the 35, 15, 10, and 7.5 μA data.

1.3 Thesis Overview

1.3.1 Goals of Thesis

The primary goal of this thesis is to investigate methods for reducing the number of recordings necessary to obtain a reasonably good estimate of the underlying evoked response. The reasonableness of this estimate is based on comparison with the mean waveform from the entire waveform population and its associated features. This mean waveform is assumed to be the best estimate of the evoked response available.

The secondary goal of the thesis is the investigation of an alternative model for the recorded waveforms and the determination of its validity in light of the recordings.

1.3.2 General Definitions

Noise refers to any portion of the recorded signal which is not part of the underlying evoked response. This includes neural activity not associated with the evoked response, other biological activity such as from the heart and lungs, and non-biological noise from the recording process.

Individual waveforms are the single recorded cortical responses due to a stimulation event, which in general include the underlying evoked response and additive noise.

Mean waveforms are those found from averaging the individual waveforms across waveform number, i.e. for each discrete point in time the individual waveform amplitudes are added and the sum is divided by the number of individual waveforms. The term *overall mean waveform* is used when the mean is taken over the entire population of individual waveforms.

Random Subset Mean (RSM) waveforms are mean waveforms where the individual waveforms used to create them are chosen randomly from the total population.

Average is used to describe operations over time for a particular waveform, and *mean* is used to describe operations over waveform number (unless specified otherwise).

1.3.3 LTI Methods for Noise Rejection

Two linear time-invariant (LTI) methods are proposed for reducing the number of recordings necessary in order to obtain an accurate estimate of the underlying evoked response. The first is an approximation to ideal bandpass filtering where individual waveforms are filtered so that the output closely resembles the desired underlying evoked response signal. The method of determining how well the output "resembles" the evoked response is in terms of waveform features, and the reduction of an appropriate error criterion. Bandpass cutoff frequencies are adjusted until this error criterion is minimized. The evoked response is estimated by the overall mean waveform for each stimulus level.

The second LTI method for noise reduction is Weiner-type filtering. Here we use an alternative form of the standard noncausal Weiner filter which is altered to account for an assumed deterministic evoked potential (as opposed to the classic Weiner filter which treats the evoked response as a random process). Implementation of this filter requires the energy spectral density of the evoked response estimate, the power spectral density of the noise estimate and a scale factor based in part on the time duration of the evoked response. The evoked response is again estimated by the overall mean waveform for each stimulus level, and the noise is estimated by independently obtained recordings when no stimulus was present.

1.3.4 Modelling the Recorded Potentials

Two models are proposed to characterize the cortical recordings. The first assumes that the cortical recordings consist of a deterministic evoked response plus an additive random process noise component. It further assumes that this noise is wide-sense stationary, that it is unaffected by the presence of the stimulus and the evoked response, and that the evoked response is not affected by the presence of the noise.

The second model assumes that the cortical recordings consist of a deterministic evoked response which is multiplied by an independent, identically distributed random variable amplitude scale factor, plus an additive random process noise component. Again the noise is assumed to be wide-sense stationary, unaffected by the stimulus and evoked

response, and independent of the amplitude scale factor, and that it does not affect the evoked response. The amplitude scale factor is also assumed to not be affected by the presence of the evoked response.

Verification of the amplitude scale model is facilitated through the use of the data covariance matrix and its corresponding eigenvalues and eigenvectors. A test based on these parameters is constructed and applied to the experimental data over finite duration time windows, allowing verification over individual segments of the waveforms. Using the second model, the amplitude scale factor is then estimated for individual waveforms, and a rough approximation of its distribution is found using a histogram approach.

1.3.5 Active Noise Cancellation

This is an example of a general class of adaptive filtering methods which utilize a searching scheme along an error surface. Typically these methods use iterative techniques where an error signal generated by the filter is fed back as information indicating the quality of the solution at each iteration. The most common searching techniques utilize the Least Mean Squares (LMS) or Recursive Least Squares (RLS) algorithms, which differ (a) in their ability to track nonstationary noise signals, (b) the speed at which they converge to their final solutions, and (c) the difficulty in how they are implemented [28, 29]. Generally the LMS algorithm does not track nonstationary signals as well as the RLS method, but it can be made to have an acceptably fast convergence rate, and it is considerably more easy to implement. For this reason it is the method used in this thesis.

Active noise cancellation attempts to remove the noise associated with the recording process (primary noise) using correlated noise from alternate sources (reference noise) such as from the heart and lungs. The active noise cancellation filter uses both the primary and reference signals as inputs, and attempts to filter the reference signal so that it resembles the primary noise. During each iteration an error signal is generated equal to the difference between the primary signal and the output to the adaptive filter, and this error signal is then used to change the weights on the adaptive filter for the next iteration. This iterative process continues until the energy in the error signal is minimized, at which time the primary noise has presumably been subtracted out from the cortical recordings. Active noise cancellation appears promising given the fact that it achieves noise reduction through subtraction, which is preferable over classical filtering in cases where the frequency content of the primary noise and the evoked response contain considerable overlap.

2. Waveform Attributes

2.1 Overall Mean Waveforms

For each of the four stimulus levels used (35, 15, 10, 7.5 μA) the overall mean waveform is assumed to be the best approximation to the underlying evoked response waveform we desire [13]. It is used as this approximation in the numerous filtering, modelling and estimation sections which follow, so its basic characteristics must be first understood. Shown in Figure 2.1 are the overall mean waveforms for each of the four stimulus levels used, when 500 individual waveforms were used to form the mean (an exception in the case of 35 μA occurred where only 494 waveforms were used, and in the case of 10 μA where 499 waveforms were used, both due to corruption of some individual waveforms due to responses exceeding the minimum voltage measurable). Only the first 200 ms of the 500 ms total duration are shown to emphasize the temporally early portions which are comparatively dominated by the evoked response in relation to the noise.

Examination of the plots in Figure 2.1 shows that there appears to be a distinctive shape of the evoked response estimate for each stimulus level consisting of an early peak followed by a trough and then an second smaller peak, with a tail which tends towards zero. Note also that the relative amplitudes of the peaks and troughs tend to decrease with decreasing stimulus, motivating the use of features which are dependent on some measure of waveform "size".

Figure 2.1 Overall Mean Waveforms for the Characteristic Waveform Set (5-200 ms)

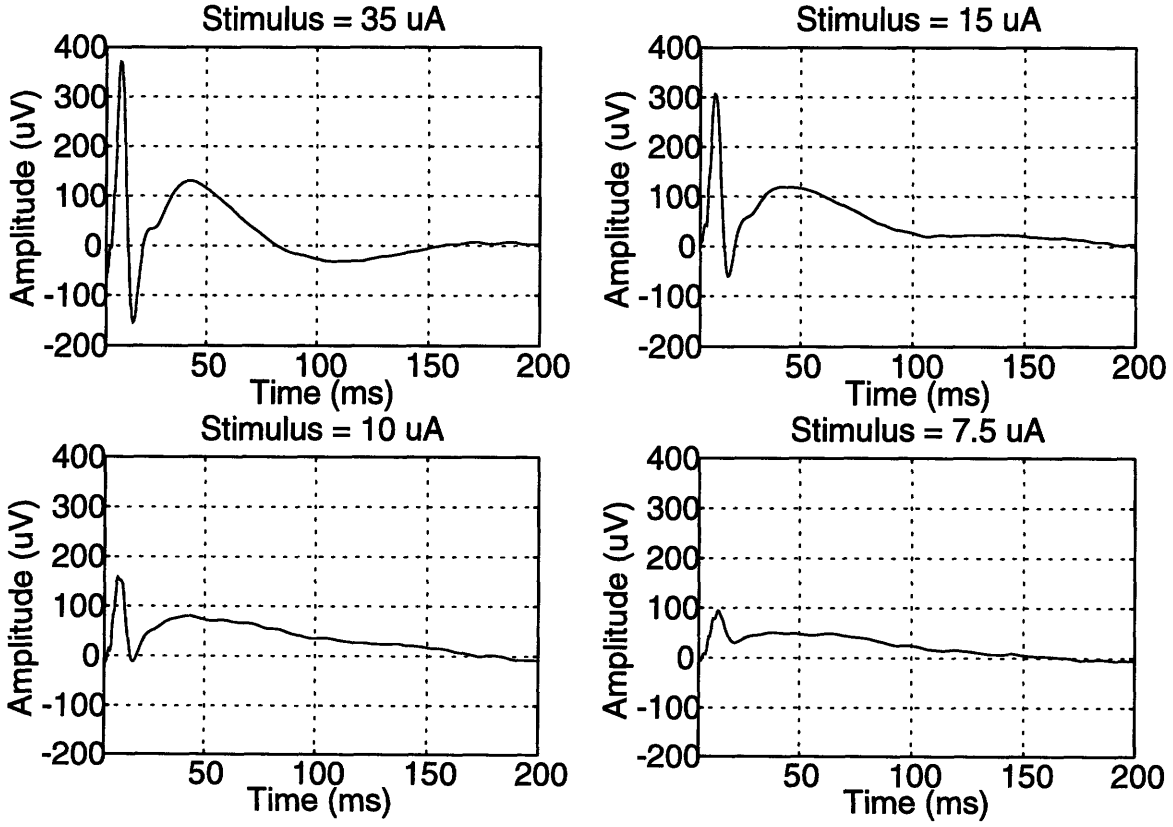

2.2 DC Shifting of Individual Waveforms

For each of the four stimulus levels discussed above, a phenomenon of DC shifting occurs for individual waveforms. This phenomenon is exhibited by individual waveforms having a wide range of time average (DC) values. It is desirable to understand (a) the range of DC values for which this occurs, (b) the distribution of DC values, and (c) any temporal dependencies.

For each of the four stimulus levels, the DC value for each of the individual waveforms was calculated and plotted in a histogram using 25 bins, showing the relative distribution of DC values. Figure 2.2 below summarizes these results. From the figure it is apparent that the range of DC values ranges roughly from ±400 μV for each of the four stimulus levels. Furthermore they appear to have a somewhat Gaussian distribution with a mean of roughly zero. The same analysis was performed on 500 noise waveforms when no stimulus was present, and the same general range and distribution shape of DC values was obtained.

Figure 2.2 DC Value Histograms for the Four Stimulus Levels

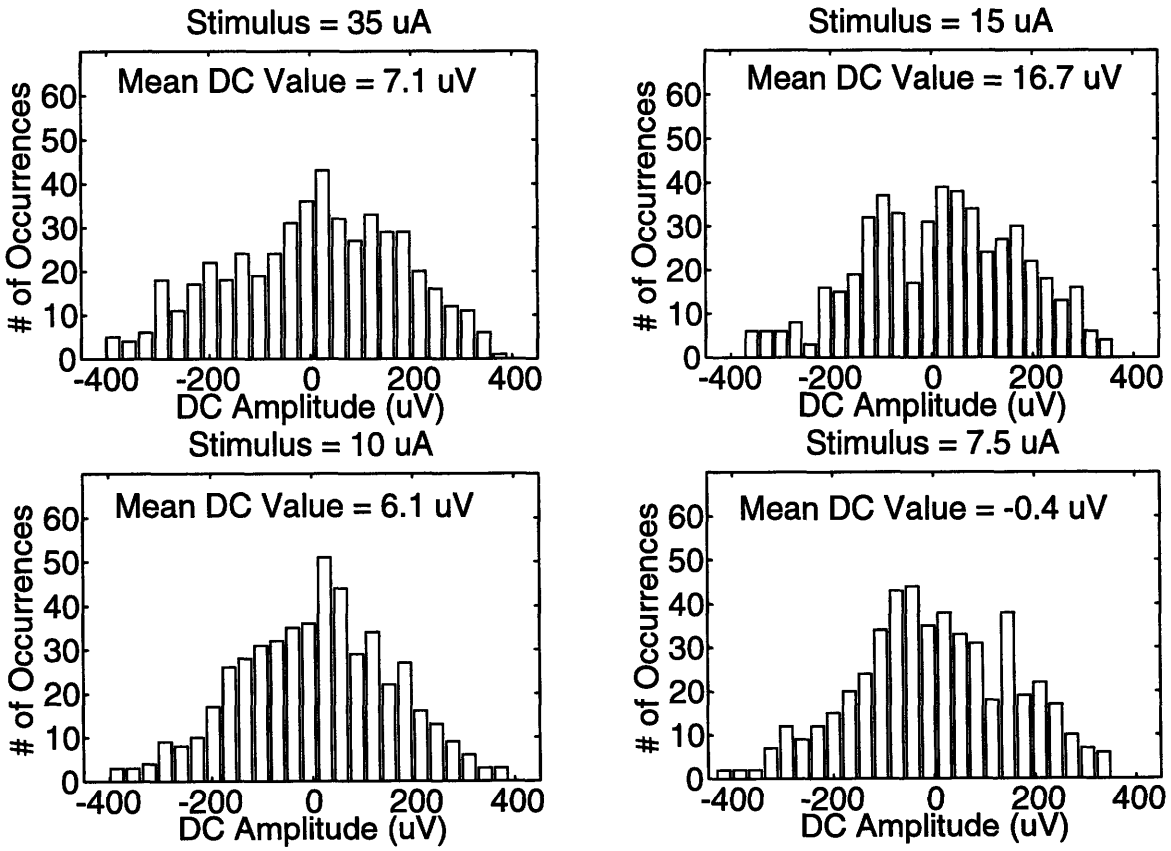

In an effort to determine if any temporal dependence exists, the autocorrelation of DC values was calculated, where any periodicity should show up as periodic peaks at non-zero lag values. These results are summarized in Figure 2.3 below, where a normalized autocorrelation was calculated, the peak at zero lag being unity. Examination of the respective autocorrelation plots reveals no discernible periodicity. Non-periodic temporal dependencies also do not appear to exist, qualitatively determined from examining plots

of DC values as a function of waveform number.

Figure 2.3 Autocorrelation of DC Values (-100 to 100 lags)

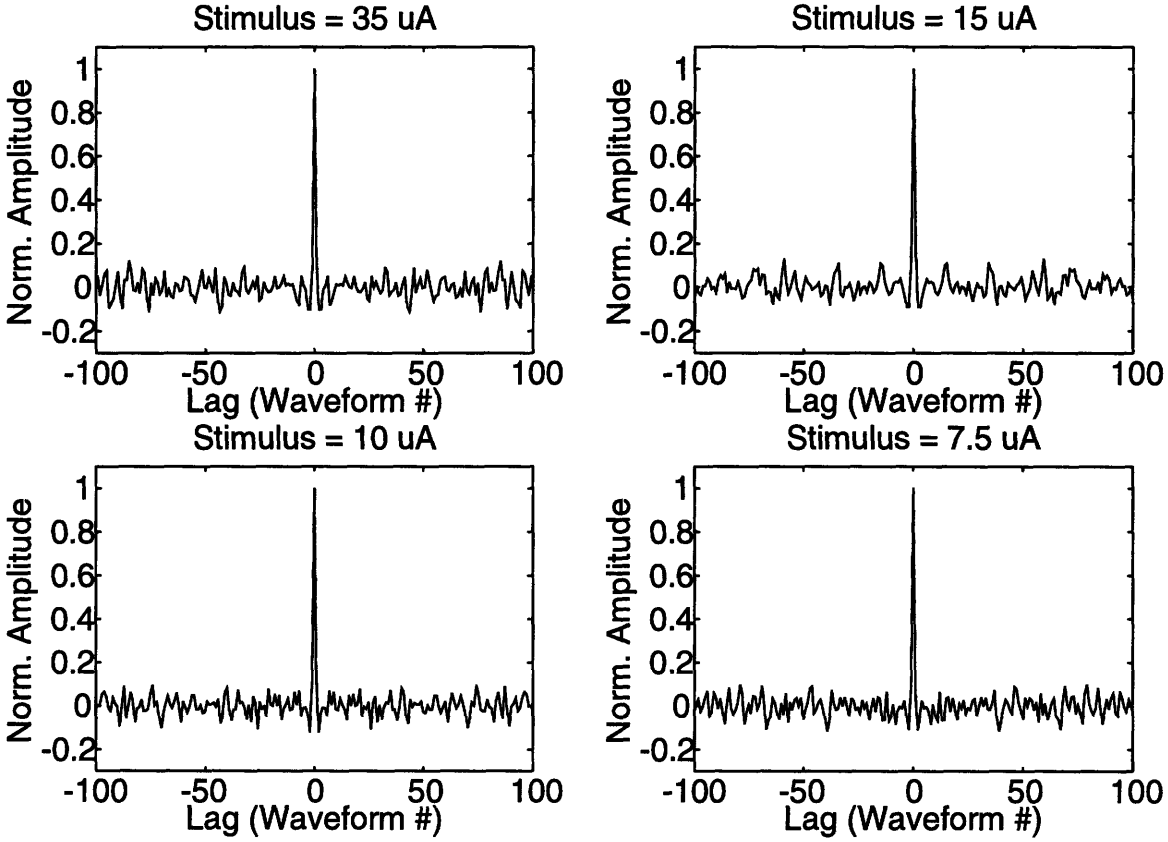

2.3 Frequency Spectrum of Evoked Potential and Noise Estimates

Using the overall mean waveform as our best estimate of the underlying evoked response, we can obtain the frequency spectrum of both the evoked potential estimate and the assumed additive noise. There are at least two methods we can use to obtain an estimate of the noise. In the first, we subtract the overall mean waveform from each of the individual waveforms used to compute it, and assume that the result is purely noise. This assumes that our recorded signal model consists of the deterministic evoked potential plus additive noise, and that we are not using the amplitude scaled deterministic signal model. Since the modelling issues are dealt with in later sections, we do not use this method here to estimate the noise. The second method for estimating noise uses independently obtained recordings when no stimulus is present, and is the method used here.

The power spectral density of the noise is estimated by the Bartlett method which suffers from poor frequency resolution, but which has lower estimator variance than most other nonparametric methods [12, 21]. This asymptotically unbiased estimator is found by finding the magnitude squared of the Fourier transform of each individual waveform divided by the length of the waveform, and then averaging these terms to yield the power spectral density estimate. This method and others are discussed in more detail in the Weiner-type filtering section elsewhere.The evoked response energy spectral density is estimated by the magnitude squared of the 4096 point FFT of the overall mean waveform.

Figure 2.4 below summarizes these results. From Figure 2.4 it is apparent that the majority of both the evoked potential and noise estimated signal components lie in an overlapping low frequency range of less than 10 Hz.

Figure 2.4a Energy Spectral Density Estimates of Evoked Responses (0-100 Hz)

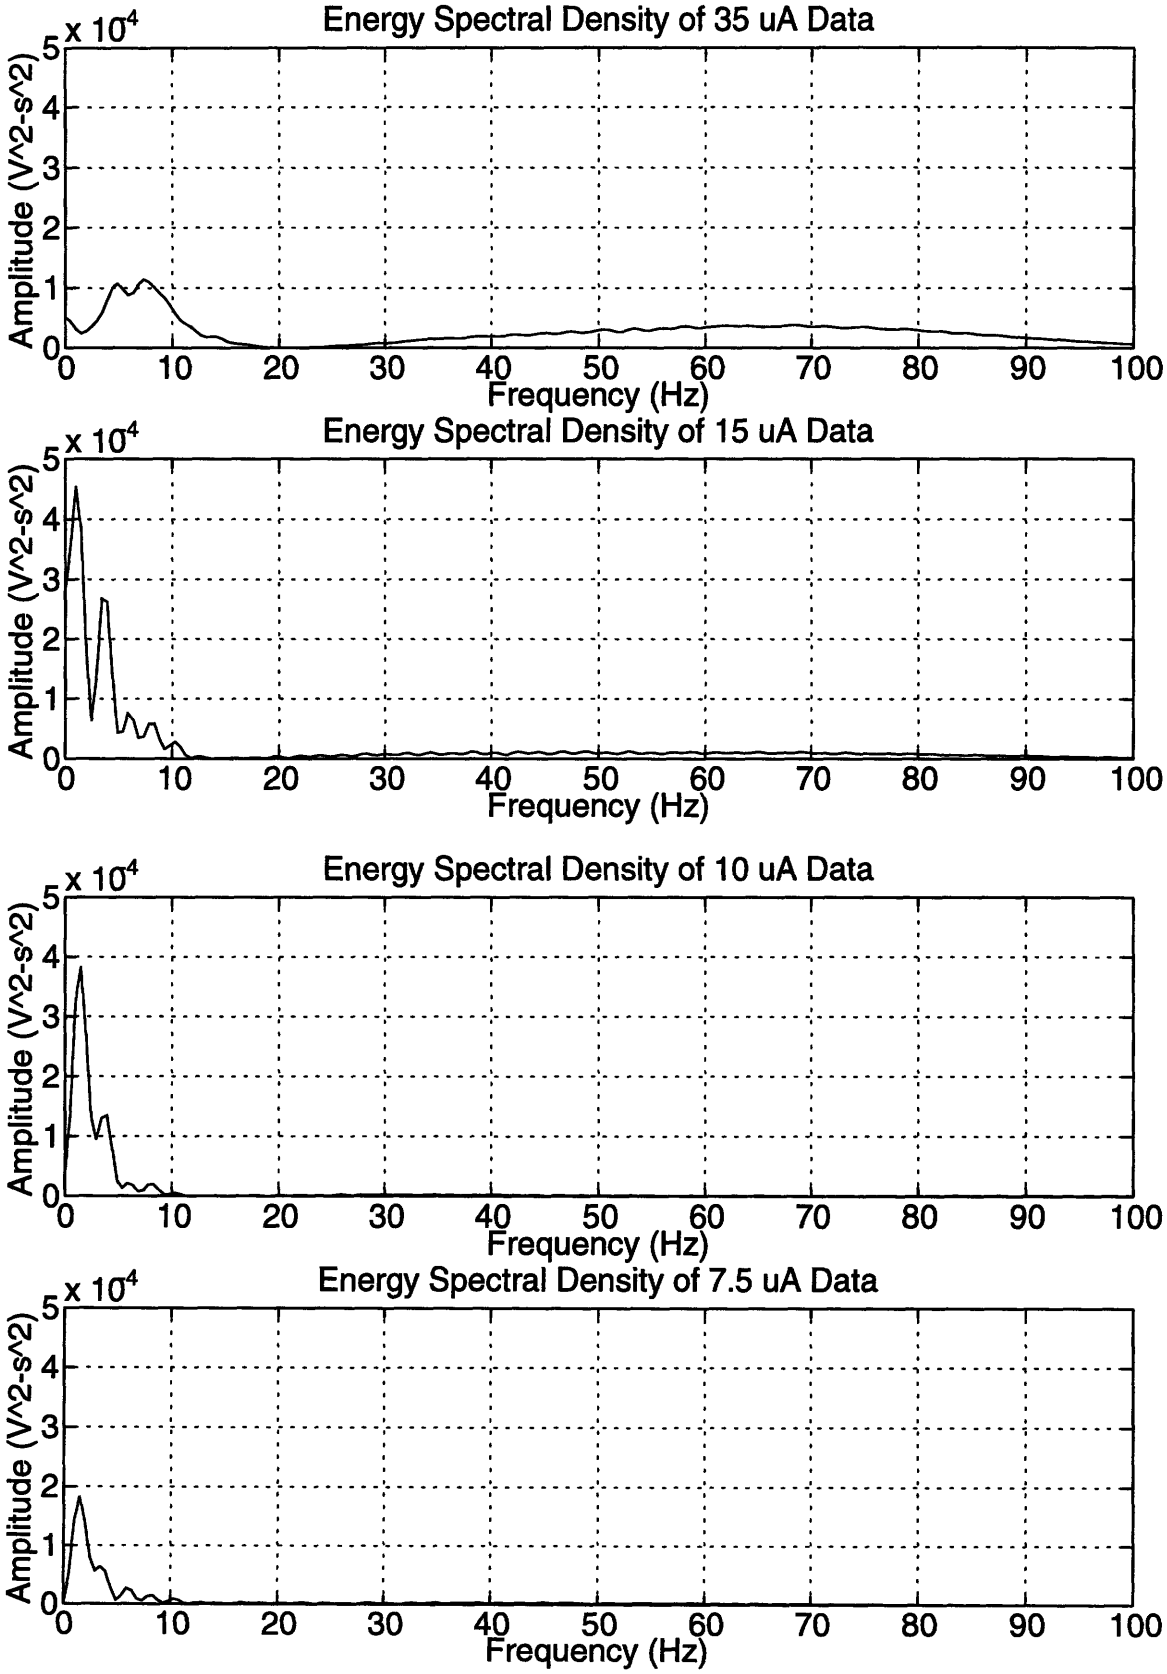

Figure 2.4b Power Spectral Density Estimate of the Noise (0-100 Hz)

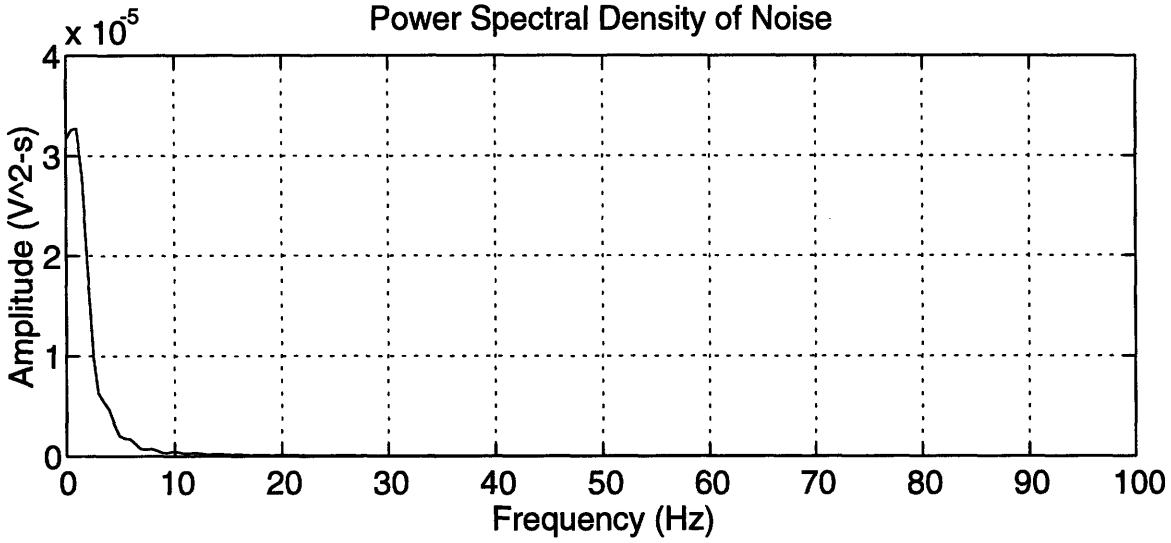

3. Feature Extraction and Optimization of Parameters

3.1 Choice of Features

Given the characteristic shape of the overall mean waveforms for each of the four stimulus levels whose "size" appears to be positively correlated to stimulus level, and the presence of DC shifting of individual waveforms, we seek some features whose measure of "size" are independent of the DC levels. Furthermore, since only the earliest portions of the overall mean waveforms appear to contain the desired information from the evoked response, we seek some features which only examine the earliest portion. Finally, given the substantial overlap in frequency components of the evoked potential and the noise, and the similarity in evoked response spectra for different stimulus levels, our features should not depend on differences in the frequency spectra. Two features which satisfy all these criteria are the peak-to-trough amplitude and the root-mean-square (RMS) amplitude about the time average. Both of these features seek some measure of size in a finite length window, and are both independent of DC shifting of individual waveforms.

3.1.1 Peak-to-Trough Amplitude Feature

The peak-to-trough amplitude over the window $\{W_l, W_u\}$ of an RSM waveform $r_k[n]$ is defined as follows:

$$P_k = \max\{r_k[n]\} - \min\{r_k[n]\}\Big|_{\{W_l, W_u\}}$$

3.1.2 Root-Mean-Squared Amplitude about the Time Average Feature

The root-mean-squared (RMS) amplitude about the time average ($\bar{r_k}$) over the window $\{W_l, W_u\}$ of an RSM waveform is defined as follows:

$$F_k = \sqrt{\frac{1}{W_u - W_l + 1} \sum_{n=W_l}^{W_u} \left(r_k[n] - \bar{r_k}\right)^2}$$

where

$$\bar{r_k} = \frac{1}{(W_u - W_l + 1)} \sum_{n=W_l}^{W_u} r_k[n]$$

3.2 Statistical Procedures for Error Criterion

Although we wish to ultimately compare the peak-to-trough and RMS amplitude features of the RSM waveforms to those of the overall mean waveform, we would also like to compare the results of the present work to the analysis of recordings made in the future, when presumably fewer individual recordings are obtained. We thus desire an error criterion which is somewhat independent of the number of recordings made, but which presents results which are comparable across different experiments.

One method of comparing the features between RSM waveforms and the overall mean waveform is using the mean squared error, which is the empirical average between

the RSM waveform features and the feature obtained from the overall mean waveform. An alternative error criterion, the sample variance, does not explicitly depend on the overall mean waveform (in the sense that we do not need to calculate the overall mean waveform to find it), and thus yield results which allow the results of the present work to be extrapolated to the analysis of future recordings. The similarity between sample variance S^2 and mean squared error ε^2 for the RMS feature can be seen by examining their defining formulas:

$$S^2 = \frac{1}{N_w - 1} \sum_{k=1}^{N_w} \left(F_k - \bar{F} \right)^2 \qquad \varepsilon^2 = \frac{1}{N_w} \sum_{k=1}^{N_w} (F_k - F_o)^2$$

where

N_w = number of random subset mean waveforms

F_k = root-mean-square (RMS) amplitude of waveform k about its time average, $k = 1, 2, \ldots, N_w$

$\bar{F}$ = mean of the F_k taken over k

F_o = RMS amplitude about the time average of the overall mean waveform

From the above it is apparent that the sample variance and mean square error are very similar in the limit as N_w becomes large. In this case the factor of $N_w - 1$ in the sample variance formula (which makes this an unbiased estimator of the true variance [21]) is approximately equal to the factor of N_w in the mean squared error formula. Furthermore, $\bar{F} \approx F_o$ as N_w becomes large since the average feature from the RSM waveforms approaches the feature of the overall mean waveform under the assumption that the RSM waveforms were found from a sufficiently large number of individual waveforms.

However, using the sample variance as the error criterion may produce some erroneous results. It is possible that it can be reduced to zero arbitrarily, as would be the case if the RSM waveforms were passed through a no-pass filter. To circumvent this problem, and to give error results which are more easily interpreted, our error criterion is the ratio $2\sigma/\mu$, where σ is the sample standard deviation (the square root of the sample variance of the features), and μ is the mean of the features. This error criterion attempts to give the range of possible feature values where a large percentage of the data will fall under. For example, if the feature distribution is Gaussian, then $\pm 2\sigma/\mu$ represents the range where 95% of the feature values will fall [7]. For the above reasons, this is our chosen error criterion, which will hereafter be referred to as the percent spread about the mean (when it is multiplied by 100%).

3.3 Optimization of Parameters

The calculation of both features for RSM waveforms requires optimization of four parameters for each feature. First we must determine the number of individual waveforms N_i used to form each of the RSM waveforms. We seek the smallest number of individual waveforms such that enough noise can be averaged out so that the measured features are not primarily dependent on the additive noise, but not so large as to subvert our desire to use the fewest number of individual recordings. We must also determine the number of

RSM waveforms N_w, where we seek the fewest number such that our sample variance estimate becomes roughly constant for larger values of N_w. This number is merely the minimum number necessary to ensure that our measure of sample variance is not strongly dependent on which waveform set is arbitrarily chosen. Although in general a larger value of N_w helps ensure accurate statistical measurements, in practice the computational time for large N_w becomes excessive. Finally, we must determine the appropriate lower and upper window times, W_l and W_u for feature extraction. This window time reflects our interest in only a finite portion of the evoked response.

The following optimization process is initially applied to the 35 μV data only, and the results are then verified for the other data sets.

3.3.1 Determination of the Number of Individual Waveforms Used to Form RSM Waveforms and the Number of RSM Waveforms

These two variables are found by applying the feature extraction processes for both features to RSM waveforms when no bandpass filtering is performed, and when the largest window is used (from 5 to 50 ms). This window was chosen because the onset of the evoked potential starts at 5 ms, and because the clinically significant portion of the evoked potential ends at about 50 ms [20]. The number of individual waveforms used to form each RSM waveform was varied between 20 and 200 using values of 20, 50, 100, and 200, while the number of RSM waveforms was varied between 20 and 60 in increments of 10. No bandpass filtering was performed in order to better understand the effect of changing N_i and N_w. This was the same motivation for using the widest window when performing feature extraction. The results are summarized in Figures 3.1 and 3.2.

Figure 3.1 Mean Peak-to-Trough Feature and Percent Spread Versus N_w and N_i

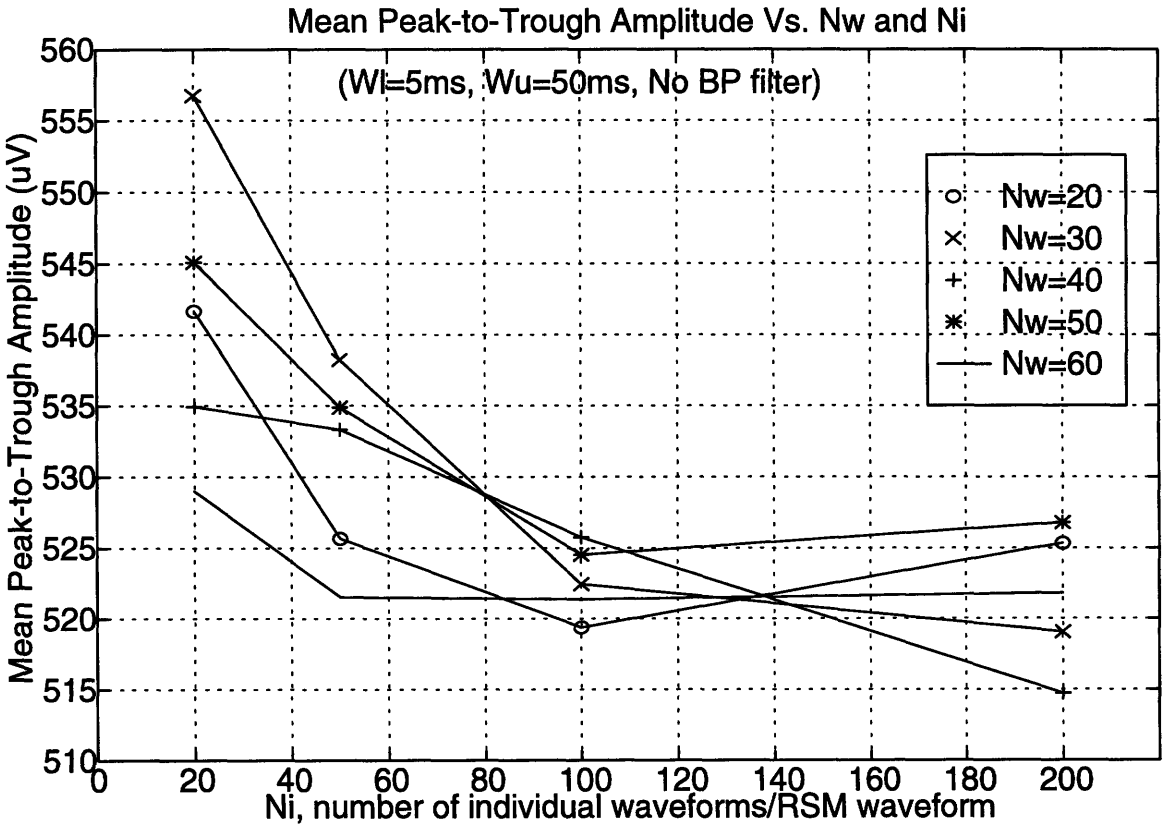

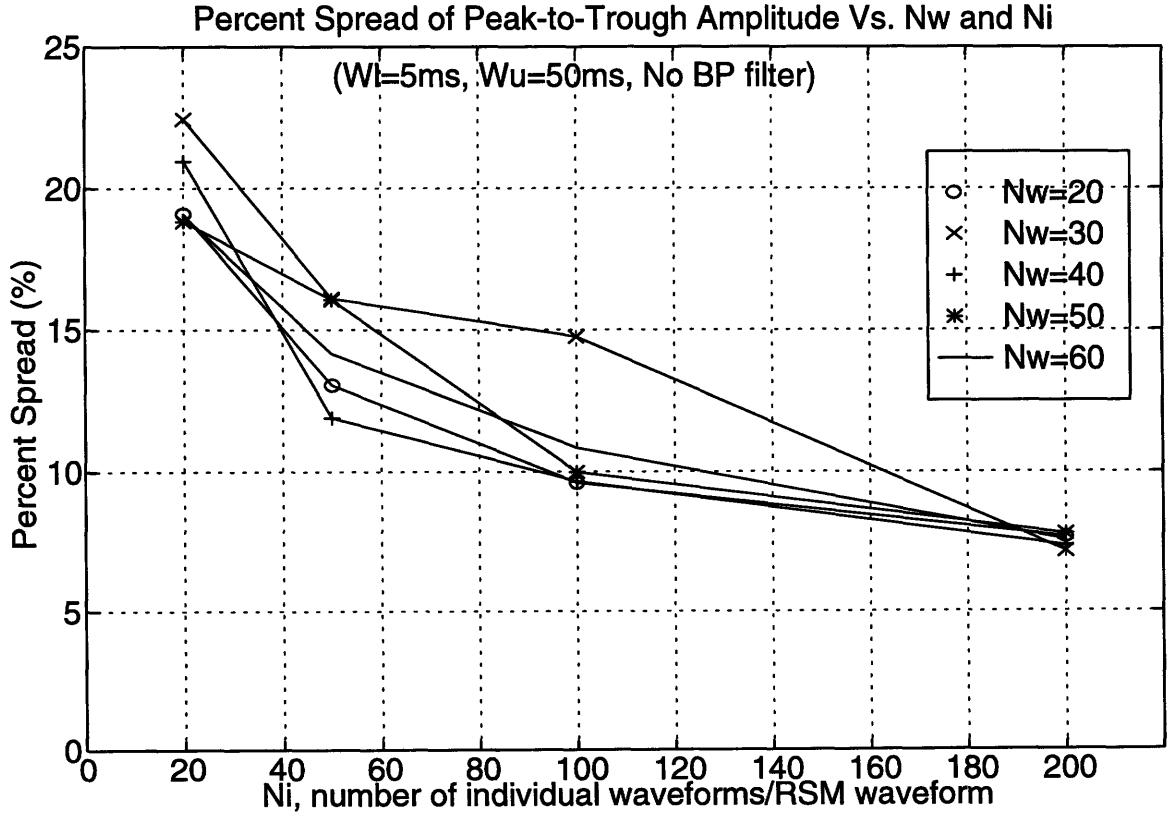

Figure 3.2 Mean RMS Feature and Percent Spread Versus N_w and N_i

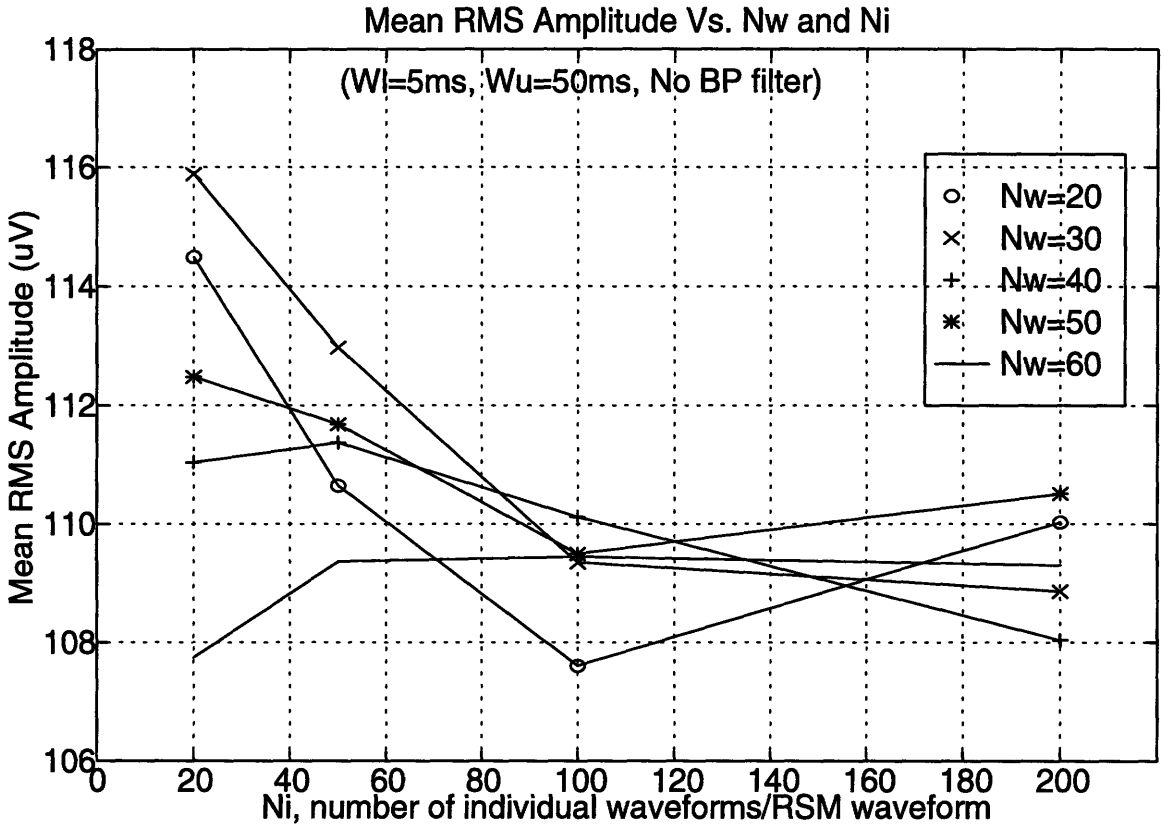

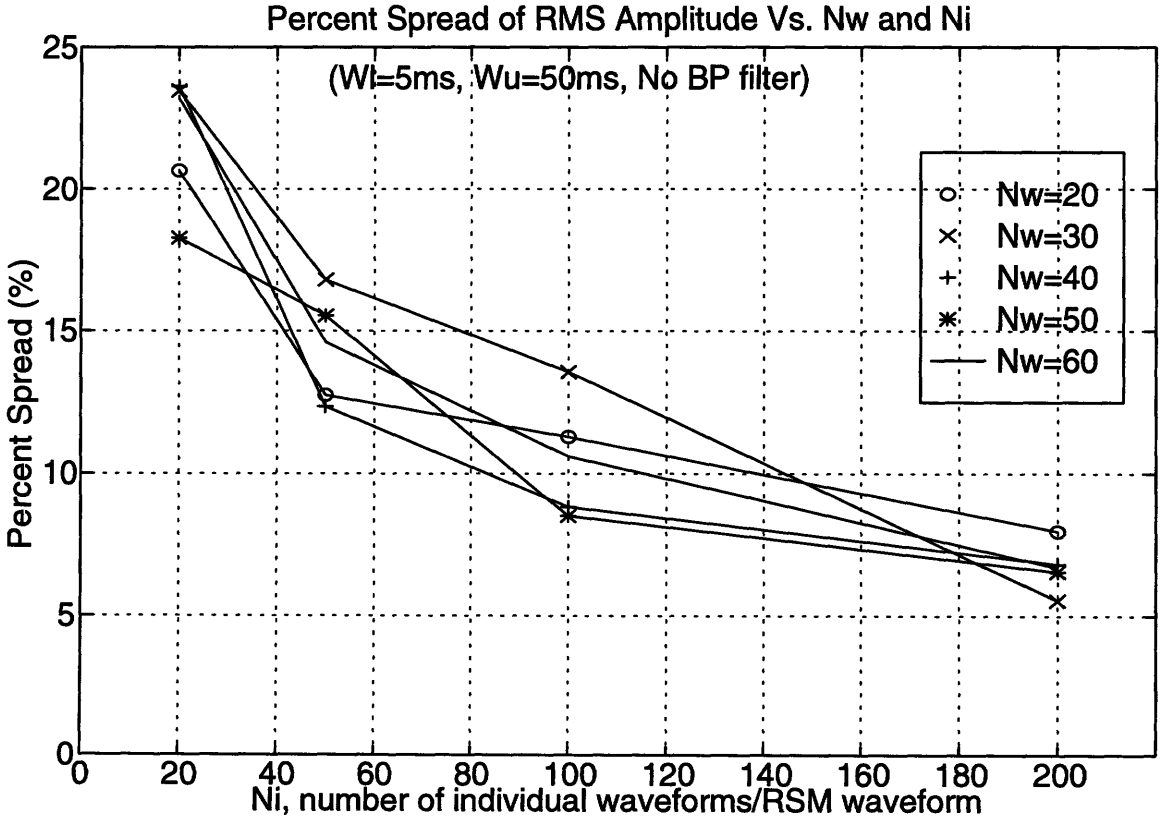

Consider first the mean peak-to-trough amplitude and corresponding percent spread plots. Note that the percent spread decreases with an increase in N_i for all values of N_w, although the difference in percent spread between $N_i = 100$ and $N_i = 200$ is relatively small when $N_w \geq 40$. This suggests that we should choose $N_i = 100$ since $N_i = 200$ represents only a modest reduction in percent spread while requiring twice the number of waveforms. Comparing the curves for the different values of N_w, the percent spread is roughly the same for $N_w \geq 40$ when $N_i = 100$ which suggests that we should choose $N_w = 40$.

Consider now the mean RMS amplitude feature and corresponding percent spread plots. Again the percent spread decreases with an increase in N_i for all values of N_w, and the difference in percent spread between $N_i = 100$ and $N_i = 200$ is relatively small for all values of $N_w \geq 40$. This confirms our choice of $N_i = 100$ and $N_w = 40$ from the peak-to-trough feature plots.

Tables 3.1 and 3.2 summarize the analysis of the 35 μV data for both features, where the case of $N_i = 1$ (RSM waveforms are identical to the individual waveforms) is added for comparison. Note the large disparity in the mean of the features when $N_i = 1$ compared to larger values of N_i due to the large amounts of noise which is not averaged out. For comparison, the peak-to-trough feature of the overall mean waveform is 522 μV, and the RMS amplitude feature of the overall mean waveform is 109 μV.

The same range values for the variables N_i and N_w was used for analysis of the other data sets, yielding the same conclusions, i.e. using $N_i = 100$ and $N_w = 40$. Based on the above analysis it was concluded that the subsequent feature extraction for both features should use $N_i = 100$ and $N_w = 40$ waveforms.

As an aside, the percent spread error criterion can be transformed into a measure of accuracy in our estimates by assuming that the distribution of the two features takes a Gaussian shape. The efficacy of the Gaussian assumption can be verified qualitatively by examination of the histograms of the RSM waveform features in Figure 3.3 when $N_w = 1000$, $N_i = 100$, $W_l = 5$ ms, $W_u = 50$ ms and no bandpass filtering was performed. Both histograms were fit to a Gaussian waveform with the same mean and sample variance as the feature distribution.

Figure 3.3 Histograms of Peak-to-Trough and RSM Features

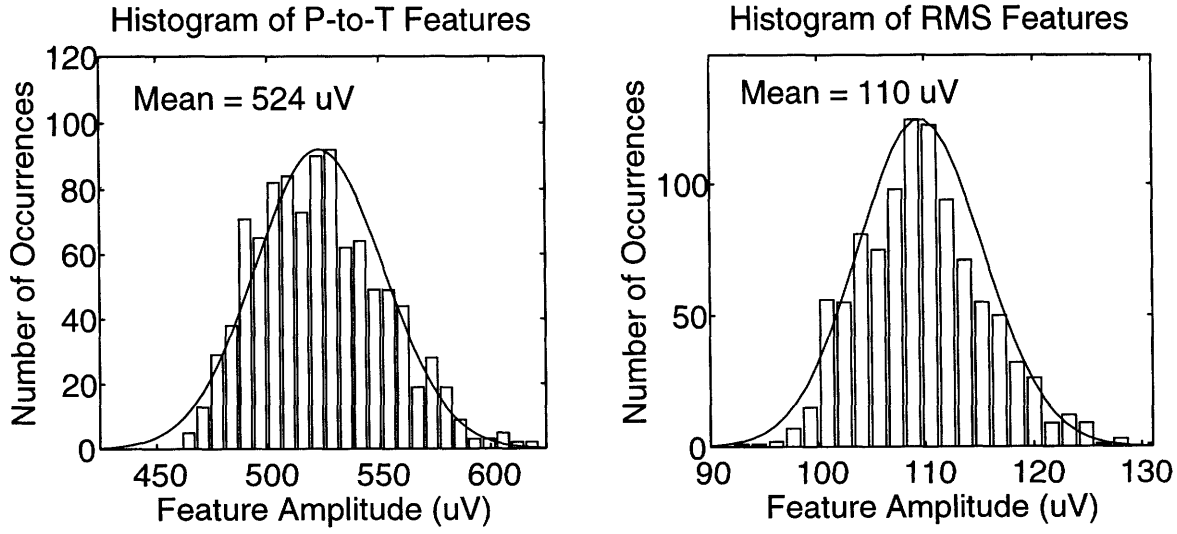

Table 3.1 Effects of N_i and N_w on Peak-to-Trough Feature Analysis

		Mean Feature (μV)	Sample Standard Deviation (μV)	Percent Spread (%)
$N_w = 20$	$N_i = 1$	719	195.35	54.4
	$N_i = 20$	542	51.70	19.1
	$N_i = 50$	526	34.28	13.0
	$N_i = 100$	519	24.94	9.6
	$N_i = 200$	525	20.19	7.7
$N_w = 30$	$N_i = 1$	744	150.47	40.4
	$N_i = 20$	557	62.50	22.5
	$N_i = 50$	538	43.33	16.1
	$N_i = 100$	522	38.52	14.7
	$N_i = 200$	519	18.64	7.2
$N_w = 40$	$N_i = 1$	702	146.97	41.8
	$N_i = 20$	535	56.06	21.0
	$N_i = 50$	533	31.69	11.9
	$N_i = 100$	526	25.37	9.7
	$N_i = 200$	515	18.97	7.4
$N_w = 50$	$N_i = 1$	707	183.66	51.9
	$N_i = 20$	545	51.31	18.8
	$N_i = 50$	535	42.93	16.0
	$N_i = 100$	525	26.19	10.0
	$N_i = 200$	527	20.53	7.8
$N_w = 60$	$N_i = 1$	740	183.52	49.6
	$N_i = 20$	529	50.32	19.0
	$N_i = 50$	522	36.92	14.2
	$N_i = 100$	521	28.25	10.8
	$N_i = 200$	522	19.78	7.6

Table 3.2 Effects of N_i and N_w on RMS Feature Analysis

		Mean Feature (μV)	Sample Standard Deviation (μV)	Percent Spread (%)
$N_w = 20$	$N_i = 1$	157	50.05	63.7
	$N_i = 20$	114	11.82	20.6
	$N_i = 50$	111	7.05	12.7
	$N_i = 100$	108	6.07	11.3
	$N_i = 200$	110	4.39	8.0
$N_w = 30$	$N_i = 1$	166	38.33	46.2
	$N_i = 20$	116	13.60	23.5
	$N_i = 50$	113	9.49	16.8
	$N_i = 100$	109	7.41	13.6
	$N_i = 200$	109	3.02	5.5
$N_w = 40$	$N_i = 1$	156	44.27	56.7
	$N_i = 20$	111	13.11	23.6
	$N_i = 50$	111	6.87	12.3
	$N_i = 100$	110	4.86	8.8
	$N_i = 200$	108	3.68	6.8
$N_w = 50$	$N_i = 1$	158	47.33	60.1
	$N_i = 20$	112	10.28	18.3
	$N_i = 50$	112	8.67	15.5
	$N_i = 100$	109	4.66	8.5
	$N_i = 200$	111	3.62	6.6
$N_w = 60$	$N_i = 1$	163	45.21	55.5
	$N_i = 20$	108	12.49	23.2
	$N_i = 50$	109	7.98	14.6
	$N_i = 100$	109	5.80	10.6
	$N_i = 200$	109	3.66	6.7

3.3.2 Determination of Waveform Window for Feature Extraction

Here we wish to find the lower and upper window times for feature extraction such that the percent spread is minimized. The results from section 3.3.1 above are utilized (N_i = 100 and N_w = 40), and no bandpass filtering was performed so the effect of windowing can be understood alone. Note that the results in this section for the given values of N_i and N_w will not be identical to the corresponding results in the previous section because a new set of RSM waveforms were formed (from the same set of 500 individual waveforms).

Due to the presence of the stimulus artifact which ends at 5 ms, the range of possible lower window times W_l starts at 5 ms. Due to the clinical significance of only the earliest portion of the evoked potential [1, 4] the range of possible upper window times W_u ends at 50 ms. Referring to Figure 3.4 below which shows the overall mean waveform, we note that the lower window time should occur prior to the large peak at about 13 ms. For this reason the range of lower window times was chosen to be from 5-10 ms, which allows for some variation in latency before the initial peak begins. Additionally, the lower end of the upper window time should occur after the large negative peak occurring at about 18 ms, which yields a range of upper window times from 20-50 ms, again allowing for latency variability, and the clinical interest in the second peak [20].

Based on the above, the lower window time W_l was varied between 5 and 9 ms in unit increments, and the upper window time W_u was varied between 20 and 50 ms in 5 ms increments. The results of the feature extraction process for both features is summarized in Figures 3.5 and 3.6 which show plots of the mean features and corresponding percent spread for the RSM waveforms.

Figure 3.4 Lower and Upper Window Ranges for Feature Extraction

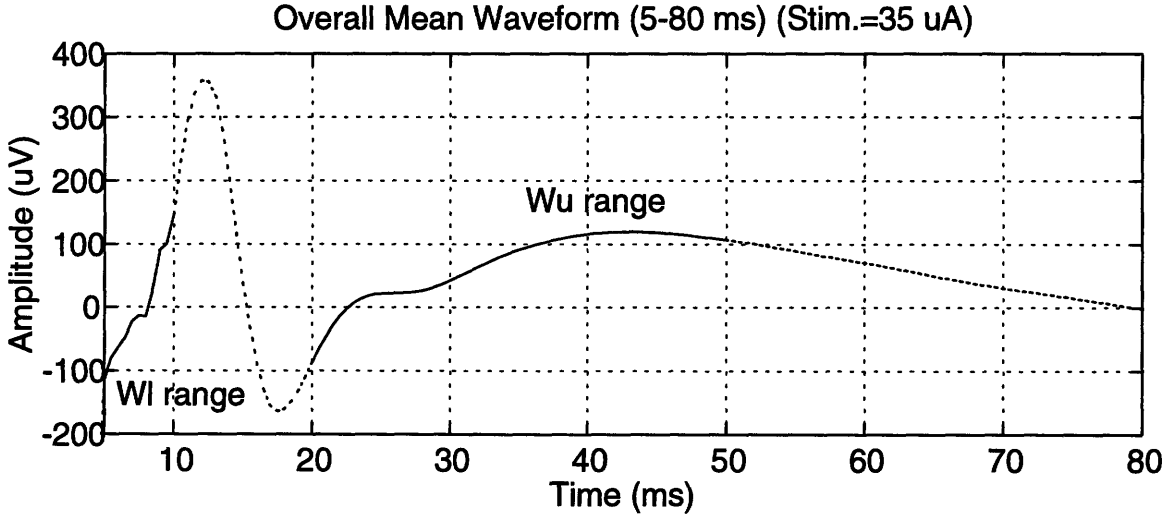

Figure 3.5 Mean Peak-to-Trough Feature and Percent Spread Versus W_l and W_u

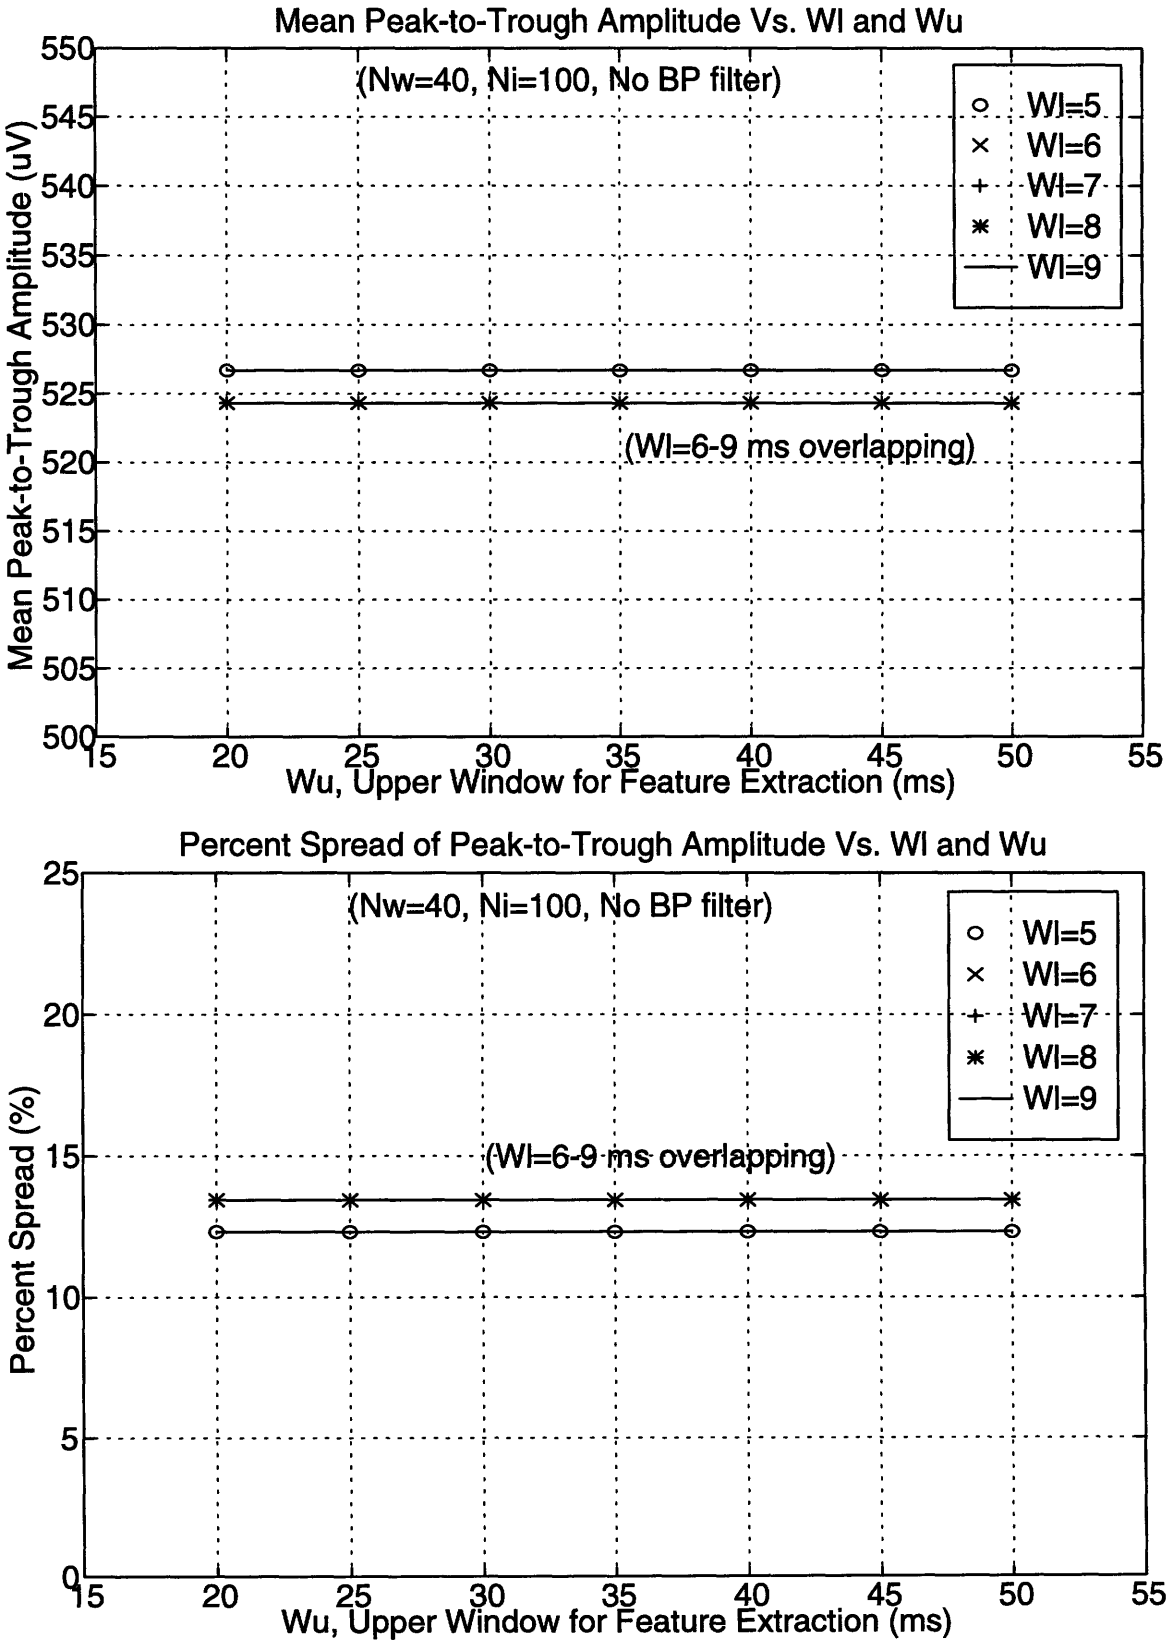

Figure 3.6 Mean RMS Feature and Sample Variance Versus W_l and W_u

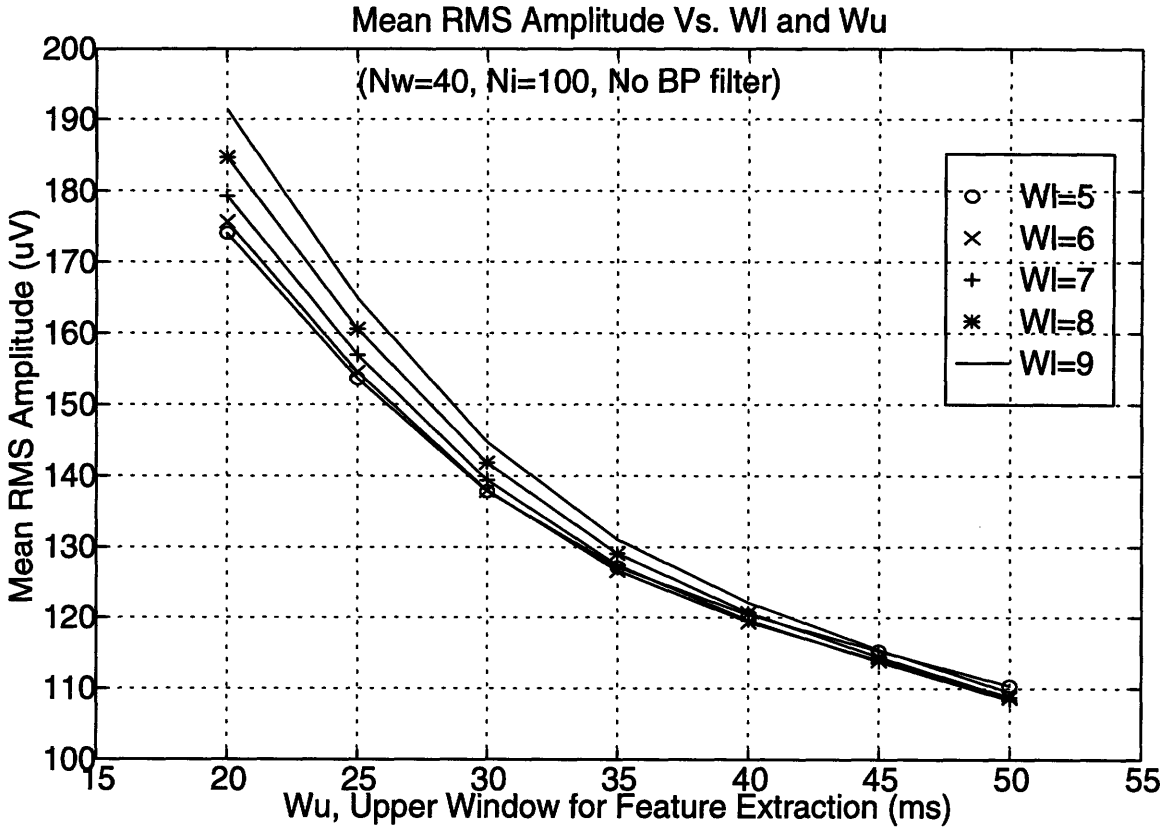

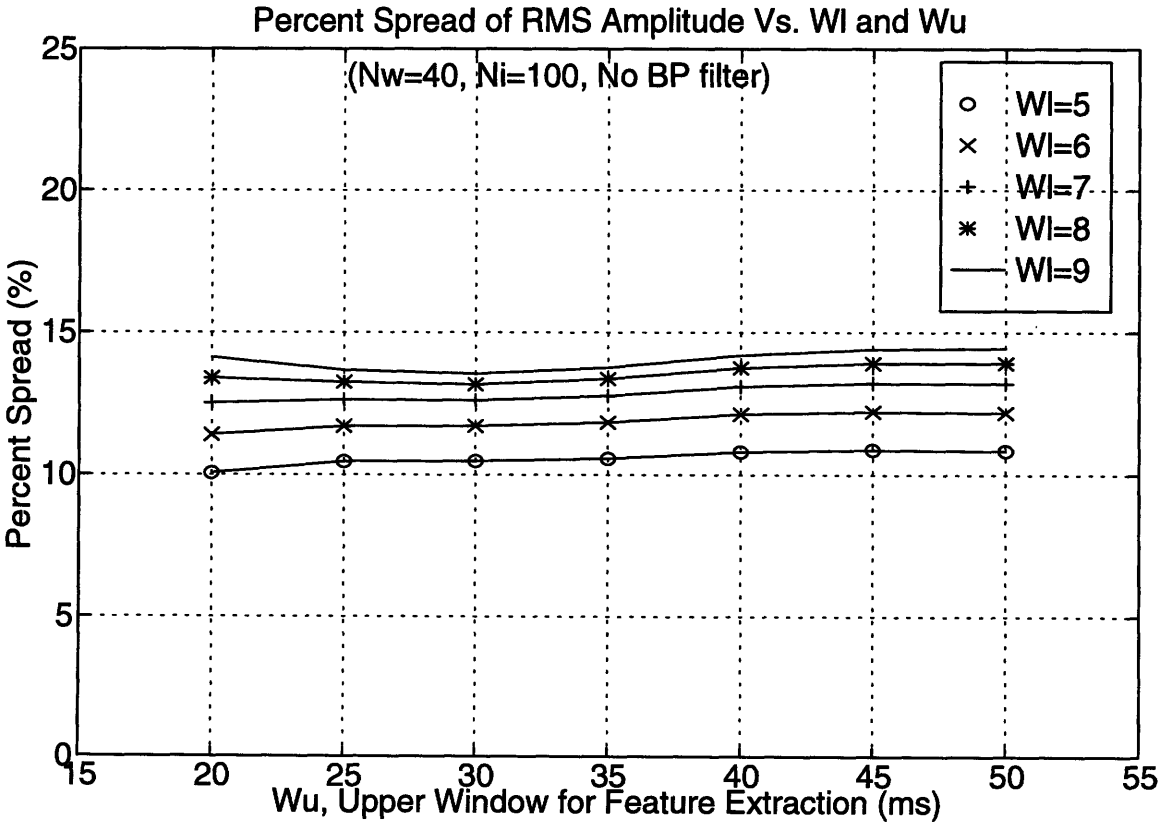

Consider first the mean peak-to-trough amplitude and corresponding percent spread plots. The mean feature size is shown to be independent of the upper window times used, which is expected when the peak and trough of the RSM waveforms occur prior to 20 ms. Note that the lower window time of $W_l = 5$ ms yields a slightly lower percent spread than larger values of W_l, suggesting that we should choose $W_l = 5$ ms. The choice for the upper window time is inconclusive from this data.

Consider now the mean RMS amplitude and corresponding percent spread plots. The mean feature size decreases with increasing W_u, for each value of W_l, which is again expected when the largest peak and trough occurs before 20 ms. However, note from the percent spread plot that the earliest lower window time of 5 ms yields the lowest percent spread for each upper window time used, suggesting that we should choose $W_l = 5$ ms, confirming our results from the peak-to-trough data discussed above. Note also that the percent spread is roughly constant across all values W_u, making the choice of upper window times inconclusive from this data.

Tables 3.3 and 3.4 summarize these results showing the mean feature, sample standard deviation and percent spread of the data about the mean for each feature, similar to Tables 3.1 and 3.2. For the sake of brevity, only the values of W_u equal to 20, 30, 40 and 50 ms are shown.

One interesting aspect of Table 3.4 is that the percent spread has a tendency to increase with an increase in W_u, despite a decrease in the sample standard deviation. This apparent anomaly can be explained by noting from the overall mean waveform plot in Figure 3.4 that as W_u is increased above 20 ms, the RMS amplitude feature decreases since it is averaging a larger portion of the waveform occurring after the maximum peak and trough. Although the sample standard deviation is also decreasing, it is not doing so as quickly as the RMS amplitude feature is decreasing, hence the ratio of two standard deviations to the mean feature value (i.e. the percent spread) is increasing.

Another interesting aspect of Table 3.4 is the fact that the sample standard deviation and percent spread tends to decrease with a decrease in W_l. This can be explained by the fact that the evoked response begins at the earliest lower window time examined (5 ms) immediately following the large stimulus artifact which is roughly uniform in shape and size for all individual waveforms. It is possible that the stimulus artifact limits the recorded waveform to a narrow range of values immediately after its occurrence, and at later times the recorded waveform is more subject to deviations from additive noise or amplitude scaling of the deterministic evoked potential.

Despite the slightly larger percent spread when $W_u = 50$ ms compared to $W_u = 20$ ms for $W_l = 5$ ms, we chose the values of $W_l = 5$ ms and $W_u = 50$ ms as our window width for both features. Although the choice of the upper window time is inconclusive from these data, we note that variations of EEP waveform shapes occur in different animals under different experiments, which includes variable latency of the initial large peak and trough. To avoid the possibility of missing these features, and because of the clinical significance of the later region [20], we decided to use the longer window time of $W_u = 50$ ms.

Table 3.3 Effects of W_l and W_u on Peak-to-Trough Feature Analysis

		Mean Feature (μV)	Sample Standard Deviation (μV)	Percent Spread (%)
$W_l = 5$	$W_u = 20$	527	32.40	12.3
	$W_u = 30$	527	32.40	12.3
	$W_u = 40$	527	32.40	12.3
	$W_u = 50$	527	32.40	12.3
$W_l = 6$	$W_u = 20$	524	35.23	13.4
	$W_u = 30$	524	35.23	13.4
	$W_u = 40$	524	35.23	13.4
	$W_u = 50$	524	35.23	13.4
$W_l = 7$	$W_u = 20$	524	35.23	13.4
	$W_u = 30$	524	35.23	13.4
	$W_u = 40$	524	35.23	13.4
	$W_u = 50$	524	35.23	13.4
$W_l = 8$	$W_u = 20$	524	35.23	13.4
	$W_u = 30$	524	35.23	13.4
	$W_u = 40$	524	35.23	13.4
	$W_u = 50$	524	35.23	13.4
$W_l = 9$	$W_u = 20$	524	35.23	13.4
	$W_u = 30$	524	35.23	13.4
	$W_u = 40$	524	35.23	13.4
	$W_u = 50$	524	35.23	13.4

Table 3.4 Effects of W_l and W_u on RMS Feature Analysis

		Mean Feature (μV)	Sample Standard Deviation (μV)	Percent Spread (%)
$W_l = 5$	$W_u = 20$	174	8.75	10.1
	$W_u = 30$	138	7.22	10.5
	$W_u = 40$	120	6.50	10.8
	$W_u = 50$	110	5.98	10.8
$W_l = 6$	$W_u = 20$	176	10.00	11.4
	$W_u = 30$	138	8.07	11.7
	$W_u = 40$	119	7.23	12.1
	$W_u = 50$	109	6.62	12.2
$W_l = 7$	$W_u = 20$	179	11.21	12.5
	$W_u = 30$	139	8.78	12.6
	$W_u = 40$	120	7.83	13.1
	$W_u = 50$	108	7.15	13.2
$W_l = 8$	$W_u = 20$	185	12.35	13.4
	$W_u = 30$	142	9.34	13.2
	$W_u = 40$	121	8.29	13.7
	$W_u = 50$	109	7.56	13.9
$W_l = 9$	$W_u = 20$	191	13.50	14.1
	$W_u = 30$	145	9.79	13.5
	$W_u = 40$	122	8.67	14.2
	$W_u = 50$	109	7.90	14.4

4. Ideal Bandpass Filtering

4.1 Statement of Goals, Assumptions and Methods

The goal here is to use the RSM waveforms as input to an approximate ideal bandpass filter designed so that the output closely resembles the desired underlying evoked response signal. Our measure of "resembles" is in terms of the peak-to-trough and RSM amplitude about the time average features, and their similarity to the corresponding features of the overall mean waveform. The comparison of the features to the overall mean waveform is made indirectly through the use of the percent spread about the mean. The method used finds the percent spread associated with RSM waveform features for each cutoff frequency, and then varies the cutoff frequency finding the minimum percent spread. The initial analysis is done on the 35 μA data, and the results are then applied to the other data sets for comparison.

4.2 Approximation to Ideal Bandpass Filtering

True ideal bandpass filtering can only be accomplished by using an infinite length impulse response, which is not possible for computerized implementation. The truncation of the infinite length impulse response by different windowing functions yields different filters which can be distinguished in the frequency domain by passband and stopband ripple, transition widths between passbands and stopbands and by numerous other parameters [17]. The approach taken here is different in that we start with an idealized frequency response in the Discrete Fourier Transform (DFT) domain where the bandpass filter is unity in the passband and zero in the stop band, with a single sample being the transition width. Since multiplication of Fourier transforms in frequency is equivalent to the Fourier transform of the convolution of the time functions, we can consider the effect of such a process in the time domain. Essentially multiplying in frequency the boxcar filter with the Fourier transform of the time signal is the Fourier transform of a sinc function convolved with the time function. Although this method produces the Gibbs effect [17], its effects are minimally intrusive for our purposes.

4.3 Determination of the Range of Cutoff Frequencies

Here we wish to determine the effects of changing lower and upper bandpass filter cutoff frequencies f_l and f_u such that the percent spread of the features is minimized. Clearly we cannot use any arbitrary values for these cutoff frequencies since by using an extremely narrow filter we can virtually annihilate the signal, producing an output which is very uniform with a very small sample variance, but whose feature sizes will not represent the evoked response. We thus begin by establishing the range of possible cutoff frequencies such that our evoked response is kept sufficiently intact. We determine the appropriate range of cutoff frequencies by examining the individual effects of lowpass and highpass filtering applied to the overall mean waveform. Specifically this test shows the minimum lowpass cutoff frequency and the maximum highpass cutoff frequency to be used in our experiment. Figures 4.1 and 4.2 below show the effects of ideal lowpass and highpass filtering respectively for the range of cutoff frequencies indicated.

From Figure 4.1 it is apparent that lowpass filtering with a cutoff of $f_u = 200$ Hz produces only minimal effects on the overall mean waveform, although filtering with a cutoff of $f_u = 100$ Hz produces large oscillating deviations. For this reason the lowpass

cutoff frequencies examined are no smaller than 200 Hz. From Figure 4.2 it is apparent that each of the cutoff frequencies used preserves the large peak and trough between 10 and 20 ms, but the second large peak at about 40 ms is not well preserved for $f_l \geq 20$ Hz. For this reason the highpass cutoff frequencies are limited to being less than 20 Hz.

Figure 4.1 Effects of Ideal Lowpass Filtering the Overall Mean Waveform

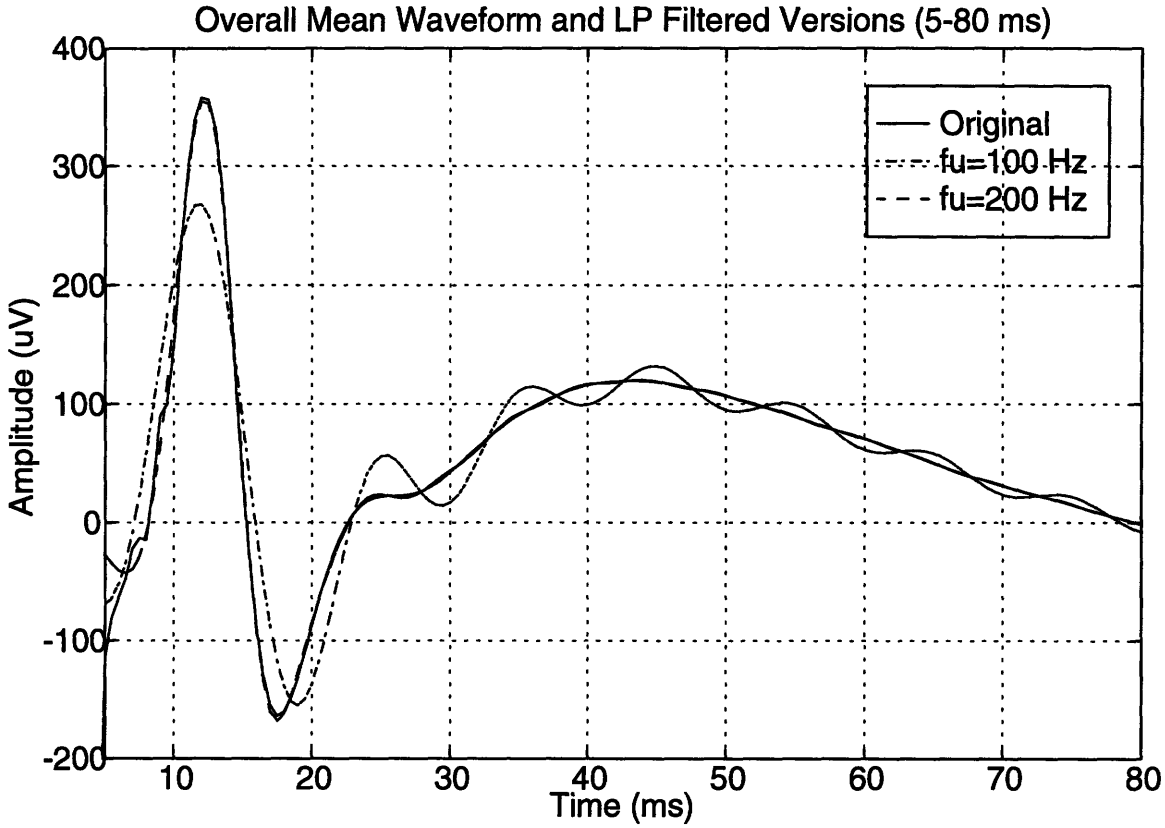

Figure 4.2 Effects of Ideal Highpass Filtering the Overall Mean Waveform

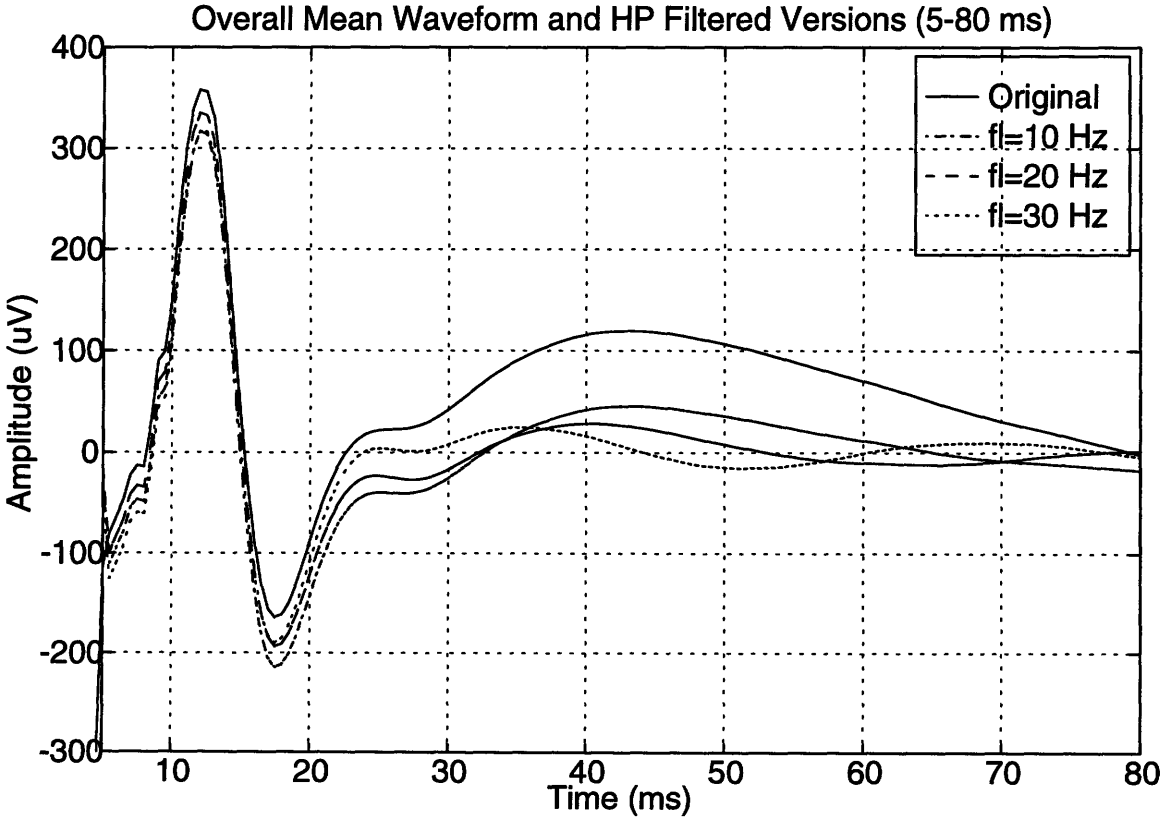

4.4 Bandpass Filtering Recorded Waveforms

Given the results of lowpass and highpass filtering of the overall mean waveform discussed above, and the previously determined values of $N_i = 100$, $N_w = 40$, $W_l = 5$ ms, and $W_u = 50$ ms, the lower cutoff frequency f_l was varied between 0 and 16 Hz in 4 Hz increments, and the upper cutoff frequency f_u was varied between 200 and 1000 Hz in 200 Hz increments. The case of $f_l = 0$ Hz, $f_u = 1000$ Hz is included for comparison when no bandpass filtering is performed. Figures 4.3 and 4.4 below summarize the results for the Peak-to-Trough and RMS amplitude feature extractions respectively.

Consider first the mean peak-to-trough amplitude and corresponding percent spread plots. From the mean feature plot it is apparent that the mean feature size is somewhat independent of both the lower and upper cutoff frequencies which is desirable since we do not want the filtering to affect the measurement of the feature. However, the percent spread plot shows that the effect of bandpass filtering produces minimal reduction of the percent spread. The smallest percent spread occurred when $f_l = 0$ Hz and $f_u = 600$ Hz. Furthermore $f_l = 0$ Hz produced the smallest percent spread for all values of f_u, and $f_u = 600$ Hz produced the smallest sample variance for all values of f_l.

Consider now the mean RMS amplitude and corresponding percent spread plots Again the mean feature size is mostly independent of both the lower and upper cutoff frequencies (note the narrow range of the amplitude scale). The percent spread plot shows virtually no effect due to lowpass filtering (varying f_u) and only minimal effect due to highpass filtering (varying f_l). The smallest percent spread occurred when $f_l = 4$ Hz and $f_u = 1000$ Hz. Furthermore $f_l = 4$ Hz produced the smallest sample variance for all values of f_u, and $f_u = 1000$ Hz produced the smallest sample variance for all values of f_l.

Tables 4.1 and 4.2 summarize this analysis in terms of the mean feature size, sample standard deviation and percent spread about the mean feature size.

Figure 4.3 Mean Peak-to-Trough Feature and Sample Variance Versus f_l and f_u

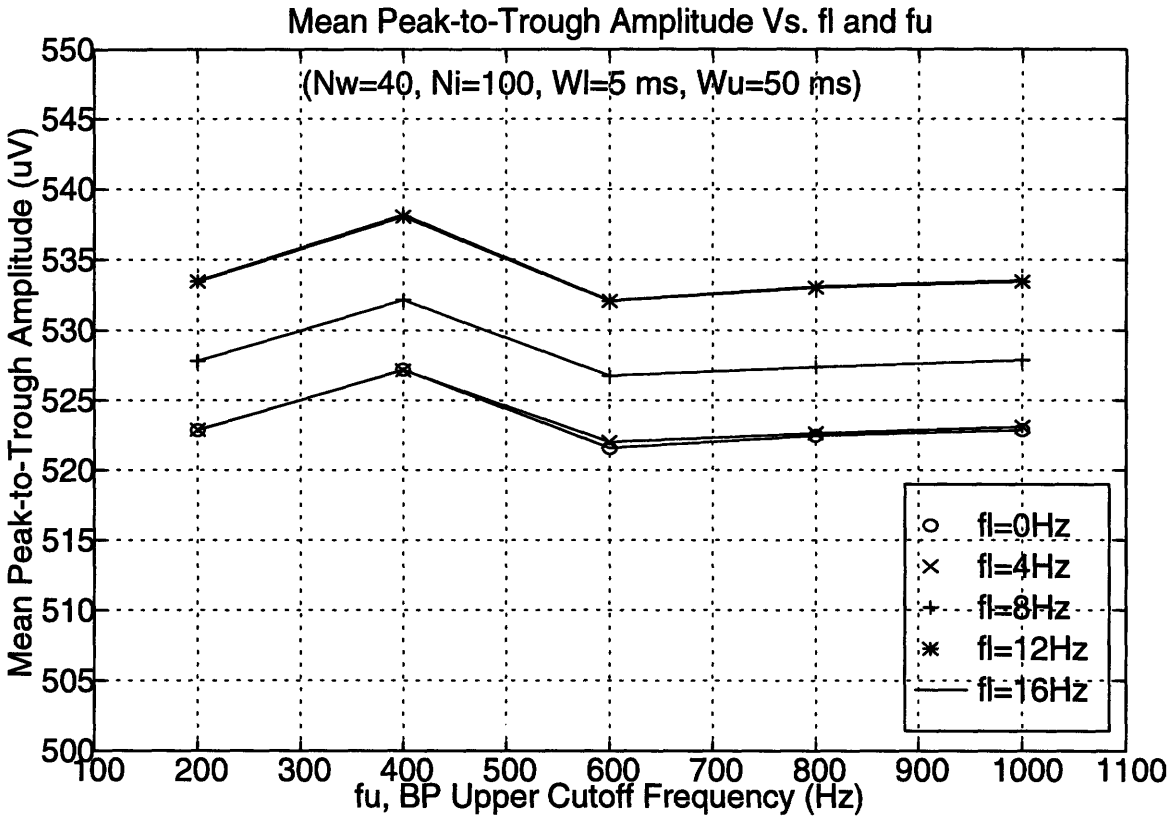

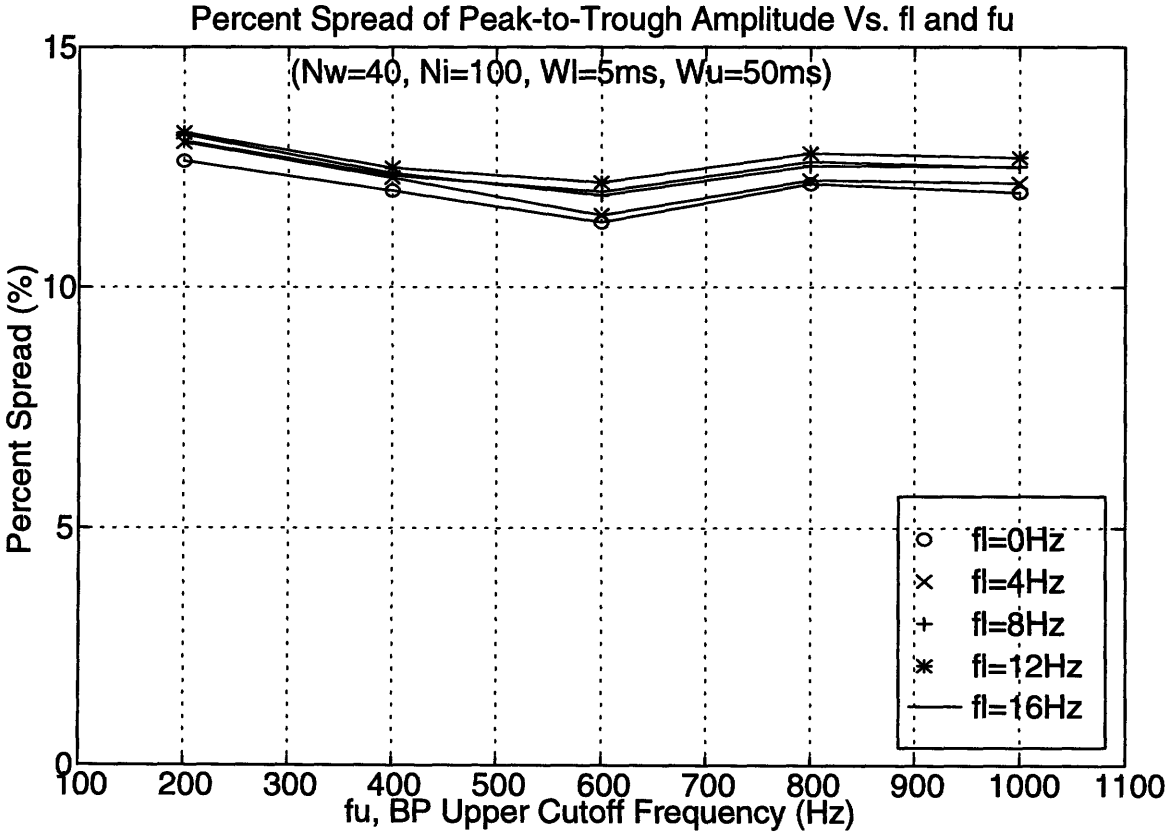

Figure 4.4 Mean RMS Feature and Sample Variance Versus f_l and f_u

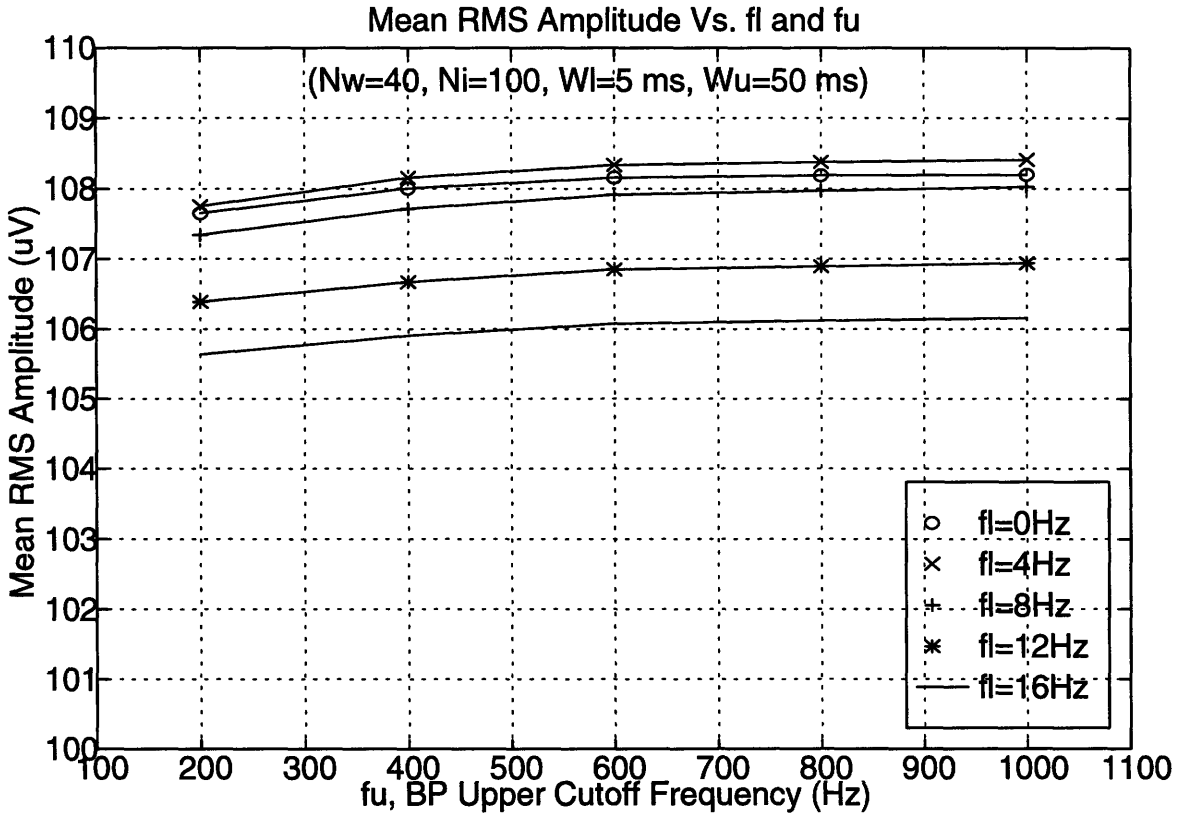

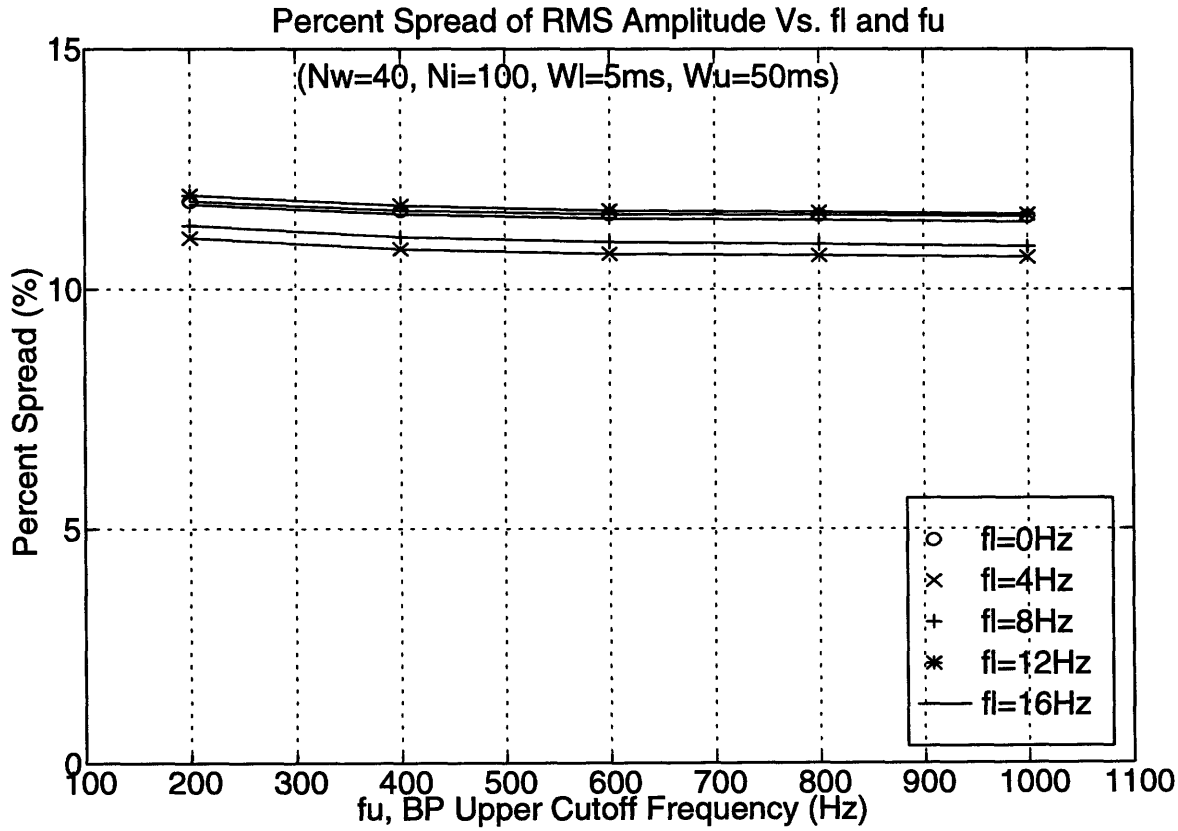

Table 4.1 Effects of f_l and f_u on Peak-to-Trough Feature Analysis (N_i = 100)

		Mean Feature (μV)	Sample Standard Deviation (μV)	Percent Spread (%)
f_l = 0	f_u = 200	523	33.02	12.6
	f_u = 400	527	31.65	12.0
	f_u = 600	522	29.60	11.4
	f_u = 800	522	31.74	12.2
	f_u = 1000	523	31.28	12.0
f_l = 4	f_u = 200	523	34.03	13.0
	f_u = 400	527	32.36	12.3
	f_u = 600	522	30.01	11.5
	f_u = 800	523	31.94	12.2
	f_u = 1000	523	31.81	12.2
f_l = 8	f_u = 200	528	34.75	13.2
	f_u = 400	532	32.91	12.4
	f_u = 600	527	31.37	11.9
	f_u = 800	527	33.03	12.5
	f_u = 1000	528	33.00	12.5
f_l = 12	f_u = 200	533	35.25	13.2
	f_u = 400	538	33.58	12.5
	f_u = 600	532	32.40	12.2
	f_u = 800	533	34.07	12.8
	f_u = 1000	533	33.85	12.7
f_l = 16	f_u = 200	533	34.78	13.0
	f_u = 400	538	33.16	12.3
	f_u = 600	532	31.90	12.0
	f_u = 800	533	33.61	12.6
	f_u = 1000	534	33.32	12.5

Table 4.2 Effects of f_l and f_u on RMS Feature Analysis (N_i = 100)

		Mean Feature (μV)	Sample Standard Deviation (μV)	Percent Spread (%)
f_l = 0	f_u = 200	108	6.37	11.8
	f_u = 400	108	6.28	11.6
	f_u = 600	108	6.25	11.6
	f_u = 800	108	6.24	11.5
	f_u = 1000	108	6.23	11.5
f_l = 4	f_u = 200	108	5.96	11.1
	f_u = 400	108	5.86	10.8
	f_u = 600	108	5.81	10.7
	f_u = 800	108	5.80	10.7
	f_u = 1000	108	5.78	10.7
f_l = 8	f_u = 200	107	6.08	11.3
	f_u = 400	108	5.97	11.1
	f_u = 600	108	5.92	11.0
	f_u = 800	108	5.90	10.9
	f_u = 1000	108	5.88	10.9
f_l = 12	f_u = 200	106	6.36	12.0
	f_u = 400	107	6.26	11.7
	f_u = 600	107	6.22	11.6
	f_u = 800	107	6.20	11.6
	f_u = 1000	107	6.18	11.6
f_l = 16	f_u = 200	106	6.21	11.8
	f_u = 400	106	6.12	11.6
	f_u = 600	106	6.08	11.5
	f_u = 800	106	6.07	11.4
	f_u = 1000	106	6.05	11.4

Given the absence of appreciable reduction in percent spread when varying the upper and lower cutoff frequencies, two different values of N_i were used in an effort to (a) determine if bandpass filtering has any affect on reducing percent spread, and (b) to potentially reduce the number of waveforms needed while still obtaining an accurate estimate of feature sizes. The feature extraction process was repeated when N_i was reduced to 50 and 20 waveforms per RSM waveform, using the same number of RSM waveforms ($N_w = 40$), using the same window ($W_l = 5$ ms, $W_u = 50$ ms) and using the same range of values for the lower and upper cutoff frequencies. Tables 4.3 and 4.4 summarize these results in terms of percent spread, including the case of $N_i = 100$ discussed above.

Again the effects of bandpass filtering produce only a negligible change in the percent spread when we reduce N_i to only 50 for both features. As expected the percent spread is increased since we are averaging fewer waveforms, but comparing all permutations of f_l and f_u to the case of $f_l = 0$ and $f_u = 1000$ (no filtering) reveal minimal changes. Despite the above, several identical conclusions can be drawn from this data when comparing it to the case of $N_i = 100$. For the peak-to-trough feature the smallest percent spread also occurred when $f_l = 0$ Hz and $f_u = 600$ Hz, and further $f_l = 0$ Hz produced the smallest percent spread for all values of f_u, and $f_u = 600$ Hz produced the smallest percent spread for all values of f_l. For the RMS feature, the smallest percent spread also occurred when $f_l = 4$ Hz and $f_u = 1000$ Hz, and further $f_l = 4$ Hz produced the smallest percent spread for all values of f_u, and $f_u = 1000$ Hz produced the smallest percent spread for all values of f_l.

Finally the effects of bandpass filtering produced only a negligible change in the percent spread when $N_i = 20$. Despite this, some conclusions can be drawn which are consistent with the analysis of data when N_i was equal to 100 and 50. First, the peak-to-trough amplitude feature again showed that the smallest percent spread occurred when $f_l = 0$ Hz and $f_u = 600$ Hz, and further $f_l = 0$ Hz produced the smallest percent spread for all values of f_u, and $f_u = 600$ produced the smallest percent spread for all values of f_l. For the RMS feature the results were slightly different for this value of N_i. Here, the smallest percent spread occurred when $f_l = 16$ Hz (contrasted with $f_l = 4$ Hz) and $f_u = 800$ or $f_u = 1000$ Hz (the same).

Table 4.3 Summary of Effects of f_l and f_u on Peak-to-Trough Feature Analysis

		Percent Spread (%)		
		N_i = 100	N_i = 50	N_i = 20
f_l = 0	f_u = 200	12.6	17.2	24.2
	f_u = 400	12.0	16.4	22.4
	f_u = 600	11.4	15.2	21.1
	f_u = 800	12.2	16.0	21.9
	f_u = 1000	12.0	15.6	21.3
f_l = 4	f_u = 200	13.0	17.7	24.4
	f_u = 400	12.3	16.5	22.7
	f_u = 600	11.5	15.3	21.3
	f_u = 800	12.2	16.0	22.2
	f_u = 1000	12.2	15.8	21.8
f_l = 8	f_u = 200	13.2	17.6	24.7
	f_u = 400	12.4	16.8	23.4
	f_u = 600	11.9	15.9	22.3
	f_u = 800	12.5	16.5	23.0
	f_u = 1000	12.5	16.4	22.8
f_l = 12	f_u = 200	13.2	17.5	24.7
	f_u = 400	12.5	16.9	23.6
	f_u = 600	12.2	16.5	22.9
	f_u = 800	12.8	17.0	23.6
	f_u = 1000	12.7	16.9	23.3
f_l = 16	f_u = 200	13.0	17.3	24.4
	f_u = 400	12.3	16.7	23.1
	f_u = 600	12.0	16.4	22.7
	f_u = 800	12.6	16.9	23.3
	f_u = 1000	12.5	16.7	23.1

Table 4.4 Summary of Effects of f_l and f_u on RMS Feature Analysis

		Percent Spread (%)		
		N_i = 100	N_i = 50	N_i = 20
f_l = 0	f_u = 200	11.8	15.0	24.2
	f_u = 400	11.6	14.8	23.9
	f_u = 600	11.6	14.7	23.8
	f_u = 800	11.5	14.7	23.8
	f_u = 1000	11.5	14.7	23.8
f_l = 4	f_u = 200	11.1	14.2	24.0
	f_u = 400	10.8	13.9	23.7
	f_u = 600	10.7	13.8	23.6
	f_u = 800	10.7	13.8	23.5
	f_u = 1000	10.7	13.7	23.5
f_l = 8	f_u = 200	11.3	15.0	24.0
	f_u = 400	11.1	14.8	23.6
	f_u = 600	11.0	14.7	23.5
	f_u = 800	10.9	14.6	23.4
	f_u = 1000	10.9	14.6	23.4
f_l = 12	f_u = 200	12.0	15.9	23.7
	f_u = 400	11.7	15.6	23.4
	f_u = 600	11.6	15.5	23.3
	f_u = 800	11.6	15.5	23.2
	f_u = 1000	11.6	15.5	23.2
f_l = 16	f_u = 200	11.8	15.7	22.4
	f_u = 400	11.6	15.5	22.1
	f_u = 600	11.5	15.4	22.0
	f_u = 800	11.4	15.4	21.9
	f_u = 1000	11.4	15.4	21.9

4.5 Summary of Bandpass Filtering Results for All Stimulus Levels

The above experiment was repeated for the other three stimulus levels of 15, 10, and 7.5 μA when N_i was equal to 100, 50 and 20, using the same number of RSM waveforms (N_w = 40),the same window (W_l = 5 ms, and W_u = 50 ms), and the same lower and upper cutoff frequencies. Although the results are not presented here in the same degree of detail as the 35 μA data, the most salient results are summarized along with the 35 μA data in Tables 4.5a-d below. The best case row in each table represents the configuration of lower and upper cutoff frequencies which yielded the lowest percent spread when filtering was performed.

Examination of Tables 4.5a-d reveal some noteworthy general trends. For the peak-to-trough amplitude feature, the minimum percent spread generally occurred when f_l = 0 Hz, the exception being for the 7.5 μA data where the minimum occurred for f_l was equal to 4 or 8 Hz. Additionally, the lowest percent spread generally occurred for f_u = 600 Hz, with one exception for the 15 μA data, and a few exceptions for the 7.5 μA data where f_u was equal to 200 Hz or 1 KHz. For the RMS amplitude feature, the lowest percent spread generally occurred for f_l = 4 Hz, with f_u equal to 600-1000 Hz. Note also that the best case configurations reduced the percent spread about the mean for both features generally by less than 2% for all data sets except the 7.5 μA data, where it was reduced by as much as about 8%, although the percent spread was much higher for this data set by comparison.

Given the results discussed above, we conclude that bandpass filtering yield no appreciable reduction in percent spread for any of the four data sets we examined, for the values of N_i equal to 100, 50, and 20. Although there were some substantial reductions in percent spread for the 7.5 μA data when N_i = 20, the percent spread still made our feature estimates highly variable. For these reasons we conclude that no ideal bandpass filtering should be performed.

A final overview of the feature sizes as a function of stimulus is presented in Figures 4.6 and 4.7 which present the results when no ideal bandpass filtering was performed and when N_i = 100 waveforms. Error bars equalling two sample standard deviations are included. As expected, there is a general trend of increasing feature size with increasing stimulus level for both features.

Table 4.5a Highlight of Bandpass Filtering Results for 35 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No filter	523	31.28	12.0	108	6.23	11.5
	Best	522	29.60	11.4	108	5.78	10.7
	(f_l, f_u)	(0,600)			(4,1K)		
Ni=50	No filter	526	40.91	15.6	109	8.00	14.7
	Best	524	39.95	15.2	109	7.50	13.7
	(f_l, f_u)	(0,600)			(4,1K)		
Ni=20	No filter	527	56.11	21.3	111	13.14	23.8
	Best	526	55.48	21.1	107	11.67	21.9
	(f_l, f_u)	(0,600)			(16,1K)		

Table 4.5b Highlight of Bandpass Filtering Results for 15 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread (2 SD's)	Mean (μV)	S. Dev. (μV)	%spread (2 SD's)
Ni=100	No filter	369	27.68	15.0	79	5.40	13.7
	Best	372	26.82	14.4	78	5.03	12.9
	(f_l, f_u)	(0,200)			(4,1K)		
Ni=50	No filter	373	43.23	23.2	80	8.73	21.8
	Best	370	42.76	23.1	79	8.25	20.9
	(f_l, f_u)	(0,600)			(4,1K)		
Ni=20	No filter	385	50.13	26.0	81	11.22	27.6
	Best	382	50.04	26.2	81	11.19	27.5
	(f_l, f_u)	(0,600)			(0,600)		

Table 4.5c Highlight of Bandpass Filtering Results for 10 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No filter	188	15.26	16.3	41	3.95	19.5
	Best	189	14.58	15.4	39	3.60	18.2
	(f_l, f_u)	(0,600)			(4,1K)		
Ni=50	No filter	190	20.24	21.3	42	5.41	25.9
	Best	192	19.90	20.8	41	4.92	24.3
	(f_l, f_u)	(0,600)			(4,1K)		
Ni=20	No filter	207	30.31	29.3	44	6.70	30.7
	Best	208	28.14	27.1	44	6.68	30.7
	(f_l, f_u)	(0,600)			(0,800)		

Table 4.5d Highlight of Bandpass Filtering Results for 7.5 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No filter	108	11.12	20.6	21	1.84	17.2
	Best	89	8.36	18.8	20	1.69	16.5
	(f_l, f_u)	(8,200)			(4,600)		
Ni=50	No filter	115	17.34	30.3	23	3.28	28.8
	Best	102	11.25	22.0	22	2.62	24.3
	(f_l, f_u)	(8,1K)			(16,1K)		
Ni=20	No filter	123	26.33	42.7	26	6.14	47.8
	Best	104	20.88	40.0	24	4.67	38.9
	(f_l, f_u)	(4,200)			(12,600)		

Figure 4.6 Peak-to-Trough Amplitude as a Function of Stimulus Level (No filtering)

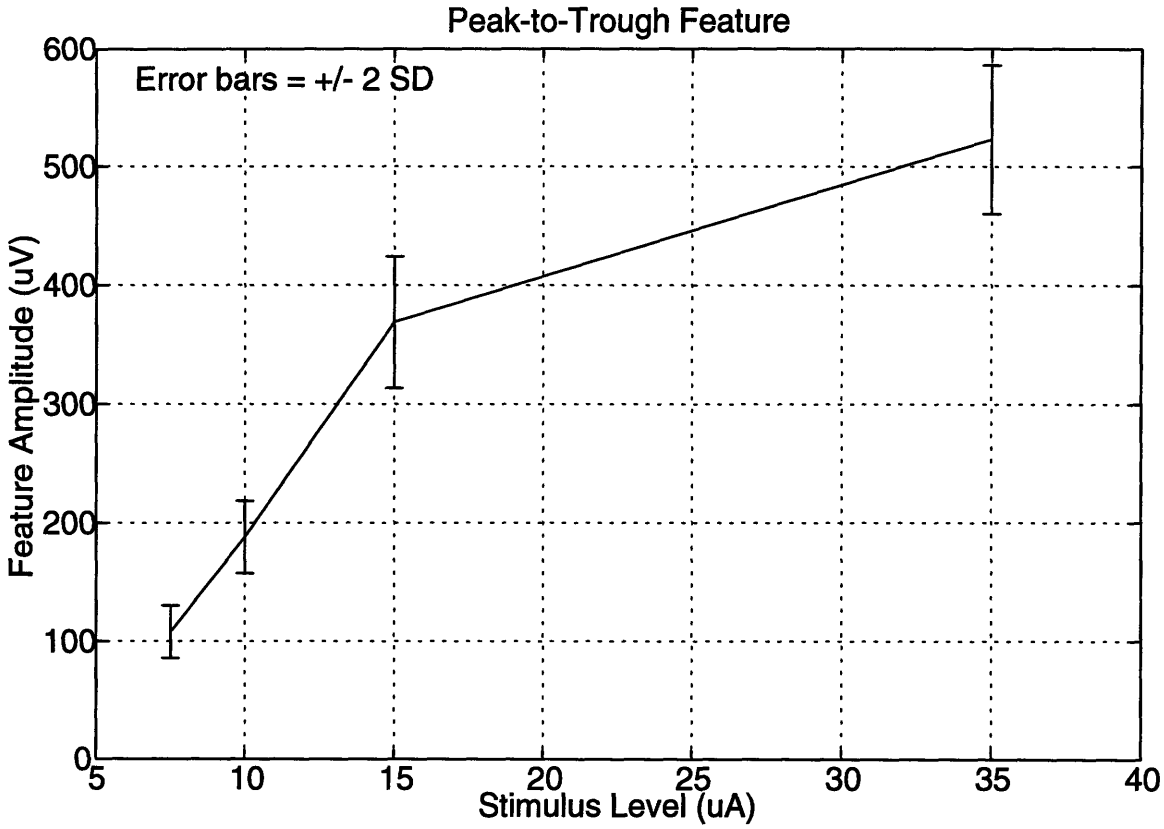

Figure 4.7 RMS Amplitude as a Function of Stimulus Level (No filtering)

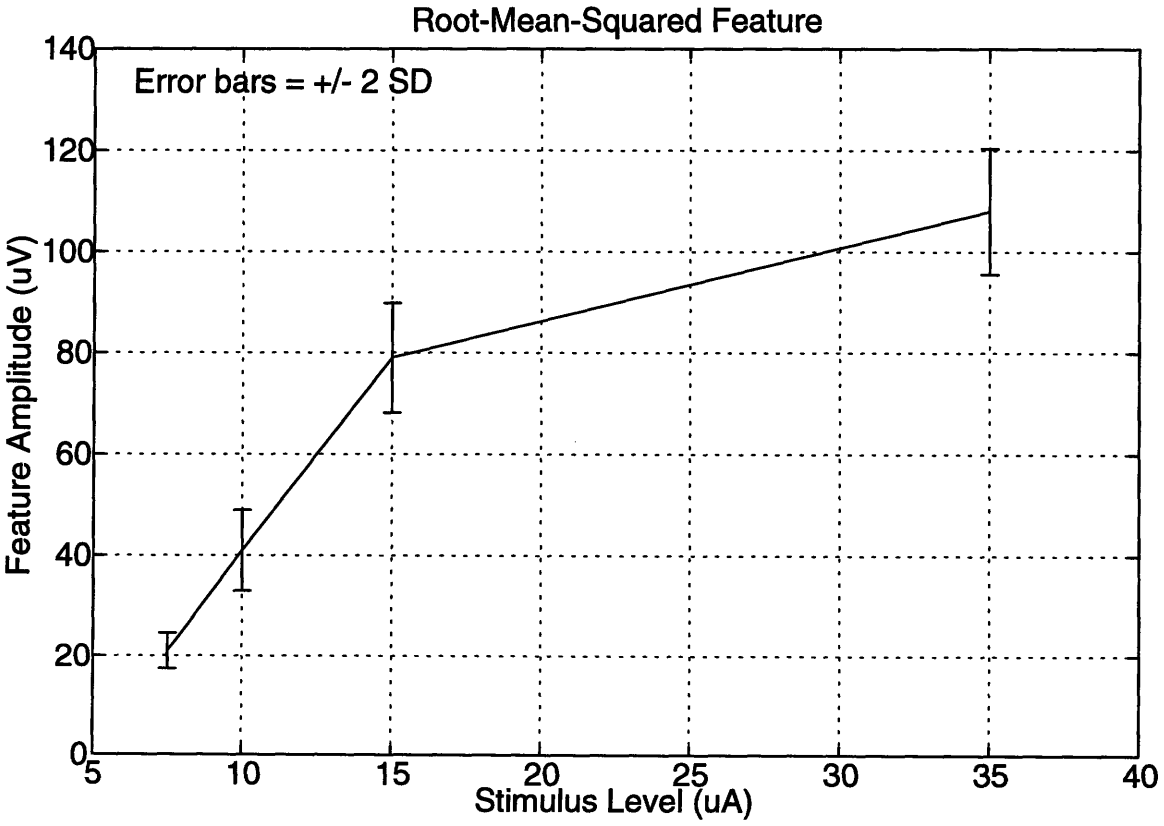

5. Weiner-Type Filtering

5.1 Statement of Goals, Assumptions and Methods

The goal here is to filter RSM waveforms with a single LTI filter such that the output of the filter best resembles the evoked potential in a mean-square sense. RSM waveforms are used to eliminate some of the assumed additive noise while not attempting to remove it all through averaging. We also use RSM waveforms in an effort to compare results obtained here with those of the ideal bandpass filtering discussed above. We again estimate the evoked potential for each set of recordings for each stimulus by the overall mean waveform. We further assume that the cortical recordings consist of the deterministic evoked response plus additive zero-mean wide-sense stationary noise, which is unaffected by the presence of the evoked response, and which does not affect the evoked response. A single filter is constructed for each of the four stimulus levels used, and a single set of 500 independently obtained noise recordings are used as the noise estimates for each of the four filters.

The traditional Weiner filter was originally designed to address such a problem, although its construction assumed that the desired signal (the evoked response in our case) was a stochastic process [4, 6, 11, 16, 21, 25]. This distinction motivated the use of our name, 'Weiner-type' filter, and was the basis for the derivation of our filter form which is similar in appearance to the traditional noncausal Weiner filter.

5.2 General Form of Weiner-Type Filter

The general form of the noncausal Weiner-type filter frequency response, which is derived in Appendix C is given by the following:

$$H_w(e^{j\Omega}) = \frac{|S(e^{j\Omega})|^2}{|S(e^{j\Omega})|^2 + TS_{vv}(\Omega)}$$

where

$|S(e^{j\Omega})|^2$ = energy spectral density of the evoked response $s[n]$

$S_{vv}(\Omega)$ = power spectral density of the noise $v[n]$ ($v[n]$ replacing $v_i[n]$ generically)

T = time factor reflecting the time duration of the deterministic signal plus some time which accounts for the nonzero convolution product between the filter and the signal beyond the range where the signal is nonzero

5.2.1 Estimation of the Energy Spectral Density $|S(e^{j\Omega})|^2$

The energy spectral density of the deterministic evoked potential is estimated by the magnitude squared of the 4096 point Discrete Fourier Transform (DFT) of the overall mean waveform for each of the stimulus levels used. This length DFT was chosen so as to give good frequency resolution, while not taking an excessively long time to compute, allowing the method to potentially be used in a real time analysis at a later date.

5.2.2 Estimation of the Power Spectral Density $S_{vv}(\Omega)$

The power spectral density of the random process noise signal is estimated by making use of 500 independently obtained noise recordings when no stimulus was present. Estimation of power spectral density is a very rich subject, usually divided

between nonparametric methods and parametric methods. Nonparametric methods such as the periodogram, Bartlett method, Blackman-Tukey and minimum variance estimators require no assumptions about the data except that it is wide-sense stationary [12, 21]. Parametric methods such as the autoregressive, moving average and autoregressive moving average are based on a rational transfer function formed from a model of the data, and require additional assumptions about the data [12]. For this reason parametric methods are not used here.

Among the different nonparamtric (and parametric) methods exists a classical trade-off between frequency resolution and estimator variance [12, 21]. On one extreme is the periodogram method which has good frequency resolution but a very large variance. This method uses the data to form a consistent unbiased estimator of the true autocorrelation. This is accomplished by using the average of the sample autocorrelation functions (defined below) formed from individual waveforms, and the Fourier transform of the average sample autocorrelation to yield the power spectral density estimate. Although it is an asymptotically unbiased estimate of the power spectral density as the record length of individual waveforms increases, its variance is very large, on the order of the power spectral density squared [12, 21]. On the other extreme is Bartlett's procedure which has poor frequency resolution, but which has lower estimator variance. This estimator is found by finding the magnitude squared of the Fourier transform of each individual waveform divided by the length of the waveform, and then averaging these terms to yield the power spectral density estimate. Again this is an asymptotically unbiased estimator, and the variance is reduced from the periodogram method by a factor equal to the number of individual waveforms. Although other nonparametric methods exist such as windowing the individual sample autocorrelation's in the periodogram method by a Hamming window, they in general represent worse frequency resolution than the periodogram method or greater estimator variance then the Bartlett method, and are not included here.

Consistent with our assumption of zero-mean noise, the mean value across waveforms for each time index n was subtracted off from each of the noise recordings. The form of the sample autocorrelation estimator for individual waveforms is as follows:

$$\hat{R}_{vv}[k] = \frac{1}{N-|k|} \sum_{n=0}^{N-|k|-1} v[n]\, v[n-|k|]$$

where

N = length of the noise waveforms (1000 samples)
k = lag variable

An example of the above discussion can be seen for the 500 noise waveforms in the lowest frequency range of interest in Figure 5.1 below, which shows each of the two estimation results discussed above. Given the excessively large variance in the power spectral density estimate using the periodogram method, we use the Bartlett method in the formation of all Weiner-type filters.

Figure 5.1 Power Spectral Density Estimates of the Noise Waveforms (0-20 Hz)

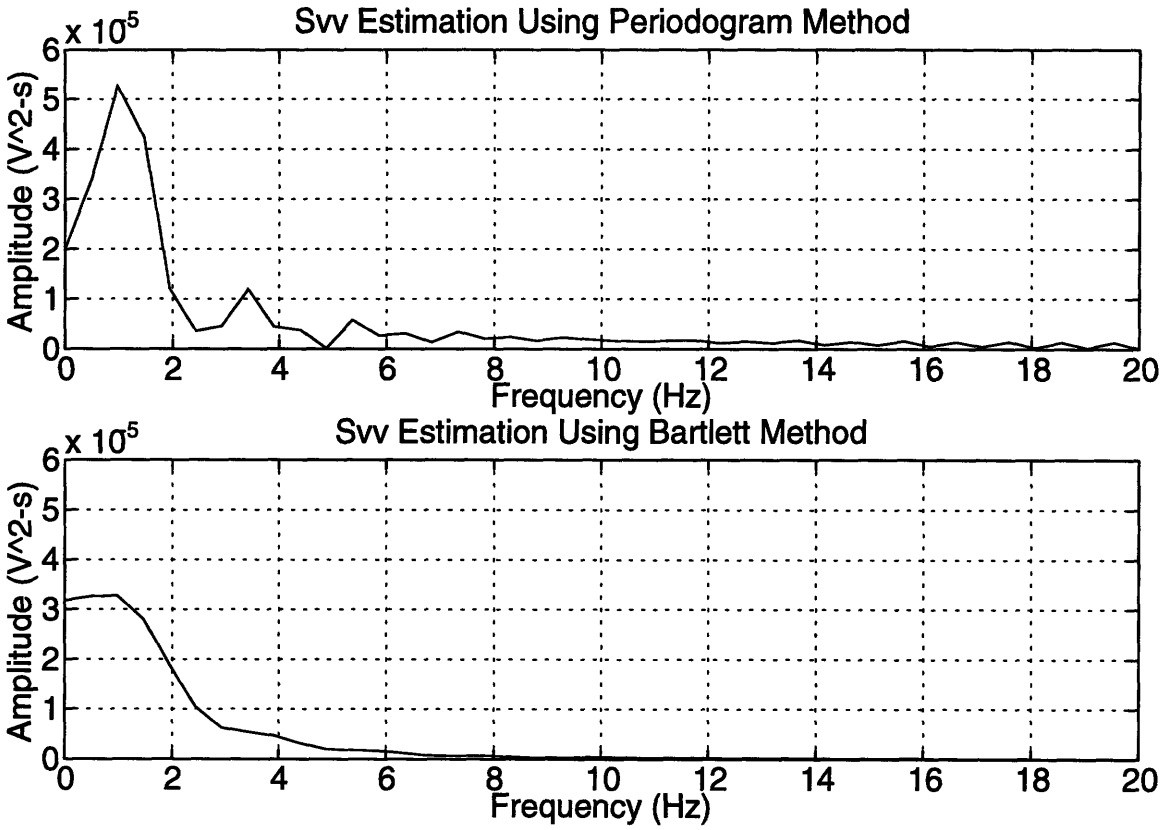

5.2.3 Determination of Time Scale Factor T

The time scale factor T reflects the time duration of the deterministic signal plus some additional time which accounts for the nonzero convolution product between the filter impulse response and the deterministic signal beyond the range where the deterministic signal is nonzero. The value of T should have a duration of no less than 500 ms (the duration of the deterministic signal), and no more than a small fraction of 500 ms longer, under the assumption that the impulse response is negligibly small outside this range. To determine the effects of various values of T on the filter's magnitude response, T was varied between 500, 550, 600 and 650 ms. The resulting Weiner-type filter magnitude responses were examined, and revealed no appreciable differences. For this reason T was chosen as 550 ms, leaving a 10% addition to the duration of the deterministic signal.

5.2.4 Weiner-Type Filter Magnitude Responses

The Weiner-type filter magnitude responses using the Bartlett method for the noise power spectral density estimate are shown below in Figure 5.2 over the frequency range of 0-100 Hz. At larger frequencies the magnitude responses are essentially allpass. Note the notch filter like character of each magnitude response at very low frequencies less than 1 Hz, and at about 15-20 Hz.

Figure 5.2 Magnitude Responses of Weiner-Type Filters (0-100 Hz.)

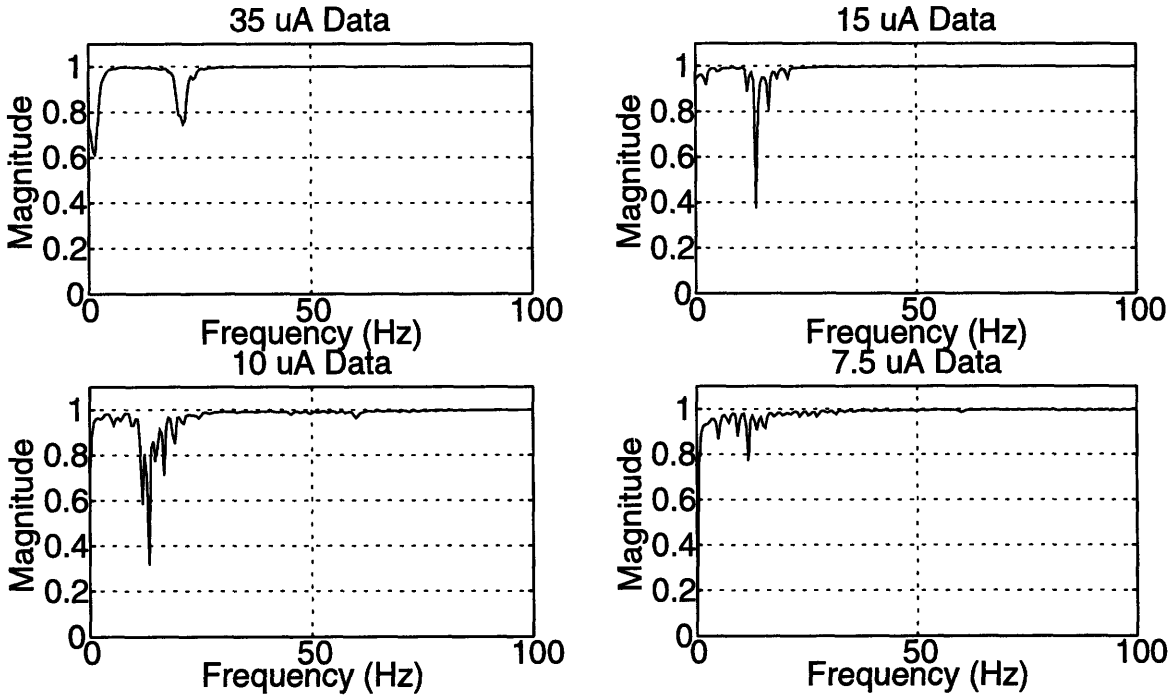

5.3 Results of Weiner-Type Filtering

In order to facilitate comparison to the results of the Ideal Bandpass filtering section, the same experimental conditions were repeated here. Specifically, we used $N_w = 40$ number of RSM waveforms formed from either 100, 50, or 20 individual waveforms as the input to the Weiner-type filter. From the output waveforms, the peak-to-trough and RSM amplitude about the time average features were calculated using a window from 5 to 50 ms, and the percent spread error criteria was calculated, again defined as the ratio of two sample standard deviations to the mean feature size. These results are summarized in Table 5.1 below, which compares the results to those obtained in the ideal bandpass section.

Table 5.1 indicates that for the 35, 15 and 10 μA data sets, the Weiner-type filter produced an output which had a modestly lower percent spread than either no filtering or the best type of ideal bandpass filtering (the ideal filter with the best choice of upper and lower cutoff frequencies for yielding the lowest percent spread), with the singular exception of when $N_i = 20$ individual waveforms per RSM waveform for the RMS feature for the 10 μA data. For the smallest stimulus level of 7.5 μA the results where mixed, although in most cases the Weiner-type filter produces lower percent spread than no filtering at all.

Table 5.1a Highlight of Filtering Results for 35 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	523	31.28	12.0	108	6.23	11.5
	Best Ideal Filter	522	29.60	11.4	108	5.78	10.7
	Weiner-Type Filter	510	23.68	9.3	106	4.63	8.7
Ni=50	No Filter	526	40.91	15.6	109	8.00	14.7
	Best Ideal Filter	524	39.95	15.2	109	7.50	13.7
	Weiner-Type Filter	518	33.71	13.0	106	7.14	13.4
Ni=20	No Filter	527	56.11	21.3	111	13.14	23.8
	Best Ideal Filter	526	55.48	21.1	107	11.67	21.9
	Weiner-Type Filter	522	51.54	19.8	108	9.66	17.9

Table 5.1b Highlight of Filtering Results for 15 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	369	27.68	15.0	79	5.40	13.7
	Best Ideal Filter	372	26.82	14.4	78	5.03	12.9
	Weiner-Type Filter	350	22.79	13.0	75	4.02	10.7
Ni=50	No Filter	373	43.23	23.2	80	8.73	21.8
	Best Ideal Filter	370	42.76	23.1	79	8.25	20.9
	Weiner-Type Filter	346	34.31	19.8	74	7.21	19.4
Ni=20	No Filter	385	50.13	26.0	81	11.22	27.6
	Best Ideal Filter	382	50.04	26.2	81	11.19	27.5
	Weiner-Type Filter	354	44.42	25.1	76	9.35	24.5

Table 5.1c Highlight of Filtering Results for 10μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	188	15.26	16.3	41	3.95	19.5
	Best Ideal Filter	189	14.58	15.4	39	3.60	18.2
	Weiner-Type Filter	166	10.34	12.5	37	2.77	15.1
Ni=50	No Filter	190	20.24	21.3	42	5.41	25.9
	Best Ideal Filter	192	19.90	20.8	41	4.92	24.3
	Weiner-Type Filter	173	16.6	19.2	37	4.23	22.8
Ni=20	No Filter	207	30.31	29.3	44	6.70	30.7
	Best Ideal Filter	208	28.14	27.1	44	6.68	30.7
	Weiner-Type Filter	183	23.52	25.8	38	6.08	31.7

Table 5.1d Highlight of Filtering Results for 7.5 μA Data

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	108	11.12	20.6	21	1.84	17.2
	Best Ideal Filter	89	8.36	18.8	20	1.69	16.5
	Weiner-Type Filter	100	8.46	16.9	20	1.51	15.5
Ni=50	No Filter	115	17.34	30.3	23	3.28	28.8
	Best Ideal Filter	102	11.25	22.0	22	2.62	24.3
	Weiner-Type Filter	105	14.16	26.8	21	2.85	27.0
Ni=20	No Filter	123	26.33	42.7	26	6.14	47.8
	Best Ideal Filter	104	20.88	40.0	24	4.67	38.9
	Weiner-Type Filter	112	25.25	44.8	23	4.33	37.5

6. System Modelling and Parameter Estimation

6.1 Signal-Plus-Additive Noise Model

6.1.1 Statement of Assumptions and General Form of Model

This model assumes that the underlying desired evoked response is a deterministic waveform, and all noise sources can be modelled by a single additive zero-mean wide-sense stationary stochastic process. It is also assumed that the presence of the additive noise does not affect the evoked response, and that it is not affected by the evoked response. Again we estimate the evoked response with the overall mean waveform assuming a sufficiently large number of individual waveforms are used to create it. Mathematically this model takes the following form:

$$x_i[n] = s[n] + v_i[n]$$

where

$x_i[n]$ = individual recorded waveform, i = 1, 2,..., I, n = 1, 2,..., N

$s[n]$ = deterministic unknown evoked response, n = 1, 2,..., N

$v_i[n]$ = WSS random process additive noise associated with $x_i[n]$, assumed unaffected by the presence of $s[n]$, i = 1, 2,..., I, n = 1, 2,..., N

6.1.2 Model Justification

This model was constructed based on observations of thousands of evoked potential waveforms using numerous stimulus levels and subjects. For a given stimulus level it was observed that as one calculates numerous mean waveforms from multiple disjoint sets of individual waveforms, there appears to be an underlying signal which has similar shape and size in each of the mean waveforms. Consequently the "size" of the waveform as defined by the feature size operations (peak-to-trough amplitude and RSM amplitude about the time average) is roughly constant across these different mean waveforms. The observed small variation in feature sizes can be attributed either to (a) additive noise which is still present despite the assumed noise reduction from the process of creating the mean waveforms, or (b) an alternative model which assumes that the underlying evoked response is amplitude scaled by a random variable (discussed below). Finally, since the noise waveforms are assumed to be unaffected by the presence of the underlying signal, this model allows us to utilize noise statistics from independent experiments when the stimulus is absent.

6.1.3 Model Limitations

First this model may be attempting to define an inherently random parameter (feature size) in deterministic terms. The implication of this is to associate the variation in estimates as being due to background noise, instead of random variation in the experimental values of a random variable. An additional implication is that the interrelationships between random variables such as the cross-correlation or cross-covariance may be overlooked. Second this model fails to account for the possible existence of a random variable scaling factor which may multiply the underlying evoked response (discussed below). The mean waveforms created may be averaging this amplitude scale factor making it approach its expected value. Third, since the noise is not fully characterized by a probability distri-

bution function (pdf) no pdf is formed for the observed signals which prevents developing theoretical estimator performance measures such as the Cramer-Rao lower bound on unbiased estimator variance [11]. Finally, the noise may in fact not be a wide-sense stationary random process, but for reasons of mathematical tractability the assumption is made.

6.2 Amplitude Scaled Signal-Plus-Additive Noise Model

6.2.1 Statement of Assumptions and General Form of Model

This model assumes that the underlying evoked response is amplitude scaled by a scalar, positive, independent, identically distributed (iid) random variable, and that all noise sources can be modelled by a single additive zero-mean wide-sense stationary stochastic process. Assuming that the amplitude scale factor is an iid random variable reduces the complexity of attempting to construct its probability density function. It is also assumed that the amplitude scale factor and additive noise are independent, and that the presence of the additive noise does not affect the evoked response, and that it is not affected by the evoked response. Finally it is assumed that the expected value of the amplitude scale factor is unity. Although this choice appears arbitrary, it is done so that the overall mean waveform can be used as an unbiased estimator of the underlying evoked response when a sufficiently large number of individual waveforms are used. Mathematically this model takes the following form:

$$x_i[n] = a_i\, s[n] + v_i[n]$$

where

$x_i[n]$ = individual recorded waveform, $i = 1, 2,..., I$, $n = 1, 2,..., N$

$s[n]$ = deterministic unknown evoked response, $n = 1, 2,..., N$

$v_i[n]$ = random process additive noise associated with $x_i[n]$, assumed unaffected by the presence of $s[n]$, $i = 1, 2,..., I$, $n = 1, 2,..., N$

a_i = random variable amplitude scale factor, $i = 1, 2,..., I$

6.2.2 Model Justification

The basic justification for this model is based on the same observations of waveforms which justified the signal-plus-noise model discussed above. The most important distinction is the supposition that the variations in feature amplitudes of mean waveforms are attributed to the presence of an amplitude scale factor, where in the signal-plus-noise model they are attributed only to additive noise. The motivation for this model came from a preliminary threshold experiment, where individual waveforms with large peak-to-trough or RMS feature sizes above some threshold were removed. From this subset a mean waveform was found, and compared to a mean waveform formed from the same number of individual waveforms chosen from the collection of waveforms whose feature sizes where below the threshold. Comparing the two mean waveforms revealed a similar underlying shape, although the feature sizes of the mean waveforms where different. This suggests that some underlying waveform (the evoked response) was amplified by a random variable scale factor.

6.2.3 Model Limitations

Many of the same limitations exist with this model as with the signal-plus-noise model discussed above. First this model may be attempting to define an inherently random parameter (feature size) in deterministic terms. Second, since the noise is not fully characterized by a pdf, no pdf is formed for the observed signals which prevents developing theoretical estimator performance measures such as the Cramer-Rao lower bound on unbiased estimator variance. Third, the noise may in fact not be a wide-sense stationary random process. Fourth, the noise and the amplitude scale factor may not be independent, and the interrelationships between them are unlikely to be found given our inability to obtain noise recordings when the evoked response is present.

6.3 Verification of the Amplitude Scaled Signal-Plus-Noise Model

6.3.1 General Form of Test

We seek some quantitative justification for the proposed amplitude scaled signal-plus-noise model. Such a justification can be made using principles of linear algebra, orthogonalization and Principle Components analysis. Since we are not interested in any DC (time average) shifting, the following test is applied to the data when the time average of the individual waveforms is first removed. Rewriting the form of the amplitude scaled model in vector notation, and defining the zero-mean random processes vectors $\mathbf{w}$ and $\mathbf{y}$, with experimental values $\mathbf{w}_i$ and $\mathbf{y}_i$ for $i = 1, 2, \ldots, I$, and the normalized signal vector $\mathbf{s}_n$ we obtain:

$$\mathbf{x}_i = a_i\mathbf{s} + \mathbf{v}_i \text{ (Amplitude Scaled Model)}$$

$$\mathbf{w}_i \equiv \mathbf{x}_i - \mathbf{s} \text{ (Zero-mean Data)}$$

$$\mathbf{y}_i \equiv \text{proj}_{\mathbf{s}_n}\mathbf{w}_i = \mathbf{s}_n\mathbf{s}_n^T\mathbf{w}_i \qquad \mathbf{s}_n \equiv \frac{\mathbf{s}}{\|\mathbf{s}\|} \text{ (unit length)}$$

where

$$\mathbf{x}_i = [x_i[1]\ x_i[2]\ \ldots\ x_i[N]]^T \ (i = 1, 2, \ldots, I)$$

$$\mathbf{s} = [s[1]\ s[2]\ \ldots\ s[N]]^T$$

$$\mathbf{v}_i = [v_i[1]\ v_i[2]\ \ldots\ v_i[N]]^T \ (i = 1, 2, \ldots, I)$$

$$\mathbf{w}_i = [w_i[1]\ w_i[2]\ \ldots\ w_i[N]]^T \ (i = 1, 2, \ldots, I)$$

$$\mathbf{y}_i = [y_i[1]\ y_i[2]\ \ldots\ y_i[N]]^T \ (i = 1, 2, \ldots, I)$$

The covariance matrices for $\mathbf{w}$ and $\mathbf{y}$ are obtained as follows:

$$\Lambda_{\mathbf{ww}} \equiv E\{(\mathbf{w}-\mu_{\mathbf{w}})(\mathbf{w}-\mu_{\mathbf{w}})^T\} = E\{\mathbf{ww}^T\}$$

$$\Lambda_{\mathbf{ww}} = E\{(a-1)^2\}\mathbf{ss}^T + E\{\mathbf{vv}^T\} = \sigma_a^2\mathbf{ss}^T + \Lambda_{\mathbf{vv}}$$

$$\Lambda_{\mathbf{yy}} \equiv E\{(\mathbf{y}-\mu_{\mathbf{y}})(\mathbf{y}-\mu_{\mathbf{y}})^T\} = E\{\mathbf{yy}^T\} = \mathbf{s}_n\mathbf{s}_n^T\Lambda_{\mathbf{ww}}\mathbf{s}_n\mathbf{s}_n^T$$

where

$$\mu_w = [\mu_w[1]\ \mu_w[2]\ \ldots\ \mu_w[N]]^T = \mathbf{0}$$

$$\mu_y = [\mu_y[1]\ \mu_y[2]\ \ldots\ \mu_y[N]]^T = \mathbf{0}$$

Finally we find the expected value of the norm squared of $\mathbf{y}$ in terms of the covariance matrix of $\mathbf{w}$, and the eigenvalues and eigenvectors of $\mathbf{w}$ as follows:

$$E\{\mathbf{y}^T\mathbf{y}\} = E\{\left(\mathbf{s}_n\mathbf{s}_n^T\mathbf{w}_i\right)^T\left(\mathbf{s}_n\mathbf{s}_n^T\mathbf{w}_i\right)\} = E\{\mathbf{w}_i^T\mathbf{s}_n\mathbf{s}_n^T\mathbf{s}_n\mathbf{s}_n^T\mathbf{w}_i\}$$

$$E\{\mathbf{y}^T\mathbf{y}\} = E\{\mathbf{w}^T\mathbf{s}_n\mathbf{s}_n^T\mathbf{w}\} = E\{\mathbf{s}_n^T\mathbf{w}\mathbf{w}^T\mathbf{s}_n\} = \mathbf{s}_n^T\Lambda_{ww}\mathbf{s}_n$$

Using Mercer's Theorem [23] we can express Λ_{ww} in terms of its eigenvalues and eigenvectors yielding:

$$E\{\mathbf{y}^T\mathbf{y}\} = \mathbf{s}_n^T\left(\sum_{i=1}^{N}\lambda_i\vartheta_i\vartheta_i^T\right)\mathbf{s}_n = \sum_{i=1}^{N}\lambda_i\left(\mathbf{s}_n^T\vartheta_i\right)^2 \text{ (Quadratic Form)}$$

where

$\lambda_i = i^{th}$ eigenvalue of Λ_{ww}

$\vartheta_i = i^{th}$ eigenvector of Λ_{ww} corresponding to λ_i

We now let $\lambda_i \geq \lambda_2 \geq \ldots \geq \lambda_N$, and since $\left(\mathbf{s}_n^T\vartheta_i\right)^2 \leq 1$ for all i, we obtain our the inequality:

$$E\{\mathbf{y}^T\mathbf{y}\} \leq \text{trace}\{\Lambda_{ww}\} \qquad \text{where trace}\{\Lambda_{ww}\} = \sum_{i=1}^{N}\lambda_i$$

Equivalently

$$R \equiv \frac{E\{\mathbf{y}^T\mathbf{y}\}}{\text{trace}\{\Lambda_{ww}\}} \leq 1$$

We can express R in terms of an additional expression which leads further insight into our choice of this test. Defining the eigenvalues Υ_i and eigenvectors ψ_i of the covariance matrix Λ_{vv} of the zero-mean additive noise random process $\mathbf{v}$, and using the quadratic form for the expansion of $\mathbf{s}_n^T\Lambda_{vv}\mathbf{s}_n$ we obtain:

$$E\{\mathbf{y}^T\mathbf{y}\} = \mathbf{s}_n^T\Lambda_{ww}\mathbf{s}_n = \mathbf{s}_n^T\left(\sigma_a^2\mathbf{s}\mathbf{s}^T + \Lambda_{vv}\right)\mathbf{s}_n = \sigma_a^2\|\mathbf{s}\|^2 + \sum_{i=1}^{N}\gamma_i\left(\mathbf{s}_n^T\psi_i\right)^2$$

$$\text{trace}\{\Lambda_{ww}\} = \sigma_a^2\text{trace}\{\mathbf{s}\mathbf{s}^T\} + \text{trace}\{\Lambda_{vv}\} = \sigma_a^2\|\mathbf{s}\|^2 + \sum_{n=1}^{N}\gamma_i.$$

Thus

$$R \equiv \frac{E\{\mathbf{y}^T\mathbf{y}\}}{\text{trace}\{\Lambda_{\mathbf{ww}}\}} = \frac{\sum_{i=1}^{N} \lambda_i \left(\mathbf{s}_n^T \vartheta_i\right)^2}{\sum_{i=1}^{N} \lambda_i} = \frac{\sigma_a^2 \|\mathbf{s}\|^2 + \sum_{i=1}^{N} \gamma_i \left(\mathbf{s}_n^T \psi_i\right)^2}{\sigma_a^2 \|\mathbf{s}\|^2 + \sum_{n=1}^{N} \gamma_i} \leq 1$$

From the above expression R = 1 if an only if all the eigenvalues of $\Lambda_{\mathbf{ww}}$ except the largest (the dominant eigenvalue) are equal to zero, and if $\mathbf{s}_n = \vartheta_i$. This provides confirmation of the model because as σ_a^2 increases, individual $\mathbf{w}_i$ tend to look more like $\mathbf{s}$, making the dominant eigenvector of $\Lambda_{\mathbf{ww}}$ look more like $\mathbf{s}$, making the numerator of the expression for R above approach the dominant eigenvalue λ_1. Furthermore the non-dominant eigenvalues tend to zero in this case because the covariance matrix $\Lambda_{\mathbf{ww}}$ can be explained mostly by the term $\lambda_1 \vartheta_1 \vartheta_1^T$ from Mercer's Theorem.

However, note the presence of a degenerate case which causes R = 1 when the model is not confirmed. This occurs when the non-dominant eigenvalues of $\Lambda_{\mathbf{vv}}$ are all equal to zero and the signal is in the direction of the dominant eigenvector (i.e. $\mathbf{s}_n = \psi_1$). We expect the degenerate case to be unlikely given our assumption that the noise is due to many additive sources.

6.3.2 Application of the Model Verification Test to the Four Data Sets

We have seen from the above discussion that the test for whether we support the amplitude scaled model depends on how close the ratio R is to unity. We also recognize that the underlying evoked potential signal $\mathbf{s}$ (and its amplitude scaled version) are likely to only occur during a finite portion of the data record, presumably during the temporally early portions of the individual waveforms, and that this time period should be roughly the same across experiments (although it is certainly possible that larger amplitude stimuli may produce an evoked response which occurs for a longer duration). Based on these assumptions, we seek confirmation for the model during windowed segments of the individual waveforms, and expect that confirmation will depend on the starting time and duration of these windowed segments.

Four time windows were examined of duration 20, 30, 50 and 100 ms, and the starting times for these windows were swept from the beginning of the individual waveforms after the stimulus artifact (at 5 ms) in 0.5 ms increments up to the end of the waveform (at 500 ms - window duration). The ratio R was calculated for the four data sets in each of these time windows, and these individual values of R were averaged across data sets for each particular time window and starting point. The window starting time when the maximum of this average R occurred was used to indicate which time window corresponds to maximal confirmation of the model across data sets. These results are summarized in Figure 6.1 below.

Figure 6.1 Average Value of R across Data Sets for Numerous Time Windows

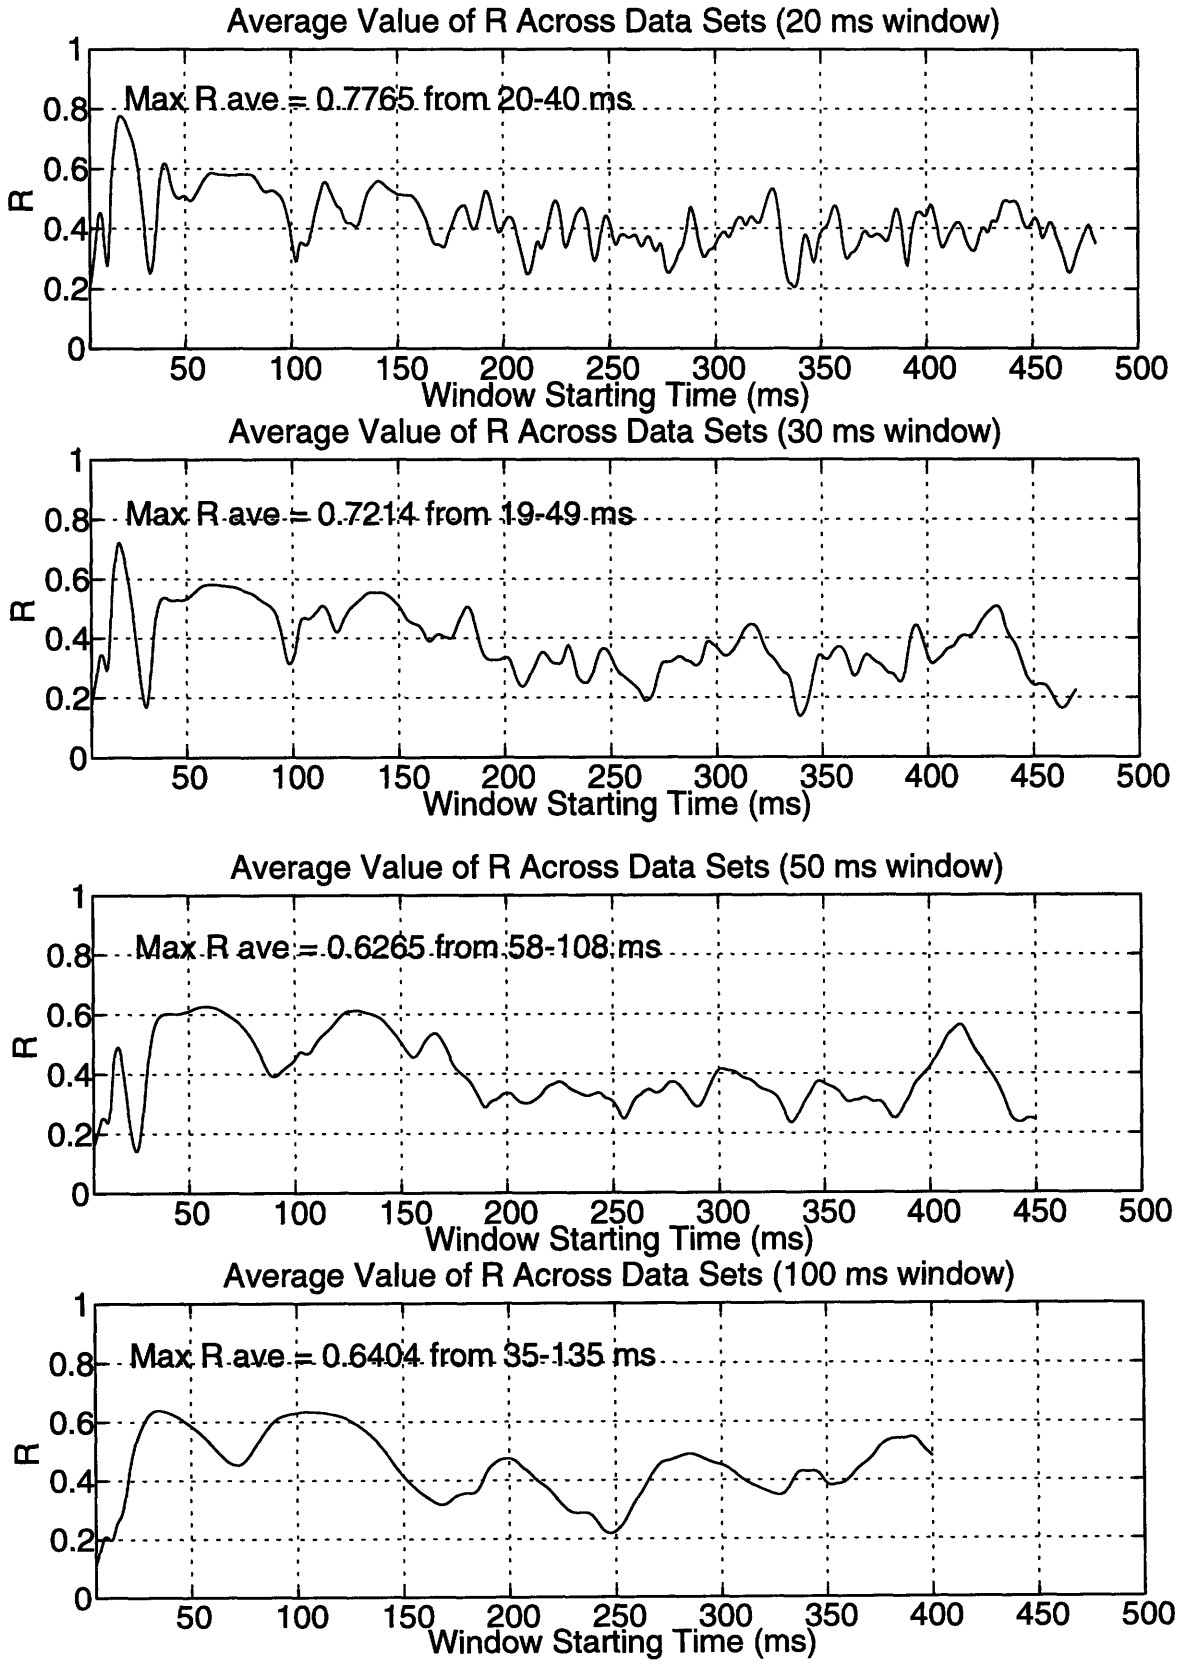

The plots in Figure 6.1 indicate a strong confirmation of the model when the time window is limited to 20 or 30 ms, occurring during the early portions of the waveform, although not in the earliest portion. This could be accounted for by a time dependent variance of the amplitude scale factor such that its variance is small during the earliest portion of the waveforms, then growing, and then dying out for later portions when the signal **s** may itself be dying out. Although the signal may be dying out after about 100 ms, we do not expect a significant loss of the signal **s** to occur before 50 ms, making the above explanation incomplete in attempting to account for the strong dip in R before 50 ms in the 20 and 30 ms time window data. It is also interesting that the lowest value of R occurs just following the highest confirmation for the 20 and 30 ms time windows. The individual values of R for the four data sets is summarized in Table 6.1 below for the time window when the average value of R across data set is a maximum.

Table 6.1 Values of R for Each Data Set When Average R across Data Sets is Maximum

	35 μA Data	15 μA Data	10 μA Data	7.5 μA Data
20 ms Window (20-40 ms)	0.7910	0.8243	0.7716	0.7190
30 ms Window (19-49 ms)	0.7668	0.8081	0.7317	0.5791
50 ms Window (58-108 ms)	0.6401	0.6401	0.6242	0.6016
100 ms Window (35-135 ms)	0.6181	0.6378	0.6594	0.6462

Further confirmation of the model can be made qualitatively by examining the eigenvector corresponding to the largest eigenvalue (the dominant eigenvector/eigenvalue pair) when the ratio R is large. We expect that this dominant eigenvector should resemble the signal **s** when properly scaled. Using simple principles of orthogonality, the best mean square fit between the signal estimate for **s** (the overall mean waveform) and a scaled dominant unit length eigenvector occurs when the scaling equals the dot product of **s** and the dominant eigenvector. Figure 6.2 below shows the comparison of the properly scaled dominant eigenvector and the overall mean waveform (made zero-mean over the window in question) for the 20 ms time window for each of the four data sets when the average ratio R across data sets was a maximum (between 20 and 40 ms). This figure provides further confirmation of the amplitude scaled signal model for the 20 ms time window over the time range given.

Figure 6.2 Signal Estimate s and Dominant Eigenvectors (20 ms Time Window)

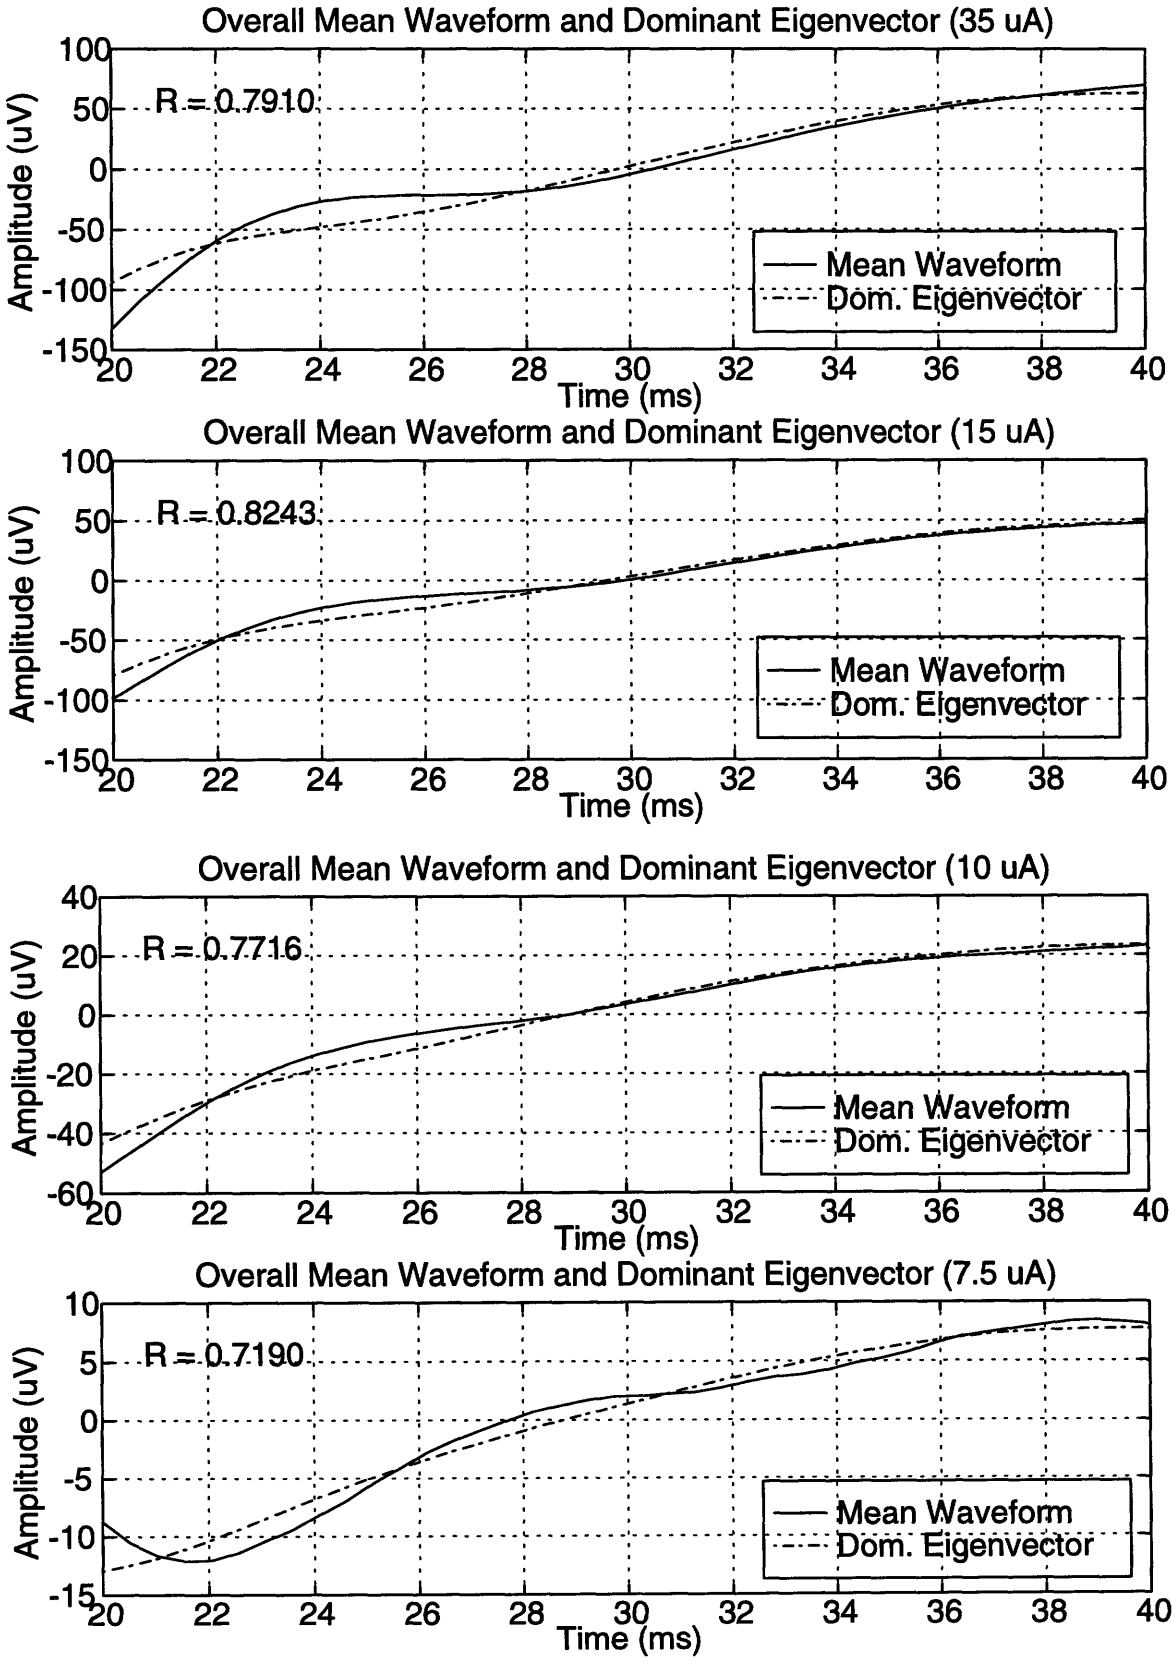

6.4 Estimation of the Amplitude Scale Factor

6.4.1 Statement of Goals

Given that the amplitude scaled model has been verified over a time window of limited duration, we seek an estimate of the amplitude scale factor of individual waveforms $\hat{a}_i$. Although the model is not completely confirmed over the waveform duration of 500 ms, we assume that if present, the amplitude scale factor multiplies the signal **s** over its duration, which may or may not be the same as the duration of the individual recorded waveforms. We thus seek to estimate the amplitude scale factor for individual waveforms using some finite portion of the waveforms, starting at 5 ms (when the stimulus artifact is over) up to some finite time less than or equal to 500 ms. Again, since we are not interested in DC shifting effects, the time average of individual waveforms over the time window in question is first subtracted off. We seek to (a) determine a rough estimate of the distribution of the amplitude scale factor, (b) determine if there is any temporal dependence for the amplitude scale factor across waveforms, and (c) to determine the reasonableness of our estimates in terms of our underlying assumptions which are enumerated below.

6.4.2 Estimator General Form and Underlying Assumptions

The estimation procedure for the amplitude scale factor for each of the individual waveforms is based on a matched filter approach [11]. We assume here that the noise does not appreciably resemble the signal and that the amplitude scale factor is a positive constant whose expected value is unity. This assumption makes unbiased our estimate of the evoked potential using the overall mean waveform. The general approach is to compare the individual waveforms to a matched filter for the evoked potential estimate. The overall mean waveform is scaled by a positive constant and compared to the individual waveform whose amplitude scale factor we wish to estimate. The positive scale factor multiplying the overall mean waveform is adjusted until the fit with the individual waveform is optimized.

An equivalent procedure to the above comparison between the amplitude scaled overall mean waveform and the individual waveforms is accomplished by minimizing the Euclidean norm squared, defined as the sum of squared differences between the two waveforms at each moment in time. Appendix B shows the derivation of this estimator, whose general form is as follows:

$$\hat{a}_i = \frac{\sum_n x_i[n]\, s[n]}{\sum_n s^2[n]}$$

where

$\hat{a}_i$ = estimate of the amplitude scale factor, i = 1, 2,..., I

$x_i[n]$ = individual recorded waveform, i = 1, 2,..., I, n = 1, 2,..., N

s [n] = deterministic unknown evoked response, n = 1, 2,..., N, estimated by the overall mean waveform

In practice, some initial signal processing of individual waveforms is performed to remove any bias due to the DC values of the individual waveforms. This processing is simply the removal of the DC component of the individual waveforms. The reason for this initial processing can be seen by examination of the amplitude scale estimator form given above. If we do not first remove the DC values of the individual waveforms, then we will effectively be adding a bias equal to the dot product of the DC values and the evoked potential estimate, normalized by the energy of the evoked potential.

6.5 Results of Amplitude Scale Estimation

The amplitude scale factor was estimated using the estimator described above, applied to the 35, 15, 10, and 7.5 μV data using a 100, 200 and 500 ms time window, starting in all cases from the point immediately following the stimulus artifact ending at 5 ms. The mean of each of these estimates is exactly one, given the form of the estimator and the fact that the overall mean waveform is used as the evoked potential estimate. These results are summarized in Figure 6.3a-c below which shows 25 bin histograms of the amplitude scale estimates for each of the stimulus levels and time windows used.

Figure 6.3a Histograms of Amplitude Scale Estimates for 100 ms Time Window

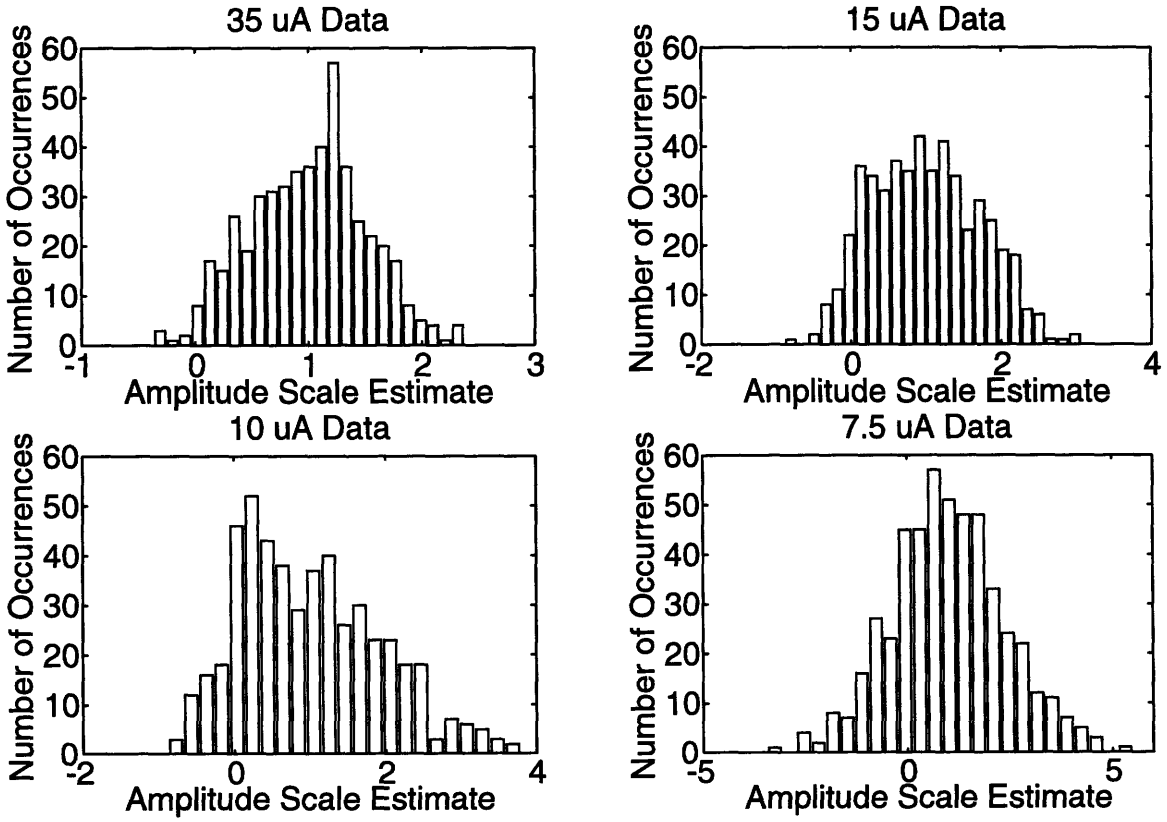

Figure 6.3b Histograms of Amplitude Scale Estimates for 200 ms Time Window

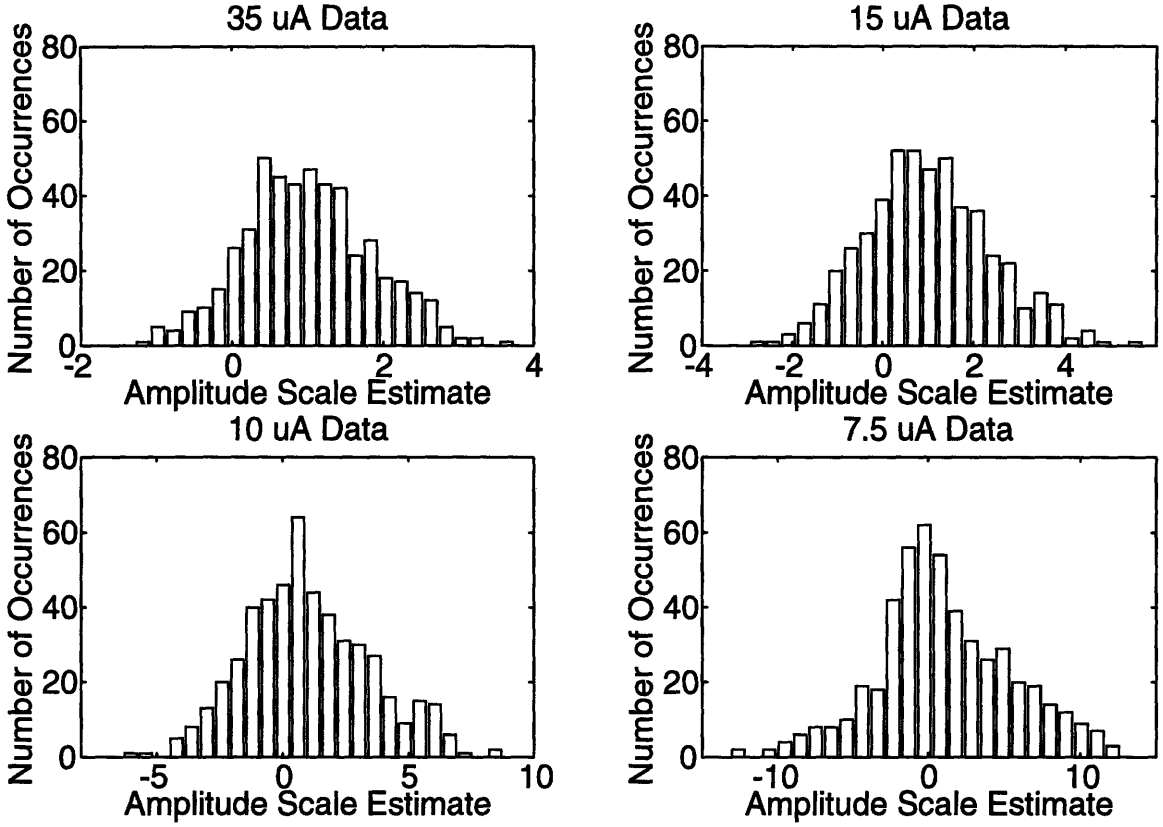

Figure 6.3c Histograms of Amplitude Scale Estimates for 500 ms Time Window

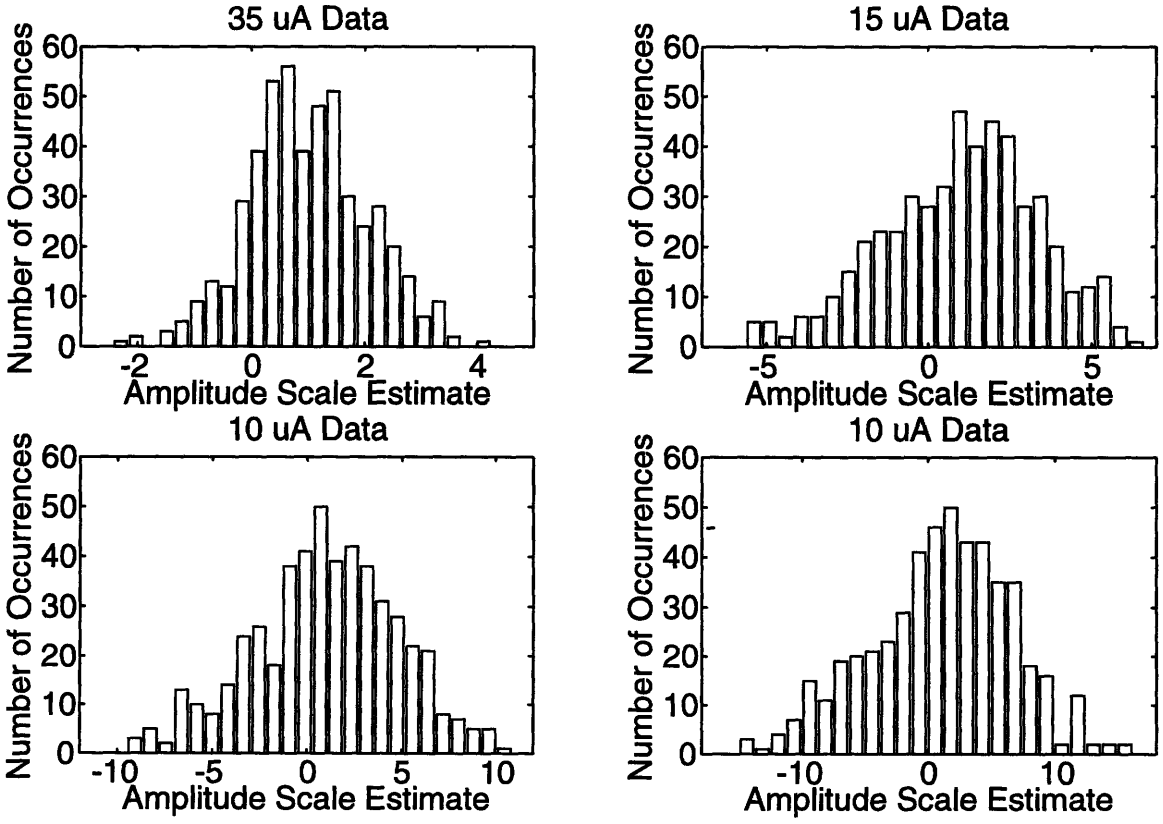

A number of salient features of the above histograms are noteworthy. First, for each time window used, some of the amplitude scale estimates are negative which is inconsistent with our assumption of only positive amplitude scale estimates. This can be explained by the fact that the amplitude scale estimator is actually equal to the sum of the actual amplitude scale factor and the dot product of the additive noise waveform v [n] and the signal estimate normalized by the energy in the signal estimate (see Appendix B). Note also that for a given time window, a larger percentage of amplitude scale estimates are negative as the stimulus amplitude is decreased. This is expected since decreasing stimulus amplitude also reduces evoked response strength, making the additive noise term stronger in comparison. Second, increasing the window time for each stimulus amplitude also increases the percentage of negative amplitude scale estimates. This can also be explained by the additive noise term since we expect that the evoked potential signal is strongest during the temporally early portions of the individual waveforms, making the longer window estimates more subject to additive noise effects. Finally, the spread of estimates is larger for smaller amplitude stimuli for a given window length, and larger for longer window lengths for a given stimulus amplitude. Again this can be explained by the effect of additive noise which is relatively stronger as the stimulus is decreased, or the window is increased. Table 6.2 below summarizes this spread in terms of sample standard deviation of the amplitude scale estimates.

Table 6.2 Sample Standard Deviation of Amplitude Scale Estimate

	35 μA Data	15 μA Data	10 μA Data	7.5 μA Data
100 ms Window	0.5073	0.7174	0.9415	1.3769
200 ms Window	0.8456	1.3774	2.4861	4.4740
500 ms Window	1.0567	2.3859	3.8296	5.6325

In an effort to determine if any periodic time dependence exists, the autocorrelation of each of the amplitude scale estimates as a function of waveform number was calculated. There appears to be no such discernible periodicity given the absence of large peaks away from zero lag. Qualitative observation of the amplitude scale estimates as a function of waveform number reveals no other temporal dependence.

7. Active Noise Cancellation

7.1 Statement of Goals and Assumptions

Active noise cancellation (ANC) is an example of a general class of adaptive filtering methods in which the noise associated with a cortical recording (the primary signal) is reduced by using a secondary noise recording (the reference signal) which is correlated to the noise in the cortical recording [29]. The goal is to reduce the noise in the primary recording to yield a more accurate estimate of the underlying evoked potential using the smallest number of waveforms necessary. We assume that the noise in the primary recording is zero-mean and additive to the deterministic evoked potential which may or may not be amplitude scaled, and that the reference noise is zero-mean and correlated to the noise in the primary recording. The earlier assumption that the additive noise is wide-sense stationary is not necessary here since the adaptive system has the ability to track the changing signal statistics [29]. Note that this method yields estimates of individual evoked potential waveforms without requiring a priori knowledge of the evoked potential or primary noise statistics, unlike the Weiner-type filtering method discussed earlier.

7.2 Active Noise Cancellation Overview

The basic active noise cancellation block diagram is depicted below in Figure 7.1. The primary signal is the sum of the desired evoked potential and the additive noise we wish to remove. The reference signal consists of noise recordings which are assumed to be correlated to the primary noise. The source of the reference signal is discussed separately below. The basic method is to filter the reference signal in such a way that the output of the adaptive filter best resembles the noise in the primary signal, so that the difference between the primary signal and the adaptive filter output is the best estimate of the underlying evoked potential. During each iteration in which a single sample of the primary and reference signals is processed, an error signal (or system output) ε_k is generated which is the difference between the filtered reference signal sample y_k and the primary signal sample p_k (note that we are using vector notation where $p_k = p[k] = x[k]$, the raw recorded individual waveforms). This error signal is then used to update the filter coefficients according to the adaptive algorithm. The adaptive algorithm used here is the Least Mean Squares (LMS) developed by Widrow, which is discussed in more detail below.

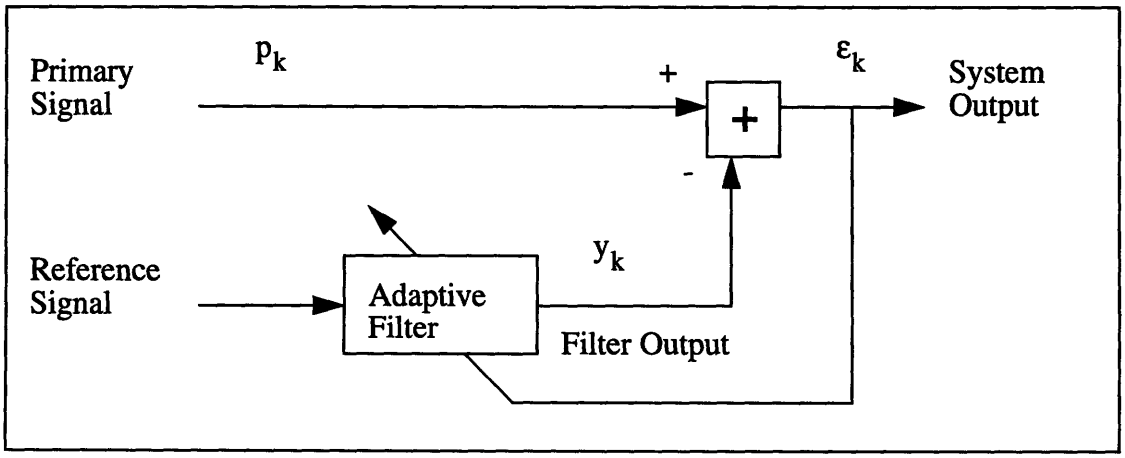

Figure 7.1 Block Diagram of Active Noise Cancellation System

7.2.1 Adaptive Filter of the ANC System and the Performance Surface

The adaptive filter of the active noise cancellation system is depicted below in Figure 7.2. It uses N_f number of filter weights which are updated after each iteration according to the LMS algorithm. The input to $N_f - 1$ of these weights are the last $N_f - 1$ samples of the reference signal, and the input to the final filter weight is merely the constant +1, which has been shown to help eliminate low frequency motion artifact [3]. Using vector notation we define the reference input at time k as $\mathbf{R}_k$ (where the bias of +1 is included for completeness) and the weight vector $\mathbf{W}_k$ as follows:

$$\mathbf{R}_k = [r[k] \; r[k-1] \; \ldots \; r[k-N_f+2] \quad 1]^T$$

$$\mathbf{W}_k = [w_{k0} \; w_{k1} \; \cdots \; w_{kN_f-1}]^T$$

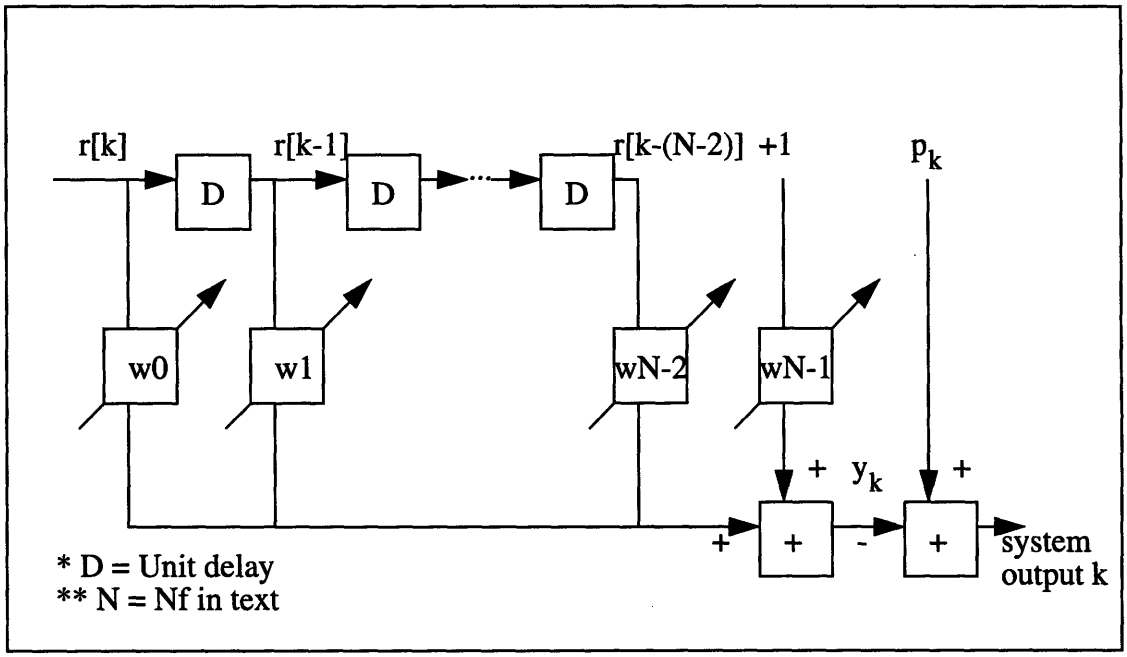

Figure 7.2 Adaptive Filter Portion of the Active Noise Cancellation System

Thus the output of the adaptive filter at time k is $y_k = \mathbf{W}_k^T \mathbf{R}_k$, and the error signal (system output) at time k is $\varepsilon_k = p_k - y_k = p_k - \mathbf{W}_k^T \mathbf{R}_k$. We seek to minimize the mean squared error ξ, which is defined as follows:

$$\xi \equiv E\{\varepsilon_k^2\} = E\{p_k^2 + \mathbf{W}_k^T \mathbf{R}_k \mathbf{R}_k^T \mathbf{W}_k - 2p_k \mathbf{R}_k^T \mathbf{W}_k\}$$

$$\xi = E\{p_k^2\} + \mathbf{W}_k^T E\{\mathbf{R}_k \mathbf{R}_k^T\} \mathbf{W}_k - 2E\{p_k \mathbf{R}_k^T\} \mathbf{W}_k$$

Dropping the k notation, and further defining the reference input correlation matrix $\Lambda_{rr} \equiv E\{\mathbf{R}\mathbf{R}^T\}$ (which is also the reference input covariance matrix since the reference noise is assumed zero-mean) and the vector of cross correlations between the primary

input sample and the reference inputs $\mathbf{P} \equiv E\{p_k \mathbf{R}_k\}$, we rewrite the MSE:

$$\xi = E\{p_k^2\} + \mathbf{W}^T \Lambda_{rr} \mathbf{W} - 2\mathbf{P}^T \mathbf{W}$$

The MSE ξ written above is a quadratic function of the components of the weight vector when the reference input components and primary signal are wide-sense stationary [28]. The goal of the adaptive system is to search this quadratic mean square error surface for the "bottom of the bowl", i.e. the value of the weight vector which yields the minimum mean squared error. With a quadratic performance surface there is only a single global minimum [28].

One method of searching the MSE quadratic performance surface is by gradient methods. Dropping the k notation and defining the gradient ∇ as the partial derivative of the MSE with respect to the weight vector, and the optimal weight vector $\mathbf{W}^*$ as that value of the weight vector when the MSE is minimum we obtain the following:

$$\nabla \equiv \frac{\partial \xi}{\partial \mathbf{W}} = \left[\frac{\partial \xi}{\partial w_o} \; \frac{\partial \xi}{\partial w_i} \cdots \frac{\partial \xi}{\partial w_{N_f - 1}} \right]^T = 2\Lambda_{rr}\mathbf{W} - 2\mathbf{P}$$

Setting the gradient to zero and using $\mathbf{W}^* = \mathbf{W}$ we obtain:

$\nabla = \mathbf{0} = 2\Lambda_{rr}\mathbf{W}^* - 2\mathbf{P}$ which implies

$\mathbf{W}^* = \Lambda_{rr}^{-1}\mathbf{P}$ (assuming Λ_{rr} is nonsingular)

This expression is an equivalent expression to the Weiner-Hopf equation in matrix form [28, 29]. Substituting this result into the expression for the MSE above, and performing algebra we obtain the following:

$$\xi_{min} = E\{p_k^2\} - \mathbf{P}^T \mathbf{W}^*$$

This expression can also be reformulated into a general expression for the MSE under nonminimized conditions [28,29], yielding the following:

$$\xi = \xi_{min} + (\mathbf{W} - \mathbf{W}^*)^T \Lambda_{rr} (\mathbf{W} - \mathbf{W}^*)$$

In general, the adaptive process k+1 iteration filter weights are related to the previous iteration filter weight by the following expression, where ∇_k is the gradient of the MSE surface at the k^{th} iteration, and μ is a convergence constant governing the rate that the weight vector converges to the optimal solution:

$$\mathbf{W}_{k+1} = \mathbf{W}_k - \mu \nabla_k$$

Various methods are used to estimate the gradient ∇_k which include Newton's method and the method of steepest descent [28]. Newton's method yields the optimal weight vec-

tor in a single iteration, but is not usable when the reference input correlation matrix is singular. The LMS algorithm is a special case of the method of steepest descent, and is discussed in the following section.

7.2.2 LMS Algorithm

The LMS algorithm is an example of a steepest descent method for estimating the gradient of the MSE surface. The LMS algorithm uses ε_k^2 as the estimate of the MSE $\xi_k = E\{\varepsilon_k^2\}$. Defining $\hat{\nabla}_k$ as the gradient estimate at the k^{th} iteration we obtain the following:

$$\hat{\nabla}_k \equiv \frac{\partial \xi_k}{\partial \mathbf{W}} = \left[\frac{\partial \varepsilon_k^2}{\partial w_0} \quad \frac{\partial \varepsilon_k^2}{\partial w_1} \cdots \frac{\partial \varepsilon_k^2}{\partial w_{N_f-1}}\right]^T = 2\varepsilon_k \left[\frac{\partial \varepsilon_k}{\partial w_0} \quad \frac{\partial \varepsilon_k}{\partial w_1} \cdots \frac{\partial \varepsilon_k}{\partial w_{N_f-1}}\right]^T$$

Therefore

$$\hat{\nabla}_k = -2\varepsilon_k \mathbf{R}_k .$$

Using the weight vector update equation developed above, and substituting $\hat{\nabla}_k$ for ∇_k we obtain the following:

$$\mathbf{W}_{k+1} = \mathbf{W}_k - \mu\nabla_k = \mathbf{W}_k + 2\mu\varepsilon_k \mathbf{R}_k \text{ (LMS Algorithm)}$$

This algorithm is easy to implement in a real time setting, which was the chief motivation for its use in this thesis. One needs to only specify three parameters, the number of weights used in the adaptive filter, the initial value for these weights, and the convergence parameter μ. We decided to use 24 weights which has been used successfully by a number of others [3] with initial weight values of zero [3]. Although some initial experiments where run with 12 and 36 weights also, the results where not significantly different. The convergence parameter represents a trade-off between convergence rate and stability. Since we desire a fast convergence rate given that the important waveform information is contained in the temporally early portions, we sought the largest convergence parameter without producing unstable results. After trial and error using convergence factors of 0.01, 0.1, 1 and 10, we decided to use $\mu = 0.1$ since the convergence was not too slow (as it was with 0.01) and the output remained stable (which was not true for values of 1 and 10).

7.2.3 Reference Noise Source

Generally, active noise cancellation uses reference waveforms which are simultaneously obtained from some secondary noise source [9, 18, 28, 29]. The secondary noise source must also be chosen so that it contains no appreciable desired signal components, or else these components may be subtracted off in the adaptive process [9, 28, 29]. For our purposes such secondary sources may include the heart, lungs or muscular activity. These sources are generally treated one at a time, and alterations in the adaptive filter structure must be made to accommodate the use of multiple simultaneous recordings.

An alternative novel approach used here is to use a priori noise recordings from the same recording location that the primary signal is later recorded. This allows the multiple reference sources to be combined as one signal from only one recording sight. The chief

disadvantages are (a) that the secondary noise statistics must be the same during the a priori recordings as during the signal recording, and (b) that synchronization of the a priori noise components must be made to match up to the noise present when the signal is recorded. Disadvantage (a) is potentially problematic for biological noise sources since the experimental procedure often stresses an animal in such a way that these noise characteristics do change over time (such as increasing heart rate or breathing rate). Disadvantage (b) can potentially be overcome if the adaptive process is able to track the phase difference in periodic noise components between the reference and primary noise, and counteract them accordingly. Despite these disadvantages, some interesting and potentially beneficial results are obtained by using a priori recordings, which are discussed below.

7.3 Results of Active Noise Cancellation

A single set of 500 a priori noise recordings was obtained prior to the presentation of any of the four stimuli. All five sets of recordings were obtained from the same subject during the same experimental period. These noise recordings form the reference signals for the active noise cancellation system for each of the four recorded data sets. All ANC input waveforms start at 5 ms to avoid the stimulus artifact. Numerous error criteria were developed to evaluate the effectiveness of the active noise cancellation system, and are each discussed separately below.

It is also instrumental to examine a typical recorded waveform used as input to the ANC system, the corresponding output waveform, and the overall mean waveform, as shown in Figure 7.3 below over two time durations. Each of these waveforms has a zero time average over the 500 ms waveform duration. From the 5-500 ms plot, note that the ANC output waveform more closely resembles the overall mean waveform from 100 to 500 ms, not exhibiting the low frequency amplitude movement characterizing the input waveform. Note also from the 5-100 ms plot that the ANC output waveform also is a better match to the overall mean waveform in the early portion of the waveform, not exhibiting the large DC shift characterizing the input waveform. However, in terms of our feature analysis (peak-to-trough and RMS amplitude), the differences between the ANC input and output tend to diminish since the features are independent of DC shifting. Another important feature, although difficult to determine from Figure 7.3, is the short duration from about 5-10 ms where the ANC adaptive filter converges. Examining numerous output waveforms reveals a large variation in amplitudes in this very early region.

Figure 7.3 Typical Input and Output Waveforms and Overall Mean Waveform (35 μA)

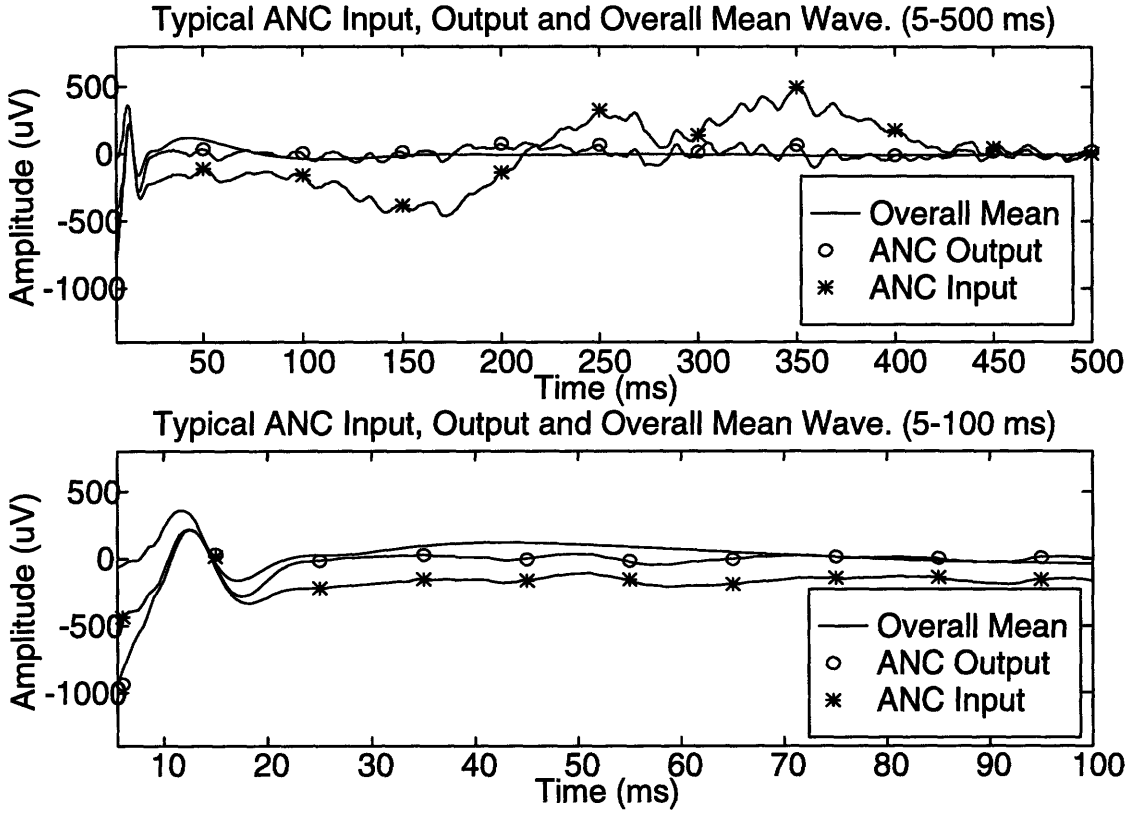

7.3.1 Normalized Root Mean Square Error of Individual ANC Input and Output Waveforms and Scaling of Output Waveforms

The normalized RMS error between each of the ANC output waveforms and the overall mean waveform (the evoked potential signal estimate) was found for each of the four data sets, and compared to the corresponding normalized RMS error of the ANC input waveforms. The normalized RMS error is defined as follows:

$$(\mathrm{RMS})_{\mathrm{Norm}} \equiv \left\{ \frac{\frac{1}{N}\sum_{n=1}^{N} (q[n] - \bar{x}[n])^2}{\frac{1}{N}\sum_{n=1}^{N} \bar{x}^2[n]} \right\}^{1/2}$$

where

$\bar{x}[n]$ = overall mean waveform

$q[n]$ = input to the ANC system ($x[n]$) or the output from the ANC system ($y[n]$)

Note that this error criterion is applied to waveform estimates, as opposed to the features of these waveforms, as was done in the earlier analysis of the feature variance (or percent spread) error criterion. This latter analysis is repeated for the ANC system output waveform in section 7.3.3 below.

To determine if the ANC system reduces the normalized RMS error more effectively over shorter segments of the waveforms, the normalized RMS error was calculated over three time windows from 5-500 ms, 5-100 ms and 5-50 ms labelled as the 500, 100 and 50 ms windows respectively. These results are summarized for the first 100 waveforms in Figure 7.4 below.

Additionally, given the likelihood that an amplitude scale factor is present for individual waveforms (see section 6) the ANC output waveforms were also rescaled by an estimate of the scale factor and by an additive constant such that the mean square error between the rescaled ANC output waveforms and the evoked potential estimate is minimized. The normalized RMS error was then calculated for the rescaled ANC output waveforms. Although this rescaling has no practical application in reducing the number of waveforms needed for an accurate estimate of the evoked potential (since the overall mean waveform is used to calculate the rescaling parameters), the rescaled output normalized RMS error does indicate the effectiveness of the ANC process.

Determining the rescaling parameter was achieved by minimizing the MSE as follows:

$$MSE = \sum_{n=1}^{N} (ay[n] + b - \hat{s}[n])^2$$

where

$y[n]$ = an individual ANC output waveform
$\hat{s}[n]$ = evoked potential estimate (the overall mean waveform)
a = amplitude scale parameter
b = additive scale parameter

$$\frac{\partial}{\partial a}MSE = \sum_{n=1}^{N} 2y[n]\,(ay[n] + b - \hat{s}[n]) = 0$$

$$a\|y[n]\|^2 + bN\overline{y[n]} = \langle y[n], \hat{s}[n] \rangle \quad (*)$$

where $\|y[n]\|^2 \equiv \sum_{n=1}^{N} y^2[n]$ (squared norm)

$$\overline{y[n]} = \frac{1}{N}\sum_{n=1}^{N} y[n] \quad \text{(time average)}$$

$$\langle y[n], \hat{s}[n] \rangle \equiv \sum_{n=1}^{N} y[n]\hat{s}[n] \quad \text{(inner product)}$$

$$\frac{\partial}{\partial b}MSE = \sum_{n=1}^{N} 2(ay[n] + b - \hat{s}[n]) = 0$$

$$aN\overline{y[n]} + Nb = N\overline{\hat{s}[n]}$$

$$b = \overline{\hat{s}[n]} - a\overline{y[n]} \quad (**)$$

Substituting:

$$a = \frac{(y[n], \hat{s}[n]) - N\overline{y[n]}\,\overline{\hat{s}[n]}}{\|y[n]\|^2 - N(\overline{y[n]})^2}$$

Note that due to the presence of noise, some the amplitude scale estimates may be erroneous, such as when $a < 0$, and therefore when this occurs the scale factor a is set to one by default, with b calculated accordingly.

From Figure 7.4 it is apparent that the vast majority of ANC unscaled output waveforms display lower normalized RMS error than the input waveforms for all stimulus levels and windows examined. This suggests that ANC system is very effective in estimating the underlying evoked potential from individual waveforms. Note also that the difference between normalized RMS error of the ANC output (scaled and unscaled) compared to ANC input tends to decrease for each stimulus level as the window is made shorter. Considering the typical ANC input, output and overall mean waveform in Figure 7.3, this can be explained by the fact that the ANC output waveforms tend to be a much better mean square match to the overall mean waveform than the ANC input waveforms in the temporally late regions. Additionally note that the scaled output waveforms universally show lower normalized RMS error then the unscaled output waveforms and hence also display lower normalized RMS error than the ANC input waveforms. Essentially this is a test of waveform shape, since discrepancies in amplitude are adjusted for in a mean-square sense by the scaling parameters a and b.

Figure 7.4a Normalized RMS Error Between Individual ANC Input, Output Waveforms and the Overall Mean Waveform (Waveform numbers 1-100) (500 ms window)

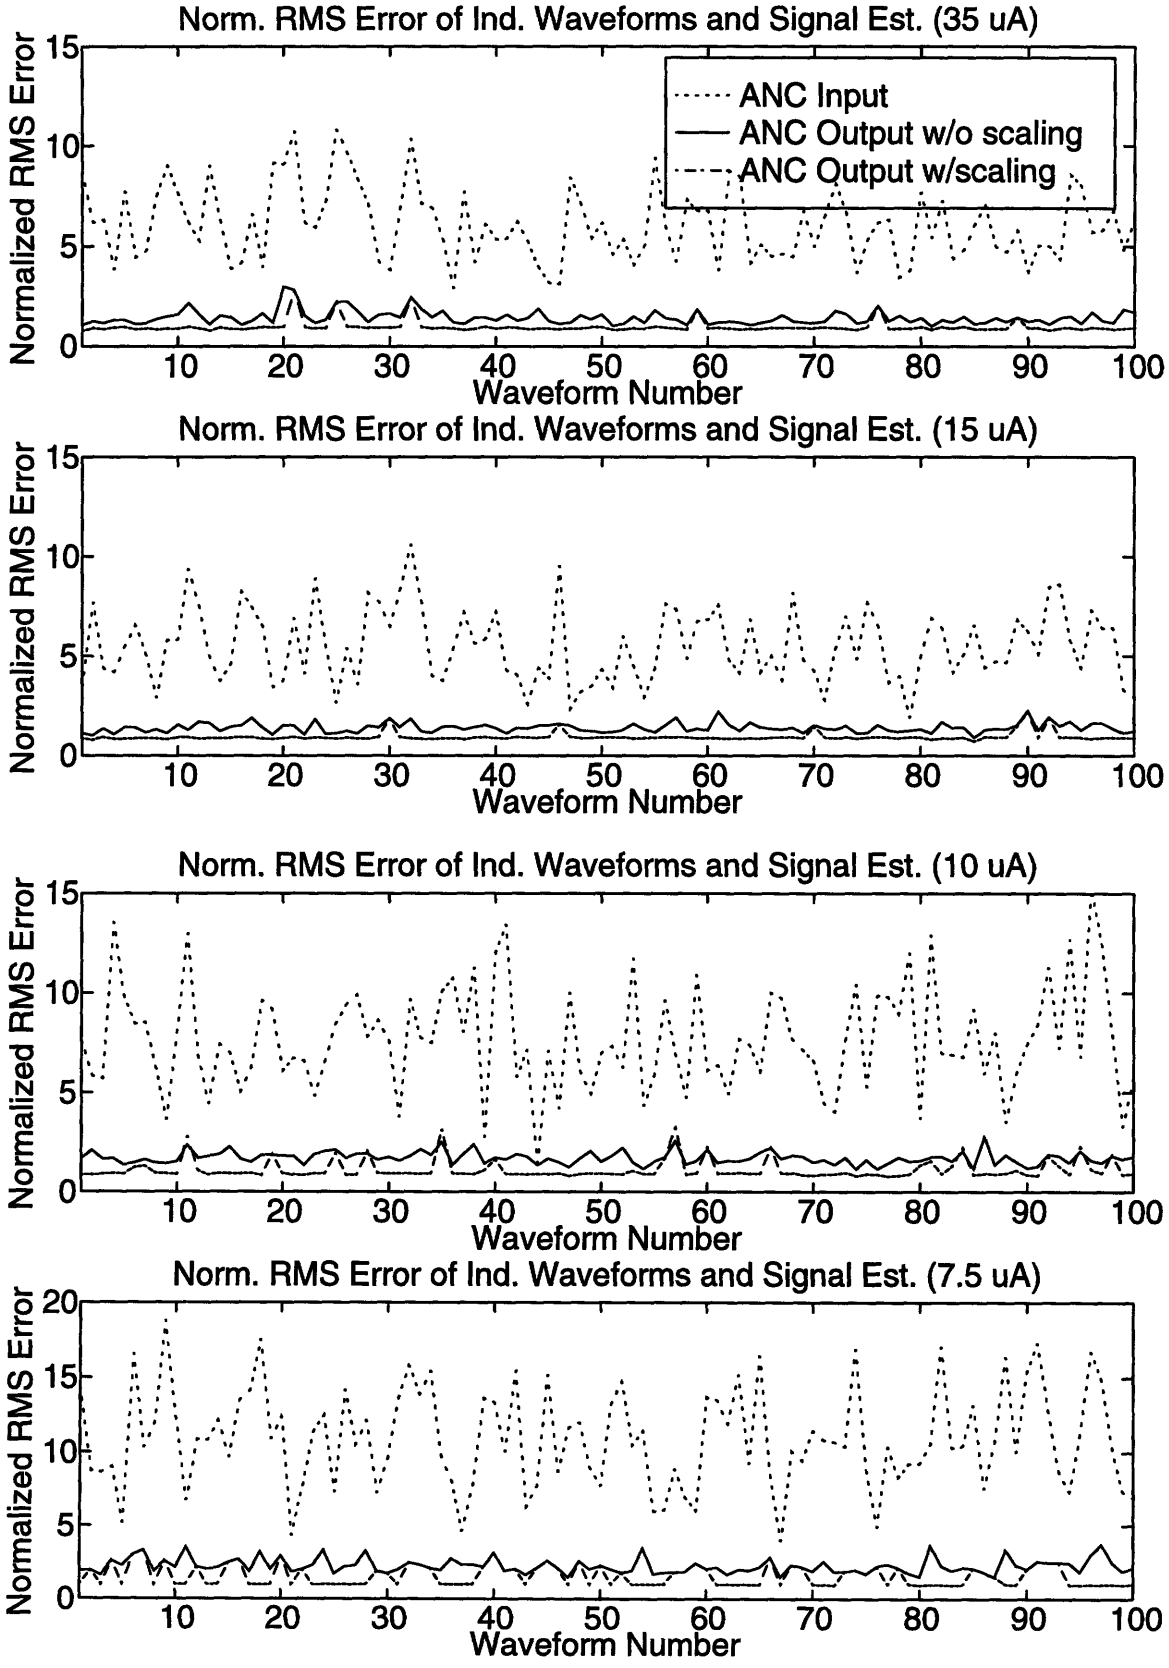

Figure 7.4b Normalized RMS Error Between Individual ANC Input, Output Waveforms and the Overall Mean Waveform (Waveform numbers 1-100) (100 ms window)

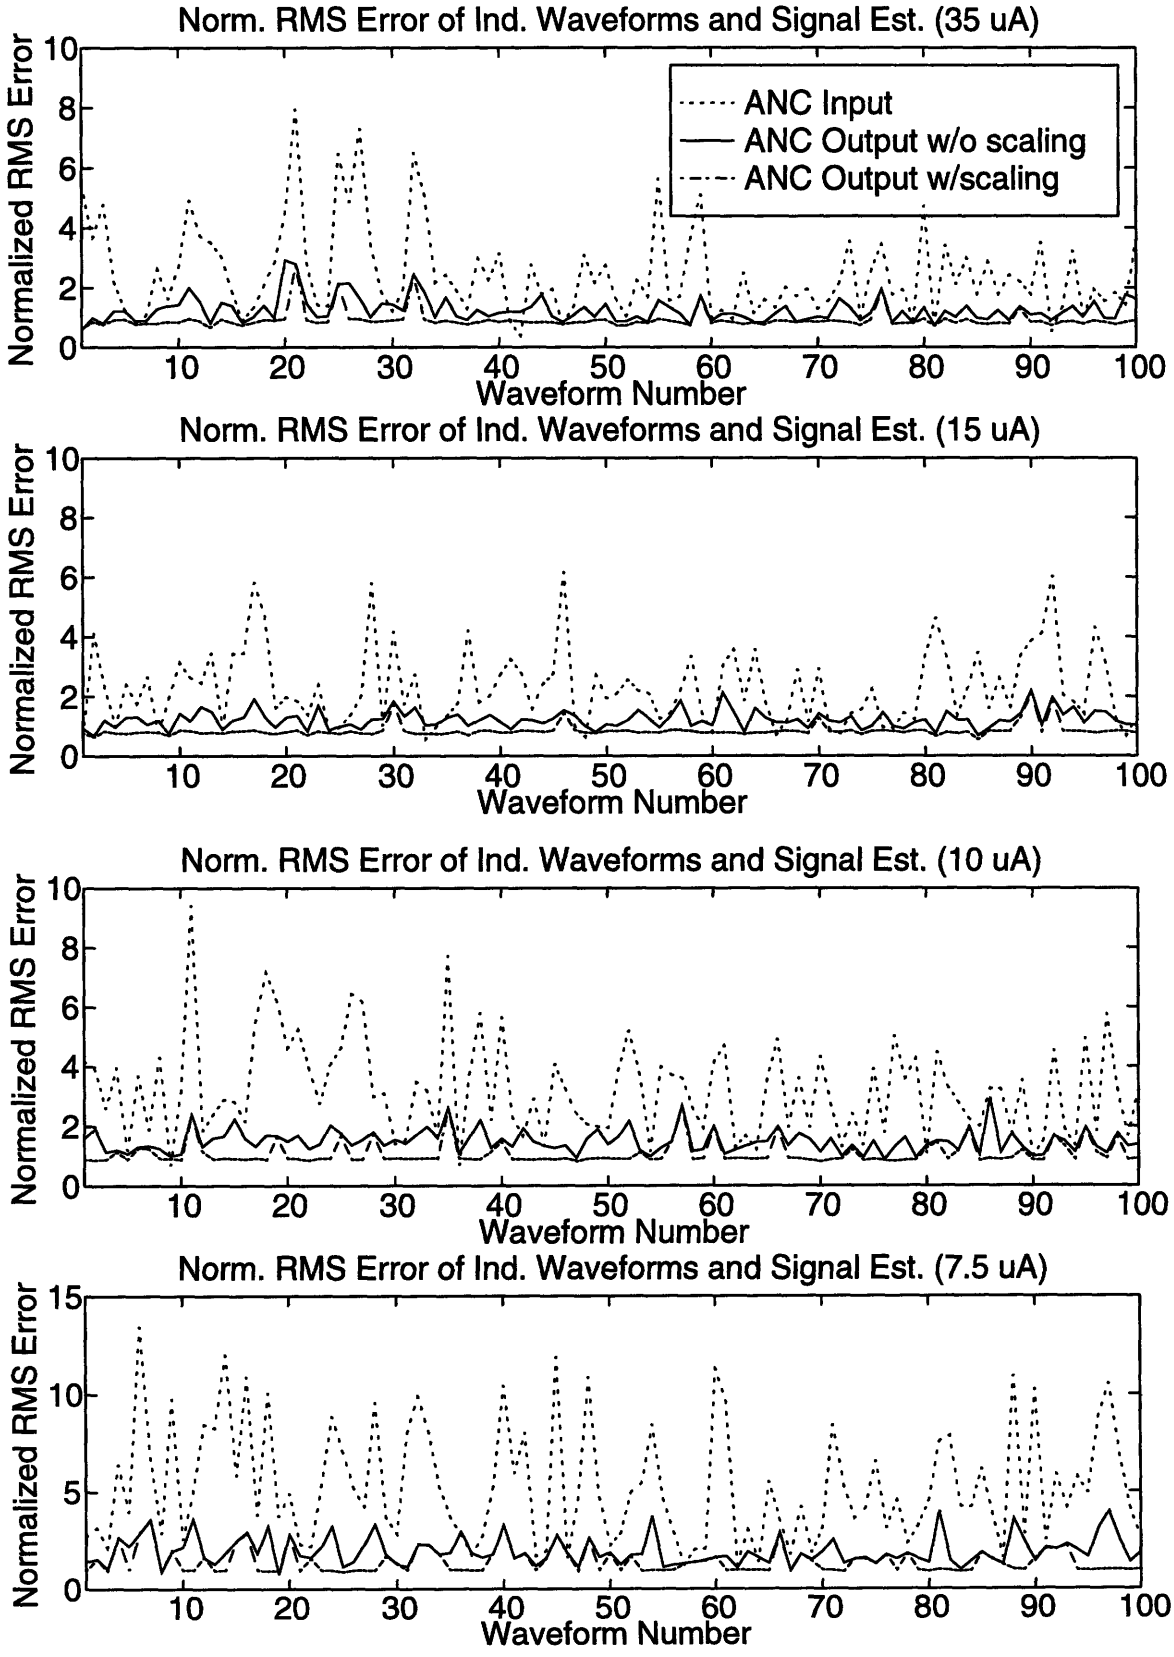

Figure 7.4c Normalized RMS Error Between Individual ANC Input, Output Waveforms and the Overall Mean Waveform (Waveform numbers 1-100) (50 ms window)

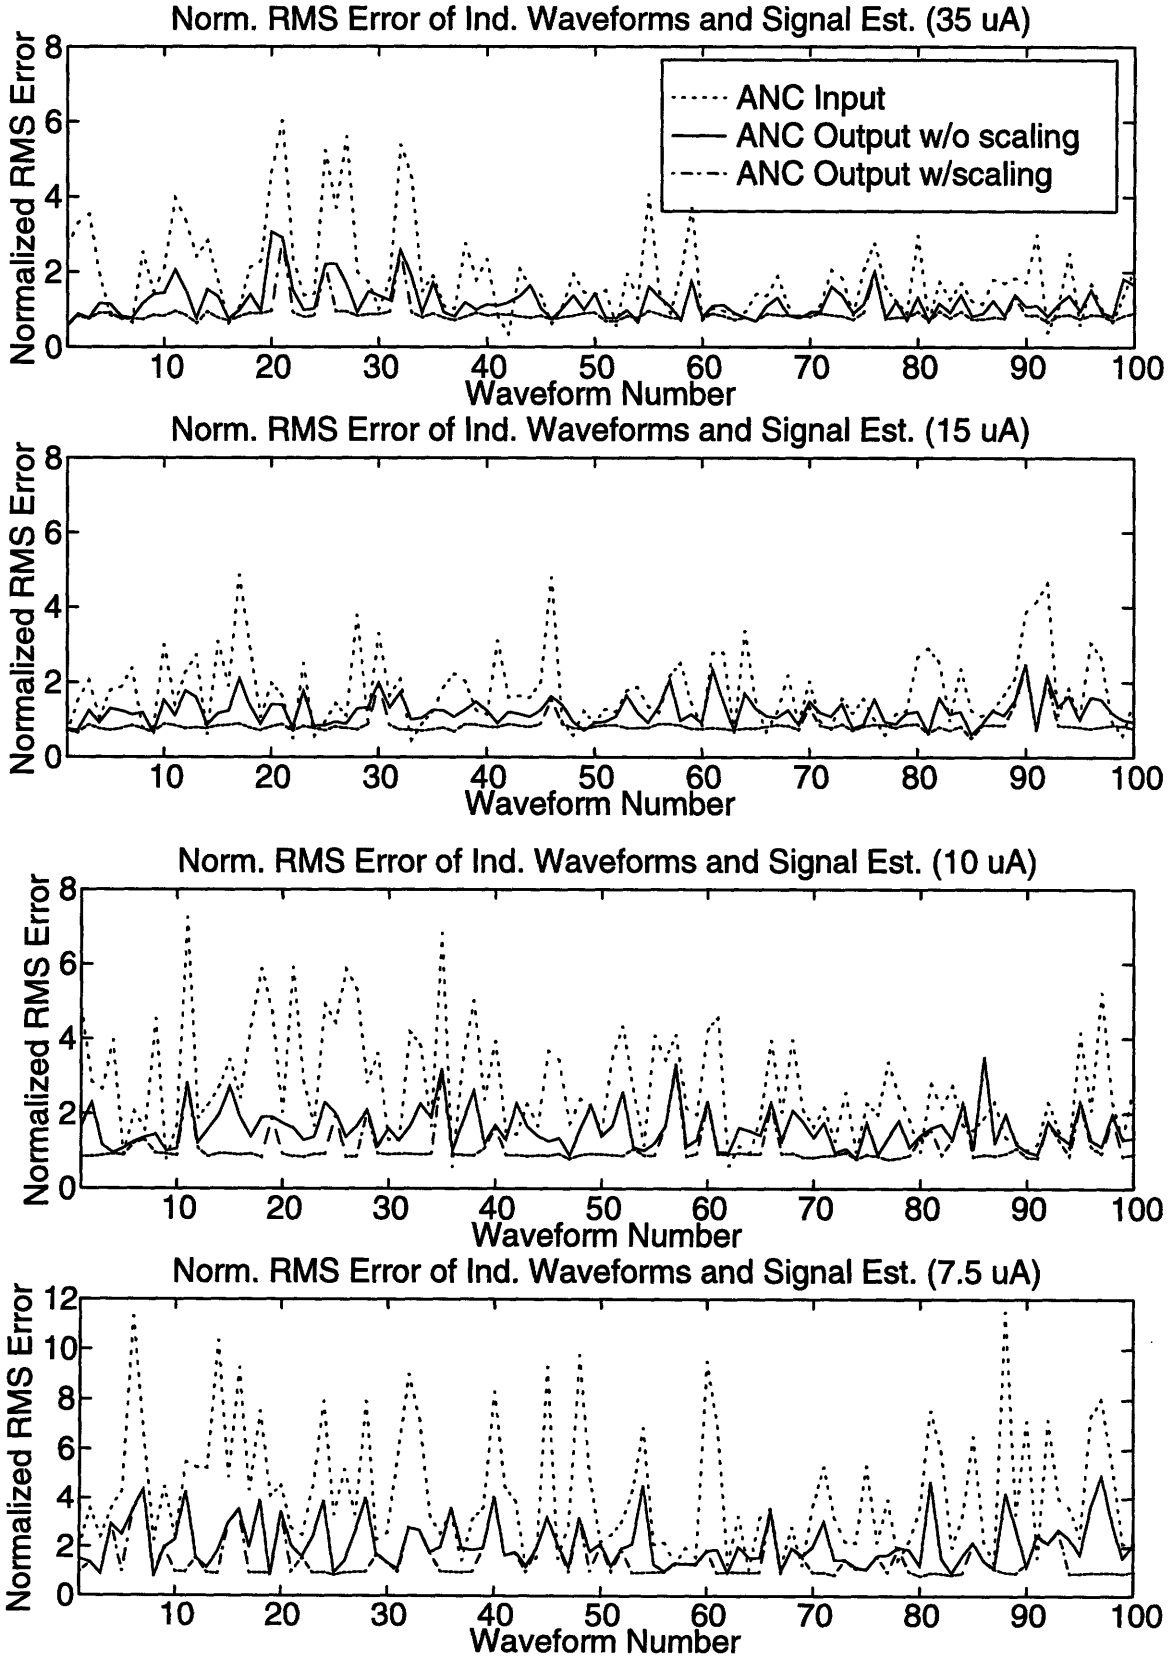

7.3.2 Normalized RMS Error of Cumulative ANC Input and Output Waveforms

Here a second error criterion based on entire waveforms was calculated, the normalized RMS error of the cumulative mean ANC input and output waveforms with the overall mean waveform. The cumulative mean waveform is the mean waveform formed from the entire past number of waveforms input to the ANC system, or output from it. This error criterion takes the same form as the normalized RMS error in section 7.3.1, except the cumulative mean waveform of the ANC input or output replaces the individual ANC input or output waveform respectively. Again the output was put into a scaled form as was done in the previous section, but the results were virtually identical to the unscaled output case, and hence are not reported here. The primary goal here is to determine how well the ANC system produces an output whose mean response is a good estimate of the underlying evoked potential. A secondary goal is a critical comparison between the ANC system evoked response estimate and the best estimate that can be made in real time using the mean waveform from the past number of recordings. Since the evoked potential is again estimated by the overall mean waveform formed from the ANC input waveforms, the normalized RMS error of the cumulative ANC input waveforms goes to zero by definition when all the waveforms are used. Thus we expect that after finding the mean of some number of waveforms, the normalized RMS error of the ANC input waveforms will be less than that for the corresponding number of ANC output waveforms. Again all four stimulus levels were used over three time windows of duration 500, 100 and 50 ms, all starting from 5 ms. These results are summarized in Figure 7.5 below.

From Figure 7.5, a number of salient features are noteworthy. First, for each stimulus level, the ANC output cumulative mean waveform produces an improved estimate of the overall mean waveform with a smaller number of waveforms than the ANC input cumulative mean waveform up to some crossover point where the estimate becomes worse. Second, this crossover point tends to decrease as the time window is decreased for all stimulus levels used. In other words the advantage in estimating the overall mean waveform from the ANC output cumulative mean waveform only exists for a small number of waveforms, and this number tends to decrease as we examine shorter waveform time windows. Third, although the normalized RMS error of the ANC input tends to decrease as more waveforms are used to form the cumulative mean for all stimulus levels and time windows (with some exceptions where the error fluctuates), the normalized RMS error of the ANC output tends to be independent of how many waveforms are used, after some initial decrease in normalized RMS error. Fourth, the steady-state normalized RMS error of the ANC output (i.e. the error as the number of waveforms in the cumulative mean increases) tends to be independent of the time window examined and the stimulus level, with a value of approximately less than or equal to one.

Given that the normalized RMS errors of individual ANC output waveforms tends to be less than that for the ANC input waveforms (section 7.3.1) we expect that the mean of the ANC output waveforms should be a better estimate of the overall mean waveform than the ANC input waveform, at least up to some crossover point, which is verified in Figure 7.5. We also expect that this crossover point should occur for a fewer number of waveforms as the time window decreases since the improvement in normalized RMS error of individual ANC output waveforms also diminishes as the time window is decreased, which is again confirmed in Figure 7.5. The independence of the normalized RMS error for the ANC output with the number of waveforms in the cumulative mean is possibly due

to the effectiveness of the ANC output in eliminating much of the additive noise. Under this scenario averaging more waveforms yields little improvement since each individual waveform is already a good estimate of the underlying evoked response. One problem with this explanation is the fact that the crossover point occurs when only a few waveforms are averaged (less than 20 in many cases), suggesting that the ANC system may have also introduced a bias which cannot be averaged out. This is consistent with the fact that the normalized RMS error is approximately the same in the steady-state for all stimulus levels and time windows used.

In summary, the ANC system produces individual waveforms which are much better estimates of the evoked potential than the individual recorded waveforms. Furthermore the cumulative mean waveform formed from the ANC output waveforms also are better estimates of the evoked potential than the cumulative mean waveform formed from the recordings for small numbers of waveforms used to form the cumulative mean, but for larger numbers the cumulative mean formed from the recordings is a better estimate. Again this is partially explained by the fact that our evoked potential estimate was formed from the overall mean waveform of recordings.

Figure 7.5a Normalized RMS Error of Cumulative Mean ANC Input and Output Waveforms (Waveform Numbers 1-200) (500 ms window)

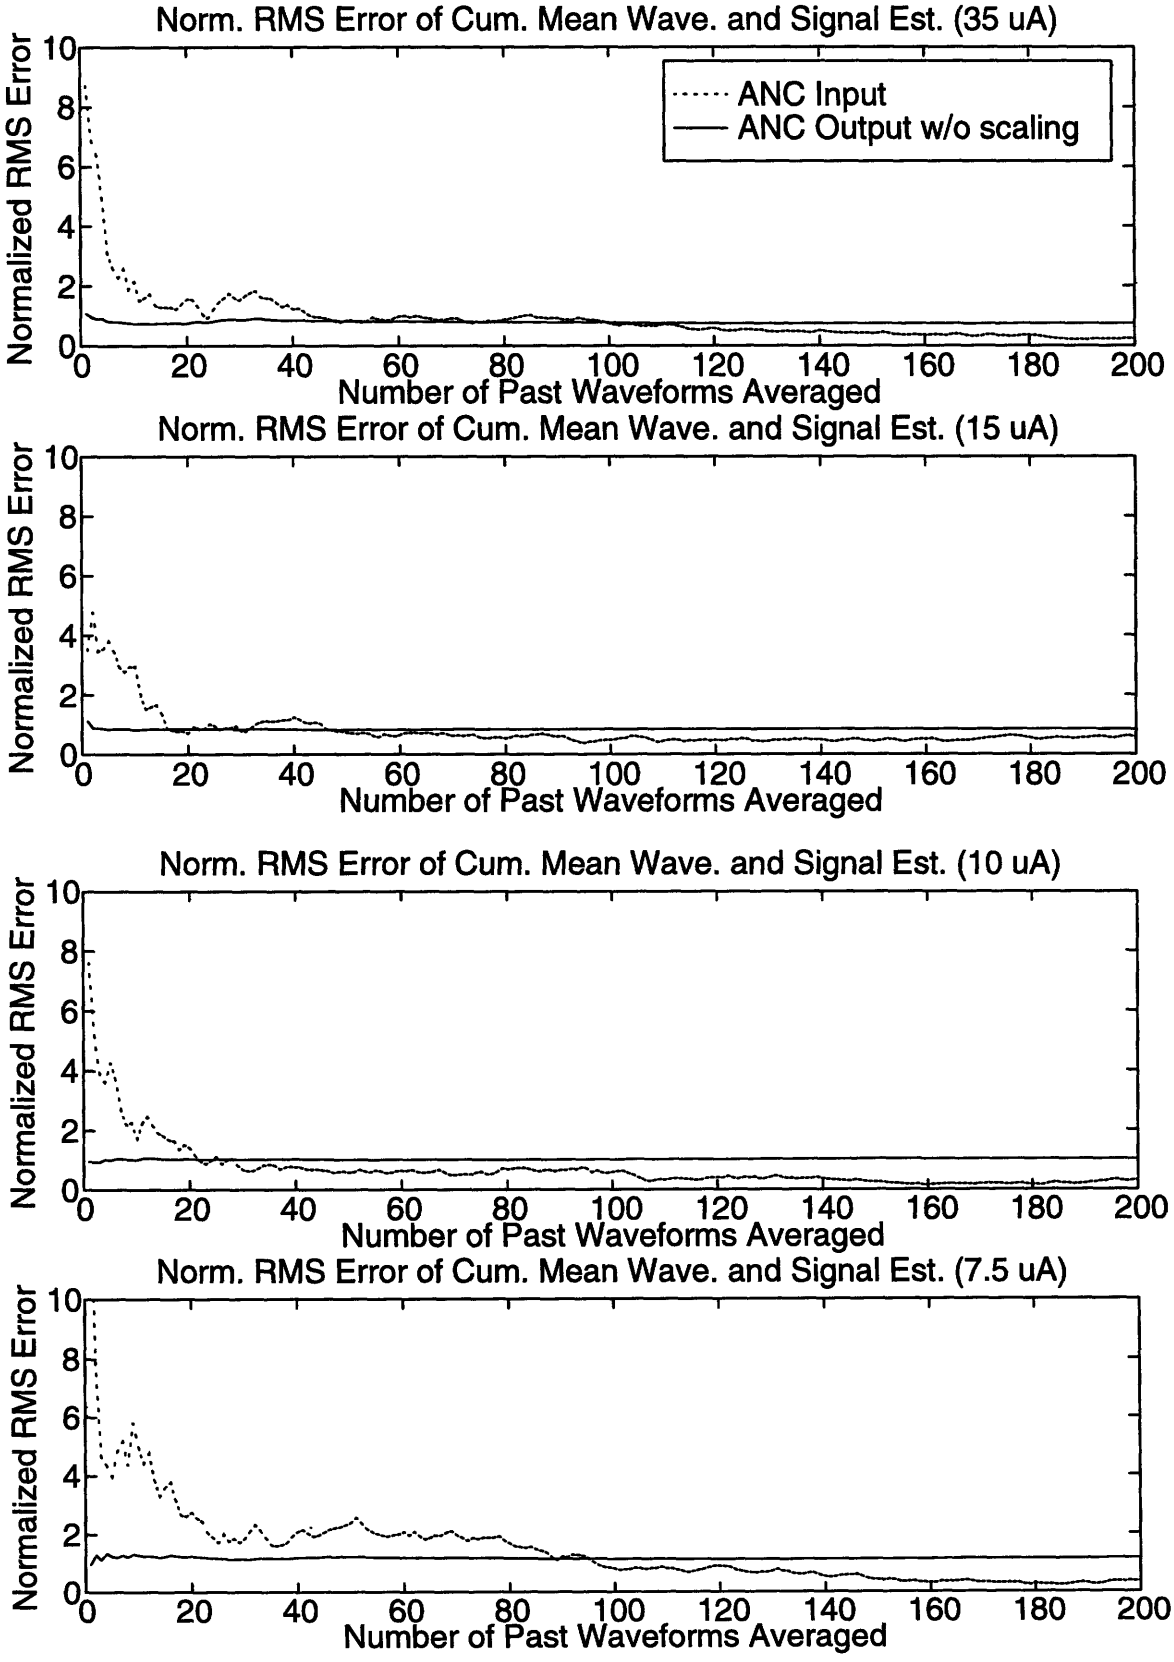

Figure 7.5b Normalized RMS Error of Cumulative Mean ANC Input and Output Waveforms (Waveform Numbers 1-200) (100 ms window)

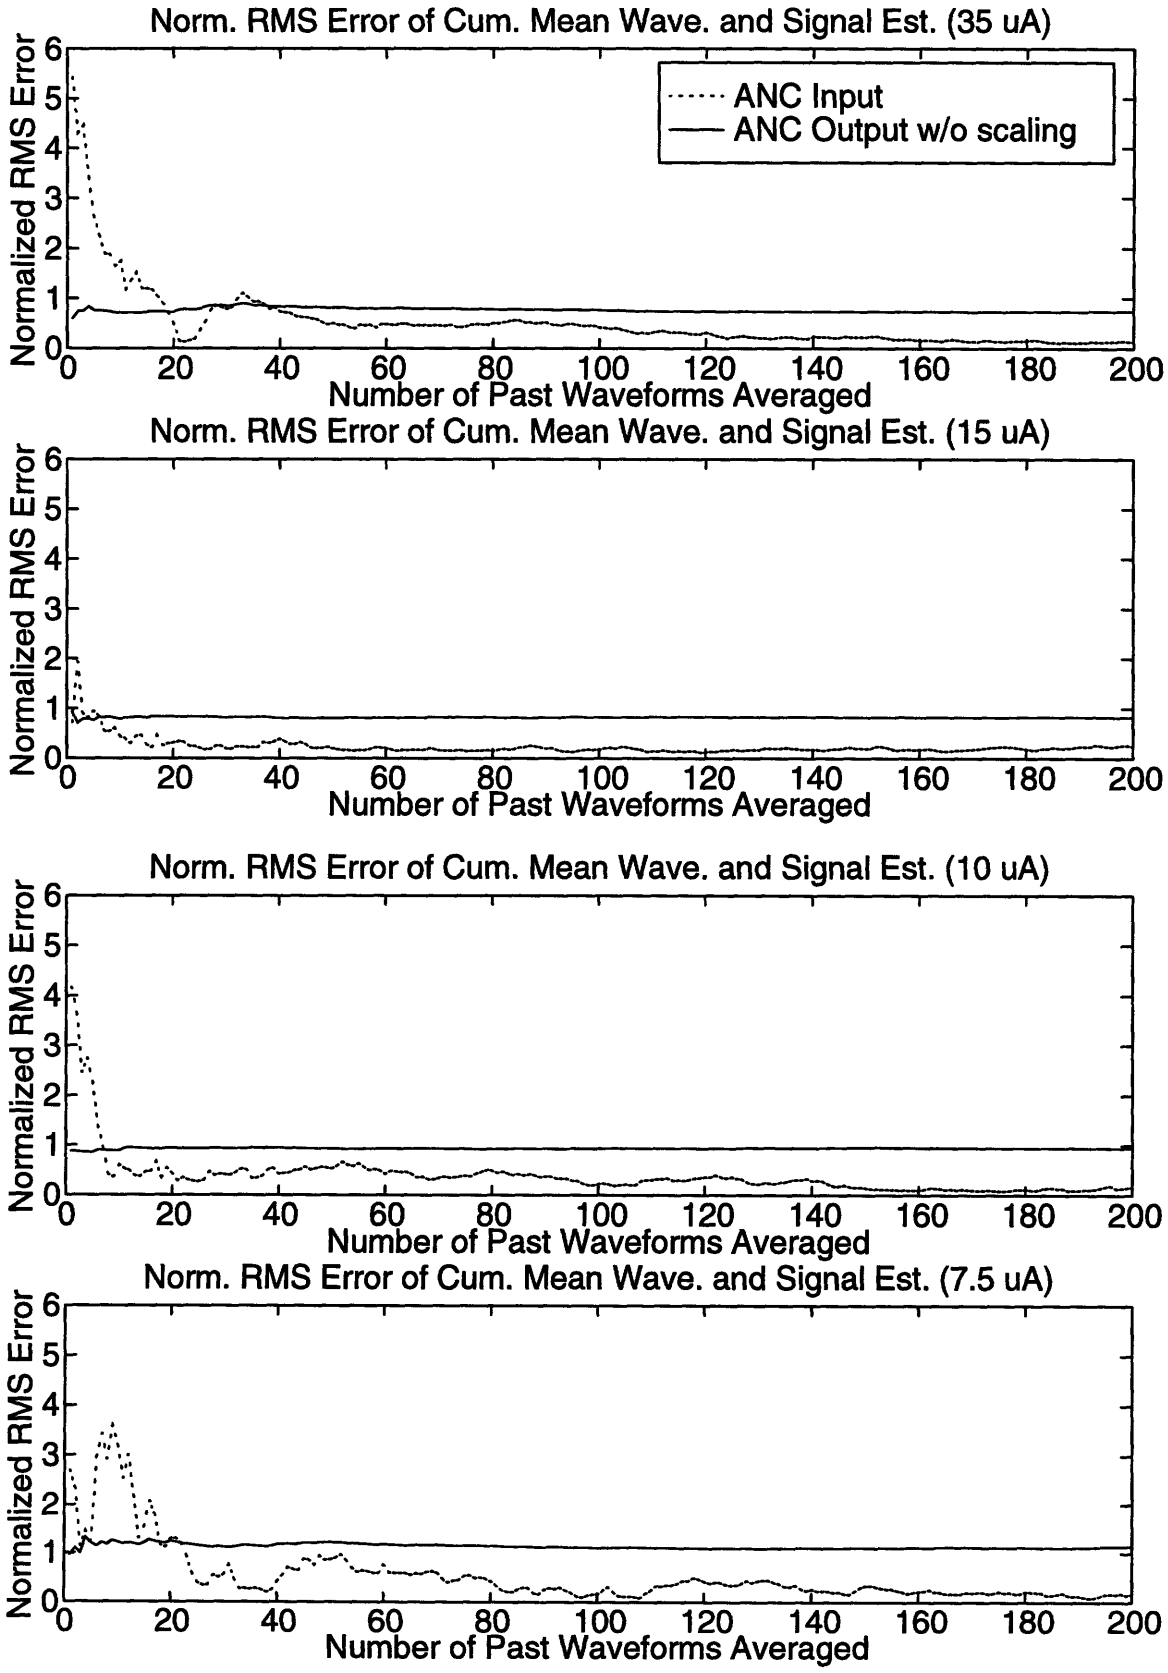

Figure 7.5c Normalized RMS Error of Cumulative Mean ANC Input and Output Waveforms (Waveform Numbers 1-200) (50 ms window)

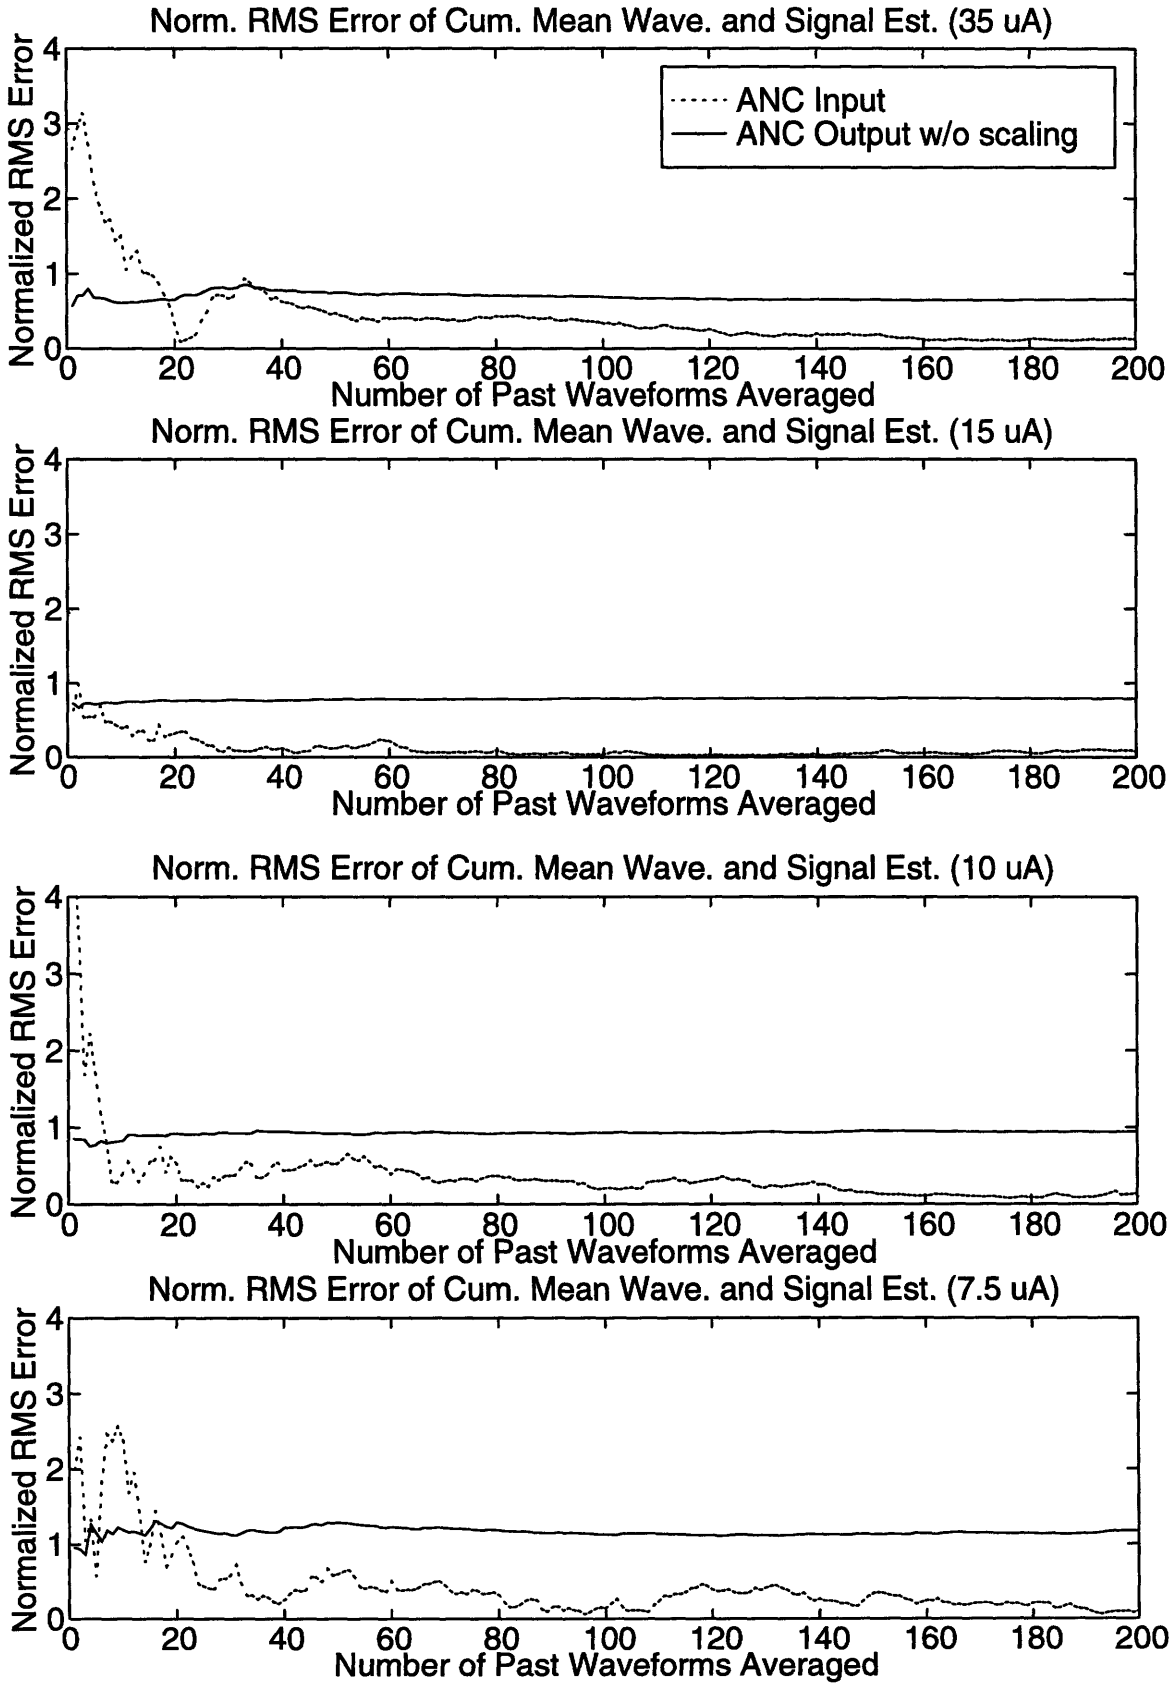

7.3.3 Feature Analysis of ANC Output Waveforms

The third and final error criterion is based on analysis of the peak-to-trough and RMS amplitude about the time average features of RSM waveforms. This error criterion attempts to quantify how well the features of mean waveforms compare to the features of the overall mean waveform, as opposed to the normalized RMS error criterion discussed above which compared individual or cumulative mean waveforms to the overall mean waveform. In order to facilitate comparison to averaging the recorded waveforms, ideal bandpass filtering, and Weiner-type filtering, the same experimental conditions were repeated here, with the exception of the lower window time for feature extraction. Instead of using a lower window time of 5 ms, here a lower window time of 10 ms was used to counteract the effect of ANC convergence during the first 5 ms after the stimulus artifact. The table is noted with an asterisk for this reason. All other experimental parameters were repeated. Specifically, for each set of output waveforms corresponding to a single stimulus level, 40 individual RSM waveforms were formed from 100, 50 and 20 individual waveforms, and the window for feature extraction extended from 10-50 ms. For each of the experimental conditions, the mean, sample standard deviation and the percent spread were calculated. These results are summarized in Table 7.1 below, where the earlier results are repeated for comparison.

Examination of Table 7.1 indicates mixed results. For the 35 μA data, the percent spread was lower for each of the experimental conditions used, although the mean RSM feature tended to be roughly 10% below the mean RMS feature for the other signal processing methods suggesting the possible introduction of a feature bias due to the ANC system. For the 15 μA data, the percent spread was lower for each of the experimental conditions except when 100 individual waveforms were used for the RMS feature, and no appreciable attenuation of the mean RMS feature was evident. However, for the 10 and 7.5 μA data, the percent spread was almost universally worse in the ANC case than the other three methods.

Numerous possible explanations exist for these mixed results. First it is possible that the statistical characteristics of the primary noise associated with the recording process changed during the experiment such that they no longer were a close match to the reference waveforms which were obtained prior to stimulus presentation. The stimuli were presented in decreasing order of amplitude, so the lowest stimulus level of 7.5 μA occurred with the greatest temporal distance from the reference waveform recordings. Although apparently the ANC system was able to track phase changes between periodic reference and primary noise components with some effectiveness, it would not be able to account for changes in breathing or heart rate as the experiment was performed. It is possible that additional a priori recordings made temporally closer to each stimulus presentation may alleviate this problem, but future work is needed to verify this.

A second possible explanation for the mixed results may be unrelated to changes in the primary noise statistics. It is also possible that the ANC system was able to track amplitude scale variations more effectively than ideal bandpass filtering or Weiner-type filtering. Under this scenario, averages of relatively few waveforms may be *expected* to be different due to the different amplitude scale factor contributions from the individual waveforms. As the number of individual waveforms used to form the RSM waveforms is increased, these differences should tend to diminish since the average of the amplitude scale factors tends towards its expected value.

A final possible explanation is that the ANC system distorts the output waveforms such that more feature variation exists among individual waveforms. From the data it is unclear if any of the above explanations or combinations of them is correct.

As an aside, an alternative error criterion of feature mean squared error was also examined here, which is the mean squared deviation between the RSM waveform features and the feature of the overall mean waveform. This error criterion was discussed in section 3.2 and compared to the sample variance where it was shown that under certain conditions these two functions yield similar results, i.e. when the feature of the overall mean waveform is approximately equal to the mean of the features from the RSM waveforms. This alternative error criterion was not adopted here because it was found to yield very similar results to the sample standard deviation.

Table 7.1a Highlight of Filtering for 35 μA Data (N_w=40, W_l=5ms, W_u=50ms)

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	523	31.28	12.0	108	6.23	11.5
	Best Ideal Filter	522	29.60	11.4	108	5.78	10.7
	Weiner-Type Filter	510	23.68	9.3	106	4.63	8.7
	Active Noise Filter *	520	21.96	8.5	97	3.62	7.5
Ni=50	No Filter	526	40.91	15.6	109	8.00	14.7
	Best Ideal Filter	524	39.95	15.2	109	7.50	13.7
	Weiner-Type Filter	518	33.71	13.0	106	7.14	13.4
	Active Noise Filter *	523	31.13	11.9	97	5.75	11.8
Ni=20	No Filter	527	56.11	21.3	111	13.14	23.8
	Best Ideal Filter	526	55.48	21.1	107	11.67	21.9
	Weiner-Type Filter	522	51.54	19.8	108	9.66	17.9
	Active Noise Filter *	535	46.42	17.3	101	8.91	17.7

Table 7.1b Highlight of Filtering for 15 μA Data (N_w=40, W_l=5ms, W_u=50ms)

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	369	27.68	15.0	79	5.40	13.7
	Best Ideal Filter	372	26.82	14.4	78	5.03	12.9
	Weiner-Type Filter	350	22.79	13.0	75	4.02	10.7
	Active Noise Filter *	384	24.83	12.9	76	4.75	12.6
Ni=50	No Filter	373	43.23	23.2	80	8.73	21.8
	Best Ideal Filter	370	42.76	23.1	79	8.25	20.9
	Weiner-Type Filter	346	34.31	19.8	74	7.21	19.4
	Active Noise Filter *	387	33.83	17.4	77	6.86	17.9
Ni=20	No Filter	385	50.13	26.0	81	11.22	27.6
	Best Ideal Filter	382	50.04	26.2	81	11.19	27.5
	Weiner-Type Filter	354	44.42	25.1	76	9.35	24.5
	Active Noise Filter *	406	49.39	24.3	81	9.72	24.1

Table 7.1c Highlight of Filtering for 10μA Data (N_w=40, W_l=5ms, W_u=50ms)

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	188	15.26	16.3	41	3.95	19.5
	Best Ideal Filter	189	14.58	15.4	39	3.60	18.2
	Weiner-Type Filter	166	10.34	12.5	37	2.77	15.1
	Active Noise Filter *	196	21.00	21.4	37	3.47	18.7
Ni=50	No Filter	190	20.24	21.3	42	5.41	25.9
	Best Ideal Filter	192	19.90	20.8	41	4.92	24.3
	Weiner-Type Filter	173	16.60	19.2	37	4.23	22.8
	Active Noise Filter *	200	23.82	23.8	40	4.29	21.7
Ni=20	No Filter	207	30.31	29.3	44	6.70	30.7
	Best Ideal Filter	208	28.14	27.1	44	6.68	30.7
	Weiner-Type Filter	183	23.52	25.8	38	6.08	31.7
	Active Noise Filter *	214	38.34	35.9	40	7.08	35.7

Table 7.1d Highlight of Filtering for 7.5 μA Data (N_w=40, W_l=5ms, W_u=50ms)

		P-to-T			RMS		
		Mean (μV)	S. Dev. (μV)	%spread	Mean (μV)	S. Dev. (μV)	%spread
Ni=100	No Filter	108	11.12	20.6	21	1.84	17.2
	Best Ideal Filter	89	8.36	18.8	20	1.69	16.5
	Weiner-Type Filter	100	8.46	16.9	20	1.51	15.5
	Active Noise Filter *	86	13.15	30.6	19	2.49	25.6
Ni=50	No Filter	115	17.34	30.3	23	3.28	28.8
	Best Ideal Filter	102	11.25	22.0	22	2.62	24.3
	Weiner-Type Filter	105	14.16	26.8	21	2.85	27.0
	Active Noise Filter *	88	19.07	43.1	20	3.13	32.1
Ni=20	No Filter	123	26.33	42.7	26	6.14	47.8
	Best Ideal Filter	104	20.88	40.0	24	4.67	38.9
	Weiner-Type Filter	112	25.25	44.8	23	4.33	37.5
	Active Noise Filter *	89	19.98	44.7	21	3.94	37.4

7.4 Amplitude Scale Model Revisited

In an effort to determine why the feature analysis results were mixed, the amplitude scale model verification test was repeated here for the ANC output data. The details of the modelling test are treated in section 6.3 and are not repeated here. We again calculated our ratio R whose values fall between 0 and 1 for each of the four ANC output data sets using windows of duration 20, 30, 50 and 100 ms, each beginning at 5 ms and sweeping upwards in 0.5 ms increments for the entire waveform duration of 500 ms. Shown below in Figure 7.6 are the average values of R across data sets for the four windows used. The individual values of R for each stimulus level corresponding to the time when R_{ave} was a maximum is summarized in Table 7.2 along with the earlier results from the ANC input data.

Figure 7.6 again shows strong evidence in favor of the amplitude scale model for some of the early portions of the waveforms before 50 ms. Comparison of these early portions across waveforms shows striking similarity in the distribution of R_{ave} for the different windows used, which was not evident when the test was applied to the ANC input data in section 6. Although the actual value of R_{ave} is lower for these tests than for the ANC input data evaluated earlier, comparison of the values of R in Table 7.2 does show similar values of R for the 15, 10, and 7.5 μA data. In fact for the 20 ms window, the 10 and 7.5 μA data value of R is larger for the ANC output data than that for the ANC input data. This suggests that the ANC system effectively removed large amounts of the additive noise in these data sets, and was able to track the effects of the amplitude scale factor. Although far from conclusive, these results suggest that the mixed feature analysis results were due in part to the effect of averaging waveforms with highly divergent amplitude scale factors, increasing the percent spread error criteria as a result.

Figure 7.6 Average Value of R across ANC Output Sets for Numerous Time Windows

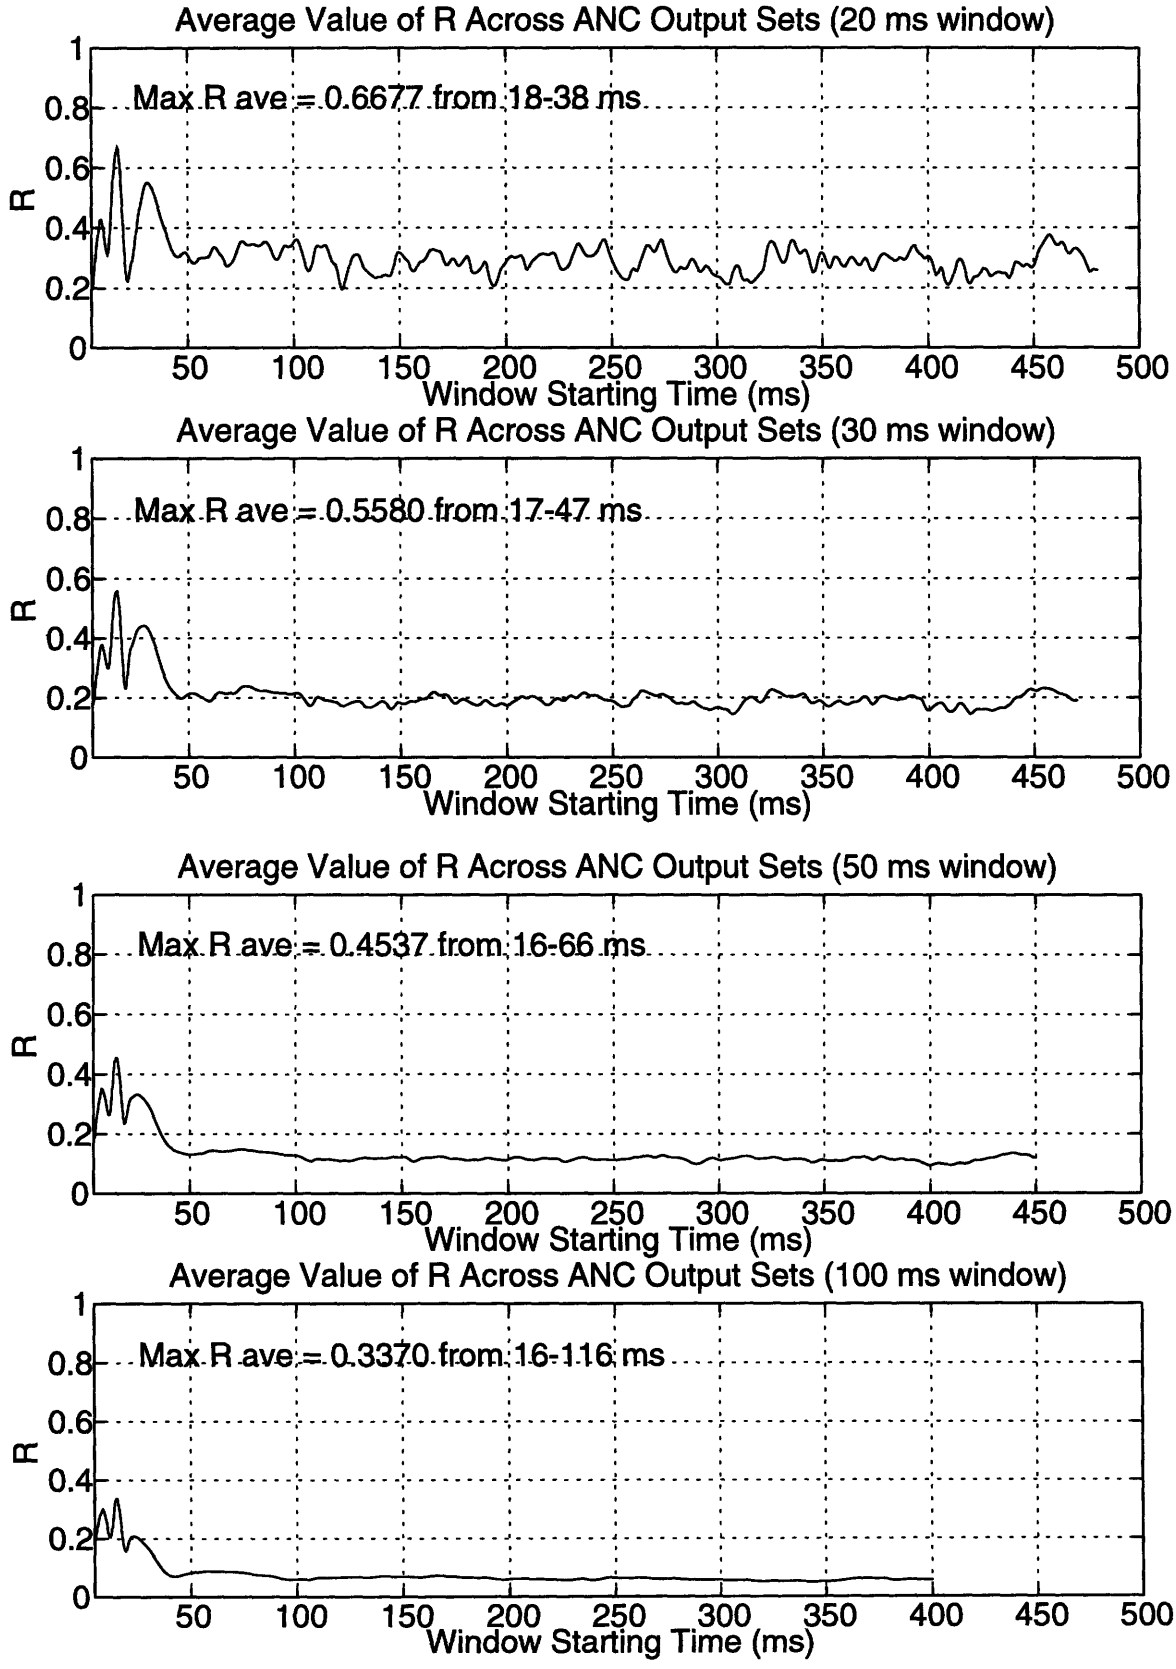

Table 7.2 Values of R for Each Data Set When Average R across Data Sets is Maximum (Raw Data and ANC Output Data)

		35 μA Data	15 μA Data	10 μA Data	7.5 μA Data
20 ms Window	Raw Data (20-40 ms)	0.7910	0.8243	0.7716	0.7190
	ANC Data (18-38 ms)	0.4184	0.7186	0.7757	0.7580
30 ms Window	Raw Data (19-49 ms)	0.7668	0.8081	0.7317	0.5791
	ANC Data (17-47 ms)	0.2618	0.6130	0.6942	0.6631
50 ms Window	Raw Data (58-108 ms)	0.6401	0.6401	0.6242	0.6016
	ANC Data (16-66 ms)	0.1059	0.5057	0.5997	0.6037
100 ms Window	Raw Data (35-135 ms)	0.6181	0.6378	0.6594	0.6462
	ANC Data (16-116 ms)	0.0637	0.3348	0.4680	0.4816

8. Conclusion

8.1 Restatement of Thesis Goals

The ultimate goals of this thesis were to (a) reduce the number of individual recordings necessary in order to obtain an accurate estimate of the underlying evoked potential, and (b) to determine the validity of an alternative signal model which assumes that the recorded waveform is the sum of an amplitude scaled evoked potential and additive noise. These goals were made in light of an overall desire to extract some characteristic features from the recorded waveforms which are a function of stimulus level, and the effectiveness of the retinal implant in producing a visual type of response. Two features were examined throughout, the peak-to-trough amplitude and the RMS amplitude about the time average. All the methods utilized four data sets corresponding to four stimulus levels, recorded during a single experiment on a single subject.

8.2 Summary of Methods

Methods used to reduce the number of recordings include ideal bandpass filtering, Weiner-type filtering and active noise cancellation. The method of averaging is standard in clinical practice, and is assumed to produce very good estimates of the evoked potential, but it suffers from requiring large number of individual waveforms to form an accurate estimate. Ideal bandpass filtering is a simple LTI method designed to remove noise components which do not share the same frequency range as the evoked potential frequency components. Weiner-type filtering is an LTI method which uses knowledge about the evoked potential and additive noise statistics to form a single filter for each data set examined. It suffers from its reliance on a priori knowledge of these signal and noise statistics. Active noise cancellation is an nonlinear, time-varying method which utilizes reference noise recordings which are assumed correlated to the noise present in the primary recordings which contain the desired evoked response. This method was here employed using a priori recordings when no stimulus was present, a departure from more traditional active noise cancellation systems which use simultaneously obtained recordings. In order to facilitate comparisons between the effectiveness of these methods, the error criterion of percent spread was adopted, which is based on the mean and the feature sample variance of waveforms formed from small averages of individual waveforms chosen randomly from the waveform population. Two additional error criteria of normalized RMS error of individual waveforms and of cumulative mean waveforms were adopted in the case of active noise cancellation, which compares individual or mean evoked potential waveform estimates to the overall mean waveform, which is assumed throughout to be the best estimate of the underlying evoked response. Normalized RMS error of both individual and cumulative mean waveforms was compared to that between the individual or cumulative mean raw recorded waveforms and the overall mean waveform. These normalized RMS error measures were made over three data lengths, from 5-500, 5-100 and 5-50 ms.

The method used to determine the validity of the alternative model was based on principles of linear algebra, orthogonality and Principle Components analysis applied to the signal statistics. This method employed an estimate of the recorded waveform covariance matrix and the evoked potential signal estimate from the overall mean waveform. The test was applied to numerous sections of the recorded waveforms and for numerous time window durations. The test was also applied to the individual evoked potential esti-

mates from the active noise cancellation system output.

8.3 Summary of Results

The methods used to reduce the number of recordings for an accurate evoked potential estimate were all evaluated in light of the number of waveforms used in averaging, which again is assumed to produce the most accurate evoked potential estimate, although requiring a large number of individual waveforms. Ideal bandpass filtering generally produced a moderate reduction in percent spread, although the required lower and upper cutoff frequencies for optimal performance were not consistent across experimental conditions or data sets. Weiner-type filtering consistently produced improvements over averaging and ideal bandpass filtering for all experimental conditions and all but the smallest stimulus data sets, with only few exceptions when RSM waveform were formed from only 20 individual waveforms. The results for the smallest stimulus data set were mixed, producing lower percent spread than averaging or ideal bandpass filtering when RSM waveforms were formed from 100 individual waveforms, but producing larger percent spread than ideal bandpass filtering when RSM waveforms were formed from 50 or 20 individual waveforms. In this case Weiner-type filtering still was an improvement over averaging.

The results of active noise cancellation also were also somewhat mixed, although these results suggest some potential flaws in our choice of percent spread as an error criterion. In terms of the normalized RMS error of individual waveforms which compares individual waveforms to the overall mean waveform in the mean-square sense, the results were very promising. The normalized RMS error of individual waveforms was consistently lower for the active noise cancellation output waveforms than that for the input waveforms, over each of the three data record lengths examined, and over all four data sets. This suggests that the active noise cancelation system produced very good individual evoked potential signal estimates. In terms of normalized RMS error of the cumulative mean waveform which compares mean waveforms to the overall mean waveform the results were also positive, where it was shown that for mean waveforms formed from small numbers of individual waveforms the ANC output produced an improved estimate of the overall mean waveform than the same number of ANC input waveforms. However, in terms of the percent spread error criterion which examines features of mean waveforms, the results were less uniform. For the 35 and 15 μA data sets, the percent spread was consistently lower (with only a single exception) than all other methods, for all experimental conditions and for both features. For the 10 and 7.5 μA the percent spread generally increased over that obtained from the other methods. A number of possible explanations were put forth, including the possibility that the amplitude scale factor model is correct, and that the active noise cancellation system tracks it well resulting in RSM waveforms which are expected to have divergent feature sizes due to the amplitude scale factor. This scenario was confirmed tentatively by the application of the alternative model test to the output waveforms which appear to show that the model is valid for this data (with the exception of the 35 μA data). This also suggests that the percent spread error criterion may be specious when RSM waveforms are formed from relatively few numbers of individual waveforms.

Application of the test for determining the validity of the amplitude scaled signal model appear to confirm the model over some of the early portions of the individual wave-

forms, both for the raw recordings and for the active noise cancellation output waveforms. Specifically it appears that the amplitude scale factor is present from roughly 20-40 ms for each of the four data sets, and that its variance is relatively large in this region. We expect that if present, the amplitude scale model should be confirmed over the temporally early region of the waveforms only since the evoked potential tends to die out during the temporally late portions.

8.4 Future Work

Future work should concentrate primarily in two areas. First development of signal processing schemes which are based on the amplitude scale model. These methods will likely borrow heavily from communication theory and the nonlinear effects of channel characteristics. These methods will likely be adaptive in nature, able to track changes in the amplitude scale factor and noise statistics over time. They should not rely on global statistics derived over long periods of time, such as those used in the Weiner-type filtering. The second concentration should be in the area of active noise cancellation. Work should be done which utilizes simultaneous recordings from multiple reference sources such as from the heart and lungs. More work should also be done using different adaptive algorithms besides the LMS algorithm employed here, such as the recursive least squares algorithm. Other algorithms may be able to track nonstationarities more effectively than the LMS has done here. Finally, more work could concentrate on modelling, preferably adaptive modelling which may reveal the dynamic nature of the system being studied. Although these future areas represent considerable advancements over the present work, their basis is formed from the advances achieved here.

Appendix A: Experimental Setup

The experimental data acquisition system essentially consists of a bandpass/gain stage, a second order filter stage and a data acquisition board and computer stage. These are depicted below in the block diagram.

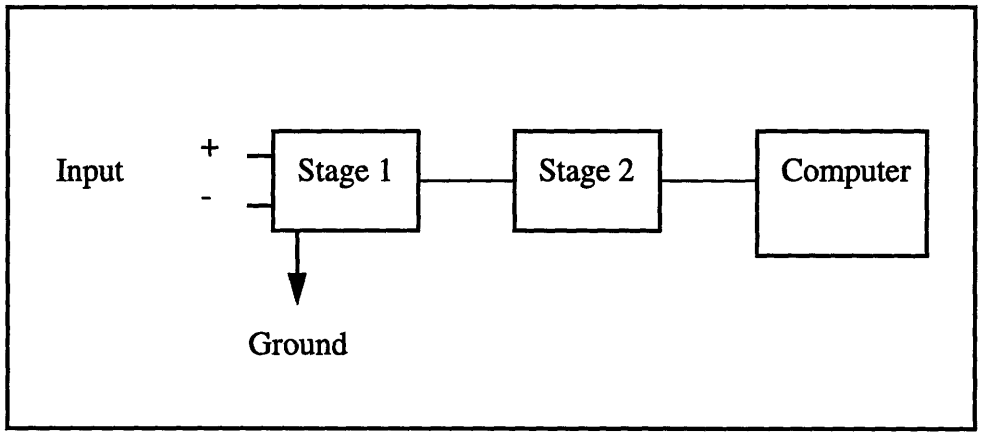

Figure A.1 Block Diagram of Data Acquisition System

Stage 1: Consists of a 10x differential instrumentation amplifier, a 100x gain sub stage, and a second order active Butterworth bandpass filter with a gain of about 1.59 and with lower and upper 3dB cutoff frequencies of about 0.5 Hz and 3.125 KHz respectively.

Stage 2: Consists of an isolator circuit used to isolate the electrode inputs from ground, and two second order active Butterworth bandpass filters whose lower and upper 3dB cut-off frequencies are about 0.33 Hz and 2.6 KHz respectively.

Computer: Consists of a 486 PC clone CPU using a Microstar data acquisition board at a sampling rate of 2 KHz using 12 bit precision. The triggering of the data acquisition system is synchronized to the stimulation pulse.

* Note that although the data used in this thesis was acquired using a sampling rate of 2KHz, and the bandpass filters described above have an upper 3dB cutoff above 1 KHz, some aliasing of signal components above 2 KHz are aliased down to the lower frequencies we are interested in. These high frequency signal components were previously examined when the data was acquired at 7.14 KHz, and were found to be very small in comparison to the low frequency signal and noise components, thus the need for new anti-aliasing filters was unnecessary.

Appendix B: Derivation of Amplitude Scale Factor Estimator

Derivation of estimator:

Using the concept of the Euclidean norm which describes the length of a vector, we seek to minimize the length of the error vector (or length squared for simplifying calculations). The error vector is defined as the difference between the recorded waveform $x_i[n]$ and the estimate of the amplitude scaled deterministic evoked potential $\hat{a}_i\, s[n]$. Written mathematically, we seek some $\hat{a}_i$ such that we minimize

$$\eta \equiv \left\| x_i[n] - \hat{a}_i\, s[n] \right\|^2 \equiv \sum_n \left(x_i[n] - \hat{a}_i\, s[n] \right)^2$$

The minimum of η can be found by setting $\frac{\partial \eta}{\partial \hat{a}_i} = 0$, solving for $\hat{a}_i$ and checking $\frac{\partial^2 \eta}{\partial \hat{a}_i^2} > 0$ verifying that the chosen value of $\hat{a}_i$ yields a minimum of η.

$$\frac{\partial \eta}{\partial \hat{a}_i} = \sum_n 2\left(x_i[n] - \hat{a}_i\, s[n] \right)\left(-s[n] \right) = 0$$

$$-\sum_n x_i[n]\, s[n] + \hat{a}_i \sum_n s^2[n] = 0$$

$$\therefore \hat{a}_i = \frac{\sum_n x_i[n]\, s[n]}{\sum_n s^2[n]} \qquad \left(\text{assuming } \sum_n s^2[n] \neq 0\right)$$

Alternate derivation of estimator:

Consider the geometric interpretation of η in two-dimensional time space (i.e. n = 1, 2) depicted below, where $x_i[n]$ and $s[n]$ are vectors, denoted $\mathbf{x}_i$ and $\mathbf{s}$ respectively, and $\hat{a}_i$ is a scalar. $\mathbf{e}$ is a vector which represents the minimum distance between $\mathbf{x}_i$ and $\hat{a}_i\,\mathbf{s}$ when the proper choice of $\hat{a}_i$ is made. Since in general $\mathbf{x}_i$ and $\hat{a}_i\,\mathbf{s}$ are not collinear (unless there is no noise, i.e. $v_i[n]$, or if the noise equals a scalar multiple of the signal) the shortest vector between $\mathbf{x}_i$ and $\hat{a}_i\,\mathbf{s}$ is perpendicular to $\hat{a}_i\,\mathbf{s}$ as shown below. Thus the figure below depicts the case when $\hat{a}_i$ is chosen so as to minimize the size of $\mathbf{e} = \mathbf{x}_i - \hat{a}_i\mathbf{s}$ which is directly analogous to the minimization of the Euclidean norm squared discussed above. Since $\mathbf{e} \perp \hat{a}_i\mathbf{s}$, $\mathbf{e} \bullet \hat{a}_i\mathbf{s} = 0$ where $\bullet$ denotes the inner product (or dot product). We assume $\hat{a}_i \neq 0$ in the following.

$$\mathbf{e} \bullet \hat{a}_i\mathbf{s} = (\mathbf{x}_i - \hat{a}_i\mathbf{s}) \bullet \hat{a}_i\mathbf{s} = 0$$

$$\mathbf{x}_i \bullet \mathbf{s} = \hat{a}_i\mathbf{s} \bullet \mathbf{s}$$

$$\therefore \hat{a}_i = \frac{\mathbf{x}_i \bullet \mathbf{s}}{\mathbf{s} \bullet \mathbf{s}} = \frac{\sum_n x_i[n]\, s[n]}{\sum_n s^2[n]}$$

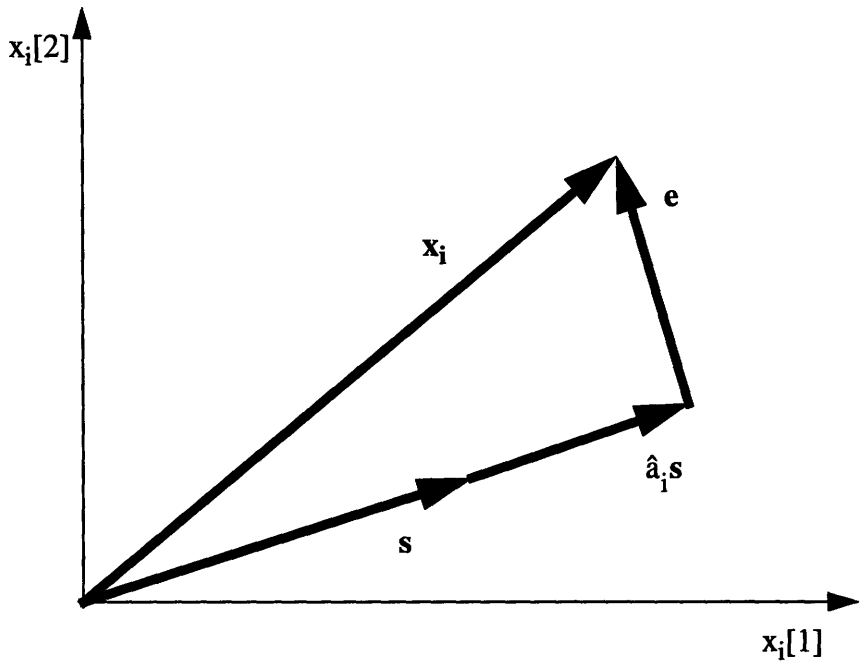

Figure B.1 Graphical Representation of Amplitude Scale Factor Estimation Problem

Estimator Bias:

The form of $\hat{a}_i$ above can be reformulated into the sum of the true amplitude scale factor a_i and an additive error term δ. Consider the appearance of $\frac{\partial \eta}{\partial \hat{a}_i} = 0$ when we substitute for $x_i[n]$:

$$\frac{\partial \eta}{\partial \hat{a}_i} = \sum_n 2\,(x_i[n] - \hat{a}_i\, s[n])\,(-s[n]) = 0$$

$$\frac{\partial \eta}{\partial \hat{a}_i} = \sum_n 2\,(a_i\, s[n] + v_i[n] - \hat{a}_i\, s[n])\,(-s[n]) = 0$$

$$\rightarrow \sum_n (\hat{a}_i - a_i)\, s^2[n] - \sum_n v_i[n]\, s[n] = 0$$

$$\therefore \hat{a}_i = a_i + \delta \qquad \text{where } \delta = \frac{\sum_n v[n]\, s[n]}{\sum_n s^2[n]}$$

Taking the expected value:

$$E\{\hat{a}_i\} = E\{a_i + \delta\} = a_i \text{ (since the noise is zero-mean)}$$

Thus the amplitude scale estimator is unbiased.

* Note that we dropped the subscript on the noise $v_i[n]$ because the only parameters of interest to us are related to the statistics of $v_i[n]$ which are presumed to be the same for all i, allowing us to determine these statistics from independent recordings when no stimulus is presented.

The error term δ can be viewed in numerous equivalent representations. First as the dot product of the noise and the deterministic evoked response divided by the dot product of the deterministic response with itself. This is a measure of the extent that the noise vector lies in the same direction as the deterministic response signal, normalized by the energy in the deterministic signal. Equivalently, the error term can be viewed as the output from a matched filter for the deterministic signal at time zero when the noise is the input, divided by the output from the matched filter at time zero when the deterministic signal is used as the input. This is a measure of how much the noise looks like the deterministic signal, compared to a "perfect" match.

Appendix C: Derivation of Weiner-Type Filter

The Weiner-type filter assumes that the recorded cortical signal consists of the deterministic evoked response and the random process additive wide-sense stationary zero-mean noise which is unaffected by the evoked response, and which does not affect the evoked response. Furthermore, the deterministic signal is nonzero over an interval $[0, T_i]$ only. The system model is as follows:

$$x(t) = s(t) + v(t) \qquad (1)$$

The following derivation first assumes that the variables of interest are functions of continuous time, and then extrapolates the results to the case of discrete time. The following variables are used in this derivation:

$x(t)$ = Individual random process recorded cortical signal used to estimate $s(t)$

$s(t)$ = Deterministic evoked response which is nonzero over $[0, T_i]$ only

$v(t)$ = Zero-mean WSS random process additive noise

$R_{vv}(\tau)$ = Stochastic autocorrelation of $v(t)$

$S_{vv}(\omega)$ = Power spectral density of $v(t)$

$\hat{s}(t)$ = Estimate of $s(t)$ using $x(t)$

$H_w(j\omega)$ = Frequency response of Weiner-type "optimal" filter

$h_w(t)$ = Impulse response of Weiner-type "optimal" filter for estimating $s(t)$ from $x(t)$, i.e. $\hat{s}(t) = h(t) * x(t)$ where $h_w(t) = h(t)$ when the error ε is minimized (defined below)

ε = Error term which the Weiner-type filter is designed to minimize

We seek to minimize the error term ε which is defined as follows:

$$\varepsilon \equiv E\{\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * x(t) - s(t)]^2 dt\} = E\{\int_{-\Delta t}^{(T_1+\Delta t)} [\hat{s}(t) - s(t)]^2 dt\} \qquad (2)$$

The time Δt is chosen so that $h_w(t)$ is negligible for $|t| > \Delta t$, i.e.

$$\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * s(t)]^2 dt \approx \int_{-\infty}^{\infty} [h(t) * s(t)]^2 dt \qquad (3)$$

The above approximation reflects the fact that $h(t)$ is finite in duration, and can be understood by considering the convolution of $h(t)$ and $s(t)$ near the boundaries of $s(t)$. As $s(\tau)$ is multiplied by $h(t-\tau)$ when t is adjusted from zero towards $-\infty$ and from T_1 towards $+\infty$, the product $s(\tau)\,h(t-\tau)$ continues to be nonzero for a finite amount of time due to the overlap between the two functions. We describe the amount of time where there is nonzero overlap before zero and after T_1 by Δt. Note that a stable $h(t)$ cannot be infinite in duration since bounded input bounded output stability requires that $h(t)$ be absolutely integrable.

Additionally, since $\hat{s}(t) = h(t) * x(t) = h(t) * s(t) + h(t) * v(t)$ uses a filter $h(t)$ which is designed to minimize the noise and to make $\hat{s}(t)$ resemble $s(t)$, we

expect that $h(t) * s(t)$ should also resemble $s(t)$. Therefore since $s(t)$ is nonzero over the interval $[0, T_1]$ only, we expect that $h(t) * s(t)$ will approximately be nonzero over the interval $[0, T_1]$ only, where the approximation is carried out in the Δt added boundary in expression (3) above. This approximation assumption in (3) is important because it allows the use of Parseval's theorem in evaluating the expression.

Alternative Description of the Error Expression:

The error expression in (2) can be rewritten as the sum of a filtered noise error term and a deterministic distortion term.

$$\varepsilon \approx E\{\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * v(t)]^2 dt\} + \int_{-\infty}^{\infty} [h(t) * s(t) - s(t)]^2 dt \qquad (4)$$

Proof:

$$\varepsilon \equiv E\{\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * x(t) - s(t)]^2 dt\}$$

$$\varepsilon = E\{\int_{-\Delta t}^{(T_1+\Delta t)} [(h(t) * x(t))^2 - 2s(t)(h(t) * x(t)) + s^2(t)]\, dt\}$$

$$\varepsilon = E\{\int_{-\Delta t}^{(T_1+\Delta t)} [(h(t) * [s(t) + v(t)])^2 - 2s(t)(h(t) * [s(t) + v(t)])]\, dt\}$$

$$+ E\{\int_{-\Delta t}^{(T_1+\Delta t)} s^2(t)\, dt\}$$

$$\varepsilon = E\{\int_{-\Delta t}^{(T_1+\Delta t)} [(h(t) * s(t))^2 + (h(t) * v(t))^2]\, dt\}$$

$$+ E\{\int_{-\Delta t}^{(T_1+\Delta t)} [2(h(t) * s(t))(h(t) * v(t))]\, dt\}$$

$$- E\{\int_{-\Delta t}^{(T_1+\Delta t)} [2s(t)(h(t) * s(t)) + 2s(t)(h(t) * v(t)) - s^2(t)]\, dt\}$$

$$\varepsilon = E\{\int_{-\Delta t}^{(T_1+\Delta t)} (h(t) * v(t))^2 dt\}$$

$$+ E\{\int_{-\Delta t}^{(T_1+\Delta t)} [(h(t) * s(t))^2 - 2s(t)(h(t) * s(t)) + s^2(t)]\, dt\}$$

$$+ E\{\int_{-\Delta t}^{(T_1+\Delta t)} 2(h(t) * v(t))[(h(t) * s(t)) - s(t)]\, dt\}$$

$$\varepsilon = E\{\int_{-\Delta t}^{(T_1+\Delta t)} (h(t) * v(t))^2 dt\} + \int_{-\Delta t}^{(T_1+\Delta t)} [(h(t) * s(t)) - s(t)]^2 dt$$

$$+ 2\int_{-\Delta t}^{(T_1+\Delta t)} E\{h(t) * v(t)\}[(h(t) * s(t)) - s(t)]\, dt$$

$$E\{h(t) * v(t)\} = E\{\int_{-\infty}^{\infty} v(\tau) h(t-\tau)\, dt\} = \int_{-\infty}^{\infty} E\{v(\tau)\} h(t-\tau)\, dt = 0$$

since $v(t)$ is zero-mean from the problem statement assumption.

$$\int_{-\Delta t}^{(T_1+\Delta t)} [\,(h(t) * s(t)) - s(t)\,]^2 dt \approx \int_{-\infty}^{\infty} [\,(h(t) * s(t)) - s(t)\,]^2 dt$$

because $\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * s(t)]^2 dt \approx \int_{-\infty}^{\infty} [h(t) * s(t)]^2 dt$ and

$$\int_{-\Delta t}^{(T_1+\Delta t)} s(t)\,dt \approx \int_{-\infty}^{\infty} s(t)\,dt \text{ since } s(t) \text{ is nonzero over } [0, T_1] \text{ only}$$

$$\therefore \varepsilon \approx E\{\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * v(t)]^2 dt\} + \int_{-\infty}^{\infty} [h(t) * s(t) - s(t)]^2 dt$$

Frequency Domain Representation of the Error Expression:

The error expression in (4) can be rewritten in the frequency domain using Parseval's theorem. Letting $n(t) = h(t) * v(t)$ and $T = T_1 + 2\Delta t$, and noting that $n(t)$ is WSS since $v(t)$ is WSS and $h(t)$ is the impulse response of an LTI system:

$$E\{\int_{-\Delta t}^{(T_1+\Delta t)} [h(t) * v(t)]^2 dt\} = \int_{-\Delta t}^{(T_1+\Delta t)} E\{n^2(t)\}\,dt$$

$$= \int_{-\Delta t}^{(T_1+\Delta t)} R_{nn}(0)\,dt = TR_{nn}(0) = \frac{T}{2\pi}\int_{-\infty}^{\infty} S_{nn}(\omega)\,d\omega$$

$$= \frac{T}{2\pi}\int_{-\infty}^{\infty} |H(j\omega)|^2 S_{vv}(\omega)\,d\omega$$

$$\int_{-\infty}^{\infty} [h(t) * s(t) - s(t)]^2 dt = \frac{1}{2\pi}\int_{-\infty}^{\infty} |H(j\omega)\,S(j\omega) - S(j\omega)|^2 d\omega$$

$$= \frac{1}{2\pi}\int_{-\infty}^{\infty} |H(j\omega) - 1|^2 |S(j\omega)|^2 d\omega$$

$$\therefore \varepsilon \approx \frac{T}{2\pi}\int_{-\infty}^{\infty} |H(j\omega)|^2 S_{vv}(\omega)\,d\omega + \frac{1}{2\pi}\int_{-\infty}^{\infty} |H(j\omega) - 1|^2 |S(j\omega)|^2 d\omega \qquad (5)$$

Frequency Response of Weiner-type Filter:

All of the terms being integrated in (5) are nonnegative, and thus ε is minimized by minimizing the sum of these terms, or equivalently by minimizing the sum of the square root of these terms which simplifies calculations. Therefore $H_w(j\omega)$ is found by choosing real and imaginary parts of $H(j\omega)$ for each ω such that we minimize the expression $T^{1/2}|H(j\omega)|\,(S_{vv}(\omega))^{1/2} + |H(j\omega) - 1||S(j\omega)|$.

Consider the geometric interpretation of the above minimization problem, where $H(j\omega)$ is written as the sum of its real and imaginary parts, $H(j\omega) \equiv H_r(j\omega) + jH_i(j\omega)$, and is described in the complex plane by the vector $\mathbf{H} \equiv \mathbf{H}_r + j\mathbf{H}_i$. Define also the vectors $\mathbf{r}_1 \equiv \mathbf{H}$ and $\mathbf{r}_2 \equiv \mathbf{H} - 1$, and weights $w_1 \equiv T^{1/2}(S_{vv}(\omega))^{1/2}$ and $w_2 \equiv |S(j\omega)|$. We thus wish to minimize the weighted sum of the lengths of $\mathbf{r}_1$ and $\mathbf{r}_2$, i.e. minimize $w_1\|\mathbf{r}_1\| + w_2\|\mathbf{r}_2\|$. Note that both the weights are real and nonnegative.

Further define the error vector $\xi \equiv w_1 \mathbf{r}_1 + w_2 \mathbf{r}_2$ whose length we wish to minimize. Using $\mathbf{r}_1 \equiv \mathbf{H}$ and $\mathbf{r}_2 \equiv \mathbf{H} - 1$, we obtain $\mathbf{r}_1 = \mathbf{r}_2 + 1$ and the error vector becomes the following:

$$\xi = w_1 (\mathbf{r}_2 + 1) + w_2 \mathbf{r}_2 = (w_1 + w_2) \mathbf{r}_2 + w_1 \qquad (6)$$

The diagram below depicts the geometric minimization problem where the circle represents the locus of points when the length of $\mathbf{r}_2$ is constant given the weights w_1 and w_2. We seek some choice of $\mathbf{r}_2$ anywhere in the complex plane yielding the smallest length of ξ, which occurs when ξ and $\mathbf{r}_2$ are coincident with the real axis. Thus since $\mathbf{r}_2$ is real, $\mathbf{H} \equiv \mathbf{H}_r + j\mathbf{H}_i$ is real, which implies $H(j\omega) = H_w(j\omega)$ is real for all ω. Furthermore this minimization implies that $\|\xi_{min}\| \le w_1$, which results in the following constraints on the frequency response vector $\mathbf{H}$, where we use $\mathbf{r}_{2min}$ to denote the value of $\mathbf{r}_2$ when $\xi = \xi_{min}$:

$$\xi_{min} = (w_1 + w_2) \mathbf{r}_{2min} + w_1 = (w_1 + w_2)(\mathbf{H}_r - 1) + w_1 = (w_1 + w_2) \mathbf{H}_r - w_2$$

$$\|\xi_{min}\| = (w_1 + w_2) \|\mathbf{H}_r\| - w_2 \le w_1$$

$$\therefore \|\mathbf{H}_r\| \le 1$$

In sum, the geometric minimization indicates that $H_w(j\omega)$ is purely real and its magnitude is less than or equal to unity.

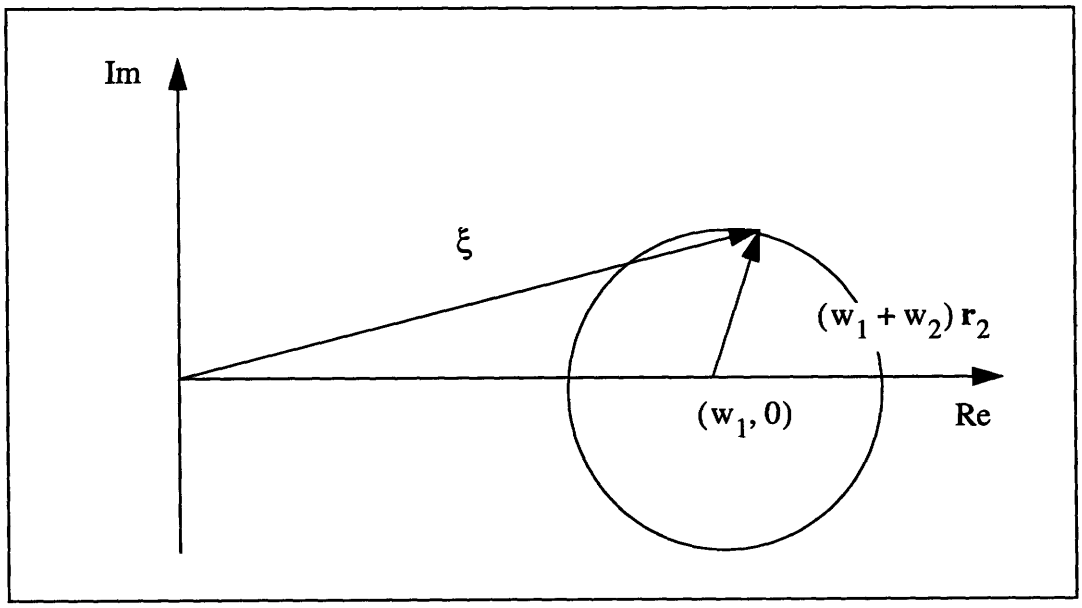

Figure C.1 Geometric Minimization Problem

Given the above constraints on $H_w(j\omega)$ we need only consider real $H(j\omega)$ in the error expression found from (5), $\zeta = T|H(j\omega)|^2 S_{vv}(\omega) + |H(j\omega) - 1|^2 |S(j\omega)|^2$, which reduces to minimization of the following:

$$\zeta = T(H(j\omega))^2 S_{vv}(\omega) + (H(j\omega) - 1)^2 |S(j\omega)|^2$$

Taking the derivative of ζ with respect to $H(j\omega)$ for a particular ω and defining $H(j\omega) = H_w(j\omega)$ when this minimum occurs:

$$\frac{\partial \zeta}{\partial H(j\omega)} = 2TH(j\omega)S_{vv}(\omega) + 2(H(j\omega) - 1)|S(j\omega)|^2$$

$$\therefore H_w(j\omega) = \frac{|S(j\omega)|^2}{|S(j\omega)|^2 + TS_{vv}(\omega)} \tag{7}$$

We verify that this frequency response actually yields a minimum of ζ by taking the second derivative and noting that it is greater than zero:

$$\frac{\partial^2 \zeta}{\partial H^2(j\omega)} = 2TS_{vv}(\omega) + 2|S(j\omega)|^2 \geq 0 \quad \forall\omega$$

Note that the restriction $|H_w(j\omega)| \leq 1$ is satisfied for all values of $|S(j\omega)|^2$, T, and $S_{vv}(\omega)$, and because all the terms on the right side in (7) are nonnegative, we can also conclude that $0 \leq H_w(j\omega) \leq 1$.

Note on Units of Energy Spectral Density and Power Spectral Density:

The units of the energy spectral density $|S(j\omega)|^2$ for the deterministic signal $s(t)$ and of the power spectral density $S_{vv}(\omega)$ for the stochastic noise process $v(t)$ can readily be determined by letting $s(t)$ and $v(t)$ be voltage signals across a 1 ohm resistor. Using convolution, the units of the deterministic autocorrelation of $s(t)$ are V^2-s, and taking the Fourier transform, the units of $|S(j\omega)|^2$ are V^2-s^2. Using the relation $\text{Power} = \frac{\partial(\text{energy})}{\partial(\text{time})}$ and noting that V^2/R are units of power, the units of $|S(j\omega)|^2$ can be seen as energy $(V^2\text{-s/R})$ per unit bandwidth, where bandwidth is in radians per second. Similarly, using expected values the units of the stochastic autocorrelation of $v(t)$ are V^2, and taking the Fourier transform the units of $S_{vv}(\omega)$ are V^2-s. Thus the units of the power spectral density can be seen as power V^2/R per unit bandwidth, where bandwidth is in radians per second as before. The factor T in the denominator of (7) adjusts the units of $S_{vv}(\omega)$ making them consistent with the units of $|S(j\omega)|^2$, and reflects the fundamental difference between energy and power spectral densities.

Formulation of Weiner-type Frequency Response for Discrete Time Signals:

Using the discrete time signals $x[n]$, $s[n]$, and $v[n]$ in place of the continuous time signals $x(t)$, $s(t)$, and $v(t)$ respectively requires a slight modification to (7). Here we must replace the continuous time angular frequency ω by the discrete time angular frequency Ω which yields the following:

$$H_w(e^{j\Omega}) = \frac{|S(e^{j\Omega})|^2}{|S(e^{j\Omega})|^2 + TS_{vv}(\Omega)} \tag{8}$$

Appendix D: Notation and Formula Summary

$x_i[n]$ = individual recorded waveform, i = 1, 2,..., I, n = 1, 2,..., N

$s[n]$ = deterministic unknown evoked response, n = 1, 2,..., N

$\hat{s}[n]$ = estimate of deterministic evoked response, n = 1, 2,..., N

$v_i[n]$ = WSS random process additive noise associated with $x_i[n]$, assumed unaffected by the presence of $s[n]$, i = 1, 2,..., I, n = 1, 2,..., N

$\hat{v}_i[n]$ = estimate of additive noise associated with $x_i[n]$, i = 1, 2,..., I, n = 1, 2,..., N

a_i = random variable amplitude scale factor (Model 2), i = 1, 2,..., I

$\hat{a}_i$ = initial estimate of the amplitude scale factor (Model 2), i = 1, 2,..., I

$\bar{x}[n]$ = overall mean waveform, n = 1, 2,..., N

$\bar{x}_{ave}$ = time average of overall mean waveform

W_l = lower window time (ms)

W_u = upper window time (ms)

$r_k[n]$ = random subset mean (RSM) waveform, k = 1, 2,..., N_w

$\overline{r_k}$ = time average of an RSM waveform over the window from W_l to W_u

N_w = number of random subset mean waveforms

N_i = number of individual waveforms used to form each RSM waveform

P_k = peak-to-trough amplitude of RSM waveform, k = 1, 2,..., N_w

P = peak-to-trough amplitude of the overall mean waveform

F_k = root-mean-square (RMS) amplitude about the time average, k = 1, 2,..., N_w

F_o = RMS amplitude about the time average of the overall mean waveform

$\bar{F}$ = mean of the F_k taken over k

S^2 = sample variance of the F_k or P_k

ε^2 = mean squared error of the F_k or P_k

$h_w[n]$ = impulse response of Weiner-type filter for estimating $s[n]$ from $x_i[n]$

$H_w(e^{j\Omega})$ = frequency response of Weiner-type filter for estimating $s[n]$ from $x_i[n]$

$|S(e^{j\Omega})|^2$ = energy spectral density of $s[n]$

$S_{vv}(\Omega)$ = power spectral density of $v[n]$

T = time scale factor which is dependent on the duration of $s[n]$ and the nonzero overlap between $h_w[n]$ and $s[n]$

$\hat{R}_{vv}[k]$ = sample autocorrelation of the random process v [n] represented by the vector **v**

x = vector representation of x [n]

s = vector representation of s [n]

$\mathbf{s}_n$ = unit length version of **s**

v = vector representation of v [n]

w = zero mean version of **s**

y = projection of **w** on $\mathbf{s}_n$

Λ_{ww} = autocovariance matrix of **w**

Λ_{vv} = autocovariance matrix of **v**

Λ_{yy} = autocovariance matrix of **y**

R = ratio used in amplitude scale model verification test

λ_i = i^{th} eigenvalue of Λ_{ww}

ϑ_i = i^{th} eigenvector of Λ_{ww} corresponding to λ_i

Υ_i = i^{th} eigenvalue of Λ_{vv}

ψ_i = i^{th} eigenvector of Λ_{vv} corresponding to Υ_i

σ_a^2 = variance of the amplitude scale factor

$\mathbf{P}_k$ = primary signal vector for the active noise cancellation system

p_k = primary signal sample at time k (identical to p [k] or x [k] in our case)

y_k = output of adaptive filter at time k

$\mathbf{R}_k$ = reference signal vector for the active noise cancellation system

$\mathbf{W}_k$ = weight vector for the active noise cancellation system

μ = convergence parameter for the active noise cancellation system

ε_k = error signal for the active noise cancellation system

$$x_i[n] = s[n] + v_i[n] \quad \text{(Model 1)}$$

$$x_i[n] = a_i\, s[n] + v_i[n] \quad \text{(Model 2)}$$

$$\bar{x}[n] = \frac{1}{I}\sum_{i=1}^{I} x_i[n] \quad (\forall n) \text{ (Overall mean waveform)}$$

$$\bar{x}_{ave} = \frac{1}{(W_u - W_l + 1)}\sum_{n=W_l}^{W_u} \bar{x}[n] \quad \text{(Time average of } \bar{x}[n] \text{ over the window)}$$

$$r_k[n] = \frac{1}{N_i}\sum_{j=1}^{N_i} x_j[n] \quad (x_j[n] \text{ chosen randomly}) \ (\forall n)$$

(Random subset mean waveform)

$$P_k = \max\{r_k[n]\} - \min\{r_k[n]\}\Big|_{\{W_l, W_u\}}$$

(Peak-to-Trough amplitude of RSM waveform)

$$P = \max\{\bar{x}[n]\} - \min\{\bar{x}[n]\}\Big|_{\{W_l, W_u\}}$$

(Peak-to-Trough amplitude of the overall mean waveform)

$$F_k = \sqrt{\frac{1}{W_u - W_l + 1}\sum_{n=W_l}^{W_u}\left(r_k[n] - \bar{r_k}\right)^2}$$

(Root-Mean-Square (RMS) amplitude about the time average)

$$\bar{r_k} = \frac{1}{(W_u - W_l + 1)}\sum_{n=W_l}^{W_u} r_k[n]$$

(Time average of an RSM waveform over the window)

$$F_o = \sqrt{\frac{1}{W_u - W_l + 1}\sum_{n=W_l}^{W_u}(\bar{x}[n] - \bar{x}_{ave})^2}$$

(RMS amplitude about the time average of the overall mean waveform)

$$\bar{x}_{ave} = \frac{1}{N}\sum_{n=1}^{N} \bar{x}[n] \quad \text{(Time average of overall mean waveform)}$$

$$\bar{F} = \frac{1}{N_w}\sum_{k=1}^{N_w} F_k \quad \text{(Mean of the } F_k \text{ (if RSM waveforms used))}$$

* $\bar{P}$ is defined equivalently

$$S^2 = \frac{1}{N_w - 1}\sum_{k=1}^{N_w}\left(F_k - \bar{F}\right)^2 \quad \text{(Sample variance of the } F_k \text{)}$$

* with a corresponding definition for the P_k

$$\varepsilon^2 = \frac{1}{N_w} \sum_{k=1}^{N_w} (F_k - F)^2$$ (Mean squared error of the F_k)

* with a corresponding definition for the P_k

$$\hat{a}_i = \frac{\sum_n x_i[n]\ s[n]}{\sum_n s^2[n]}$$ (Estimator of the amplitude scale factor)

$$H_w(e^{j\Omega}) = \frac{|S(e^{j\Omega})|^2}{|S(e^{j\Omega})|^2 + TS_{vv}(\Omega)}$$

(Weiner-type filter for estimating s[n] from x[n])

$$\hat{R}_{vv}[k] = \frac{1}{N-|k|} \sum_{n=0}^{N-|k|-1} v[n]\, v[n-|k|]$$

(Sample Autocorrelation of v[n])

$$\mathbf{y}_i \equiv \text{proj}_{s_n} \mathbf{w}_i = \mathbf{s}_n \mathbf{s}_n^T \mathbf{w}_i$$

(Random process used in modelling section)

$$\mathbf{s}_n \equiv \frac{\mathbf{s}}{\|\mathbf{s}\|}$$ (Unit length version of $\mathbf{s}$)

$$R \equiv \frac{E\{\mathbf{y}^T\mathbf{y}\}}{\text{trace}\{\Lambda_{\mathbf{ww}}\}}$$ (Ratio used in model verification test)

$$\mathbf{W}_{k+1} = \mathbf{W}_k + 2\mu\varepsilon_k\mathbf{R}_k$$ (LMS algorithm)

$$(RMS)_{Norm} \equiv \left\{ \frac{\frac{1}{N}\sum_{n=1}^{N} (q[n] - \bar{x}[n])^2}{\frac{1}{N}\sum_{n=1}^{N} \bar{x}^2[n]} \right\}^{1/2}$$ (Normalized RMS Error)

References

[1] E. B. Brooks, K. H. Chiappa, "A comparison of clinical neuro-opthalmological findings and pattern shift visual evoked potentials in multiple sclerosis." *Clinical Applications of Evoked Potentials in Neurology.* Raven Press, 1982.

[2] K. Chiappa, *Evoked Potentials in Clinical Medicine*. Raven Press, 1990.

[3] C. Davila, A. Abaye, A. Khtanzad, "Estimation of Single Sweep Steady-State Visual Evoked Potentials by Adaptive Line Enhancement." *IEEE Transactions on Biomedical Engineering*, 41(2): 197-200, February 1994.

[4] J. P. C. DeWeerd, "Facts and fancies about a posteriori "Weiner" filtering," *Biomedical Engineering*, 28: 252-257, 1981.

[5] R. Dobie, "Objective Response Detection." *Ear and Hearing*, Volume 14, Number 1, 1993.

[6] D. J. Doyle, "Some comments on the use of Weiner filtering for estimation of evoked potentials," *Electroencephalography and Clinical Neurophysiology*, 38: 533-534, 1975.

[7] A. Drake, *Fundamentals of Applied Probability Theory.* McGraw-Hill, 1967.

[8] F. H. Duffy, V.G. Iyer and W.W. Surwillo, *Clinical Electroencephalography and Topographic Brain Mapping*. Springer-Verlag, 1989.

[9] L. J. Eriksson, M. C. Allie, and C. D. Bremigan, "Active noise control using adaptive digital signal processing," Proceedings ICASSP (New York), 1988, pp. 2594-2597.

[10] M. Furst and A. Blau, "Optimal *a Posteriori* Time Domain Filter for Average Evoked Potentials." *IEEE Transactions on Biomedical Engineering*, 38(9): 827-833, September 1991.

[11] S. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory.* Prentice-Hall, 1993.

[12] S. Kay, "Spectral Estimation." *Advanced Topics in Signal Processing*, chapter 2. Prentice Hall, 1988.

[13] D. W. Klass, D. D. Daly (Editors), *Current Practice of clinical Electroencephalography*, Raven Press, 1979.

[14] D. H. Lange, H. Pratt, and G. F. Inbar, "Segmented Matched Filtering of Single Event Related Evoked Potentials." *IEEE Transactions on Biomedical Engineering*, 42(3), March 1995.

[15] G. P. Madhavan, "Comments on 'Adaptive Filtering of Evoked Potentials.'" *IEEE Transactions on Biomedical Engineering*, 35(4), April 1988.

[16] J. McClellan, "Parametric Signal Modelling," *Advanced Topics in Signal Processing*, chapter 1. Prentice Hall, 1988.

[17] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing*. Prentice-Hall, 1989.

[18] A. Oppenheim, "Single-Sensor Active Noise Cancellation." *IEEE Transactions on Speech and Audio Processing*, 2(2), April 1994.

[19] A. Owens, T. Denison, et al, "Muli-electrode array for measuring evoked potentials from surface of ferret primary auditory cortex." *J. Neurosci. Methods*, 50:209-220.

[20] J. Rizzo. Personal communication.

[21] H. Stark and J. W. Woods, *Probability, Random Processes, and Estimation Theory for Engineers*. Prentice Hall, 1994.

[22] J. Stone, w. Barlow, M. Humayun, E. de Juan Jr., and A. Milam. "Morphometric analysis of macular photoreceptors and ganglion cells in retinas with retinitis pigmentosa." Archives of Opthalmology, 110:1634-1639, November 1992.

[23] G. Strang, *Introduction to Applied Mathematics*. Wellesley-Cambridge Press, 1986.

[24] N. V. Thakor, "Adaptive Filtering of Evoked Potentials." *IEEE Transactions on Biomedical Engineering,* 34(1): 6-12, January 1987.

[25] D. O. Walter, "A posteriori "Weiner filtering" of averaged evoked responses," *Electroencephalography and Clinical Neurophysiology Supplemnent*, 27: 61-70, 1969.

[26] J. Wicke, W. Goff, J. Wallace and T. Allison, "On-line statistical detection of average evoked potentials: Application to evoked response audiometry." *Electroencephalogr Clin Neurophysiol* 1978; 44:328:343.

[27] N. Weiner, *Extrapolation, Interpolation and Smoothing of Stationary Time Series, with Engineering Applications*. Wiley, 1949.

[28] B. Widrow and S. D. Stearns, *Adaptive Signal Processing*. Prentice-Hall,1985.

[29] B. Widrow, J. R. Glover, et al, "Adaptive Noise Cancelling: Principles and Applications." *Proceedings of IEEE*, 63(12): 1692-1716, December 1975.

[30] K. B. Yu, C. D. McGillem, "Optimum filters for estimating evoked potential waveforms." *IEEE Transactions in Biomedical Engineering*, 30: 730-736, 1983.