Using Forward Masking Patterns to Predict Imperceptible Information in Speech for Cochlear Implant Subjects

by

Jill M. Desmond

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

_____
Leslie M. Collins, Supervisor

_____
Lianne Cartee

_____
Loren Nolte

Thesis submitted in partial fulfillment of the requirements
for the degree of Master of Science
in the Department of Electrical and Computer Engineering
in the Graduate School of Duke University
2011

Abstract

(Electrical Engineering)

Using Forward Masking Patterns to Predict Imperceptible Information in Speech for Cochlear Implant Subjects

by

Jill M. Desmond

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

_____
Leslie M. Collins, Supervisor

_____
Lianne Cartee

_____
Loren Nolte

An abstract of a thesis submitted in partial fulfillment of the requirements
for the degree of Master of Science
in the Department of Electrical and Computer Engineering
in the Graduate School of Duke University
2011

Abstract

Forward masking is a phenomenon that occurs when one stimulus (masker) elevates the threshold for perception of a subsequent stimulus (probe). In cochlear implant listeners, it has been hypothesized that interactions between electrodes (termed channel interactions) can have a deleterious effect on speech recognition [1], [2]. Forward masking is one measure that has been used to assess channel interactions [1], [3], [4], [5], [6], [7], and it has also been proposed as a mechanism that could be used to assess information that is being presented to but is not received by the user [2], [8]. Determining information that is not received by the user, or is masked, has the potential to provide guidance for the design of new speech processing algorithms that either work to reduce lost information or substitute alternate information.

Nogueira et al., (2005) investigated the potential for using forward masking to estimate masked information and found a limited benefit with information substitution; however, their findings were based on normal hearing psychoacoustic forward masking patterns [8]. Given that psychophysically measured forward masking patterns vary from subject to subject and electrode to electrode (e.g. [1], [3], [6]), measuring subject- and electrode-specific patterns has the potential to provide a more accurate assessment of masked information. In addition, the experimental time required to gather the full set of psychophysical forward masking patterns for all electrodes is far too long for clinical relevance [2]. Recently, it has been suggested that forward masking patterns can be measured physiologically via the electrically evoked com-

pound action potential (ECAP) (e.g. [9]) which may make it possible to measure the information necessary for a subject-specific assessment of masked information in a clinically relevant time frame.

This study utilized ECAP measurements to estimate the forward masking patterns in subjects, and these masking patterns were used to estimate the masked stimuli in speech. The estimates were validated using a speech reception threshold task to assess whether speech recognition is affected by removing "masked" pulses from a subject's stimulation pattern. The results of this study suggest that forward masking patterns measured via ECAPs could potentially determine information that is not being perceived by the user.

The masked stimuli were also used to statistically evaluate the segments of speech that are most vulnerable to masking. These results suggest that the amount of masking that occurs per phoneme depends on the characteristics of the subject-specific masking patterns.

Contents

List of Tables

List of Figures

1

Introduction

Modern cochlear implants contain an array of up to 22 electrodes, or channels. Ideally, each channel would stimulate a unique group of nerve fibers, and the stimulation of one channel would not effect the perception of stimuli on another channel. However, channel interactions do occur (e.g. [3]), and these interactions may negatively impact speech recognition (e.g. [1], [2]). One measure of channel interactions is forward masking (e.g. [3]). Forward masking increases stimulus thresholds, which in turn can prevent a subject from perceiving the associated information. This study aims to assess the lost information that potentially results from the forward masking phenomenon.

Forward masking occurs when a preceding stimulus (masker) fully or partially prevents a neural population from responding to a subsequent stimulus (probe), thus evoking less of a neural response. Forward masking occurs because nerve fibers require a period of time to recover (termed the refractory period) from the original stimulus, and the result is often a lowered sensitivity to a subsequent stimulus because of the lowered neural response. This lower sensitivity can be observed as an increased threshold of perception of the subsequent stimulus (e.g. [10]). If the

probe's presentation level is less than its masked threshold, a subject may be unable to perceive the probe stimulus.

Forward masking is thought to provide insight into auditory excitation patterns by providing a measure of overlap in neural populations stimulated by the masker and the probe [2]. (A greater forward masking threshold shift suggests a greater overlap in neural populations affected by the two stimuli, and thus increased channel interactions.) Figure 1.1 demonstrates examples of neural responses to various masker and probe electrode locations. In this figure, white and black circles represent unstimulated and stimulated electrodes, respectively, and gray ovals represent possible current fields. In the top segment of Figure 1.1, the same electrode is stimulated by the masker and the probe, resulting in the stimulation of the same neural population. In the middle segment of the figure, neighboring electrodes are stimulated, which results in the stimulation of some nerves common to both electrodes. In the bottom segment of Figure 1.1, the large distance between the two stimulated electrodes results in no overlap between the nerves stimulated by the masker and probe.

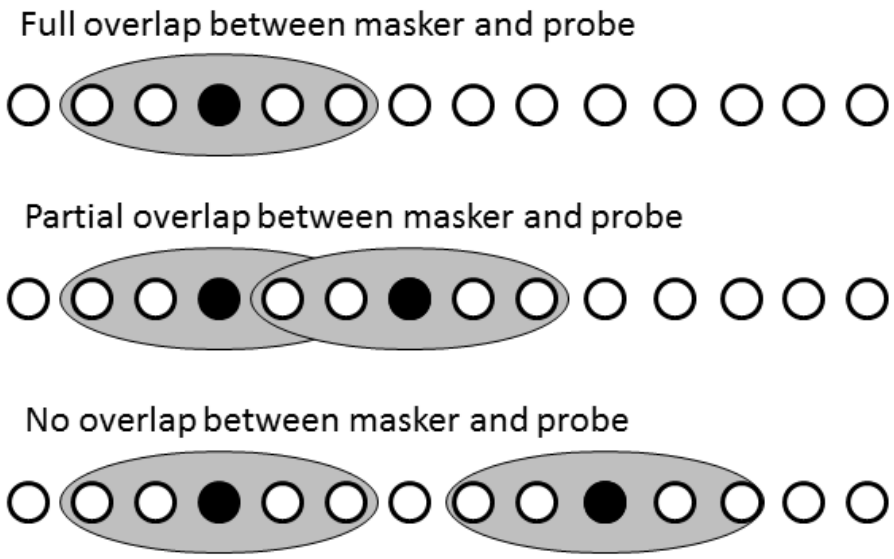

FIGURE 1.1: Diagram of neural responses when stimulating two electrodes. White and black circles represent unstimulated and stimulated electrodes, respectively, and gray ovals represent possible current fields.

While forward masking data provides an opportunity to assess channel interaction and potential for information loss, traditional psychophysical measures of forward masking require significant experimental time [2], and since such data varies across subjects (e.g. [1], [3], [6]) and electrodes (e.g. [2]), psychophysical measures of forward masking are not clinically relevant. Many studies (e.g. [9], [11], [12]) have found a promising correlation between forward masking patterns measured psychophysically and those approximated physiologically using electrically evoked compound action potential (ECAP) measurements. ECAP-based data collection requires much less experimental time, and as such could potentially be clinically relevant. Rather than requiring subject feedback on the interaction of each electrode pairing, which may require about 4 minutes per repetition for each pairing, ECAPs are measured by interpreting neural responses to presented stimuli, which requires about 1 minute per pairing for a total of 100 repetitions. Therefore ECAP measurements, rather than psychophysical measurements, were used to predict masking patterns in this

study, although a set of psychophysical measurements was used to scale the patterns estimated from the ECAP data.

Chapter 2 of this document will include background information on cochlear implants, channel interactions, forward masking, and ECAPs. Chapter 3 will then discuss ECAP measurements and detection techniques. In Chapter 4, this document will review the experimental methods required to approximate forward masking patterns. The applications of these forward masking patterns to speech stimuli will be presented in Chapter 5. Finally, Chapter 6 will discuss conclusions and plans for future work.

2

Background

2.1 Cochlear Implants

In a normal hearing ear, small bones located in the middle ear convert pressure waves received by the outer ear into mechanical vibrations. As a result, fluid inside the cochlea (located in the inner ear) vibrates and causes the basilar membrane to shift. Hair cells that are attached to the basilar membrane bend in response to this shift. As a direct result of being bent, neurotransmitters are released from the hair cells, causing neighboring neurons to fire. A common form of deafness is caused by a loss of hair cells (e.g. [13]). However, if the nerves are intact, they can be excited with electrical stimulation and caused to fire.

The cochlear implant is an approach to address the loss of hair cells. An array of up to 22 electrodes is inserted into the cochlea, and these are used to electrically stimulate the surviving neurons. In a normal hearing ear, the basilar membrane vibrates maximally at different locations given different stimulation frequencies, from low frequencies at one end (the apex) to high frequencies at the opposite end (the base) of the cochlea. Electrodes within the array of a cochlear implant are able to

make use of this tonotopic arrangement. Specifically, each electrode is responsible for transmitting the information corresponding to one frequency band (for a review, see [14]).

All cochlear implants are equipped with a microphone, which sends sound to a speech processor (see Figure 2.1). The speech processor filters the input into the available frequency bands and converts the sound into electrical signals that will be transmitted to the implanted device either transcutaneously via a radio frequency (RF) signal or percutaneously (not shown). From here, the signal is decoded, transformed from a bit stream into currents, and sent to the electrode array inside the cochlea. These electrodes, in turn, stimulate the auditory nerve, which results in the perception of sound.

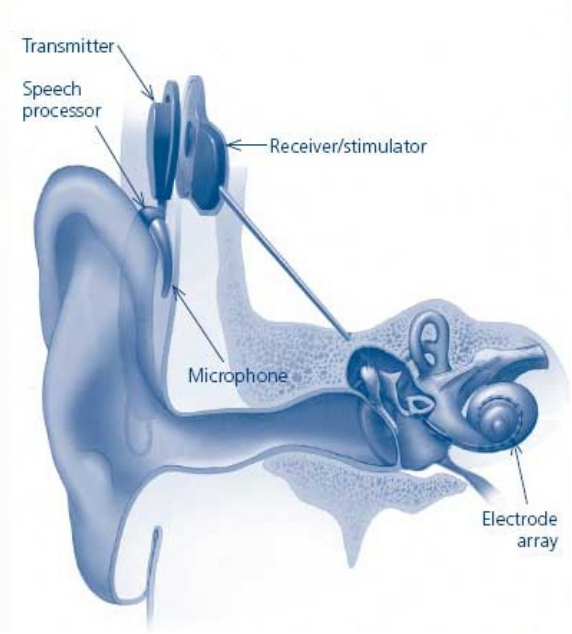

FIGURE 2.1: Diagram of a human ear with a cochlear implant. Sound is transmitted from a microphone to the speech processor. The signal, which is filtered into the available frequency bands and converted into electrical signals, is transmitted transcutaneously or percutaneously (not shown). The signal is then transformed into currents and sent to the electrode array located inside the cochlea, which is responsible for stimulating the auditory nerve (National Institutes of Health, Division of Medical Arts).

Although several speech processing algorithms are currently utilized in the cochlear implant population, the participants in this study use the Advanced Combination Encoder (ACE) strategy. In this strategy (see Figure 2.2), sound travels from the microphone to an array of M bandpass filters, each corresponding to one electrode. The signal segments are then lowpass filtered and rectified in order to extract their envelopes. Next, the electrodes corresponding to the 'N' (less than 'M') frequency band envelopes with the greatest energy in each temporal analysis window are selected for stimulation. Following this step is an amplitude compression stage, which accounts for the reduced dynamic range of electric hearing. This signal then modulates a biphasic current pulse train, which is presented to the electrodes (e.g. [15]).

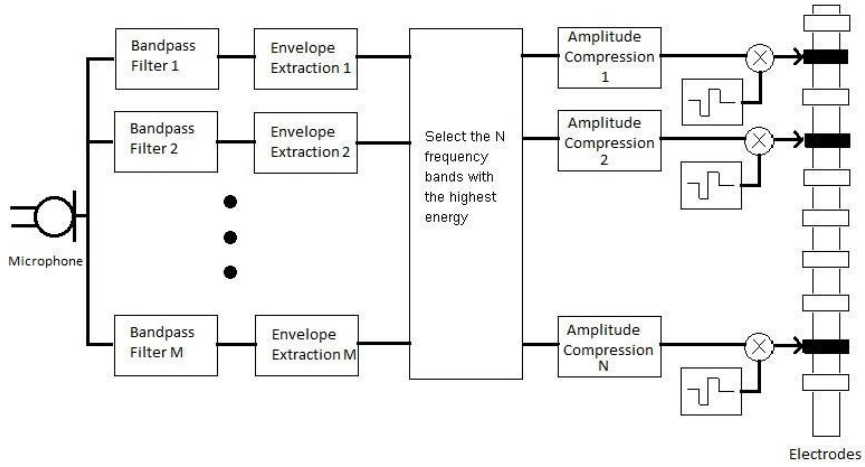

FIGURE 2.2: Diagram outlining the ACE processing strategy. Sound travels from the microphone to an array of bandpass filters. The envelopes are then extracted from the signal segments. Next, the electrodes corresponding to the 'N' frequency band envelopes with the greatest energy in each window are selected for stimulation. Following this step is an amplitude compression stage. Finally, the signal is modulated with a biphasic current pulse train, which is sent to the electrodes (e.g. [15]).

The speech processing algorithms determine the temporal and frequency information that will be presented to the cochlear implant user. This information can be

viewed in plots, termed electrodograms, as in Figure 2.3. Electrodograms are plots of electrode location as a function of time. If an electrode is to be stimulated at a given time, a "tick" will appear with amplitude corresponding to the stimulation current level. Because each electrode corresponds to a different frequency band, electrodograms are simply a method of displaying the frequency and temporal content of the stimulation pattern. High frequencies, which stimulate the base of the cochlea, are represented by lower numbers in the electrode array.

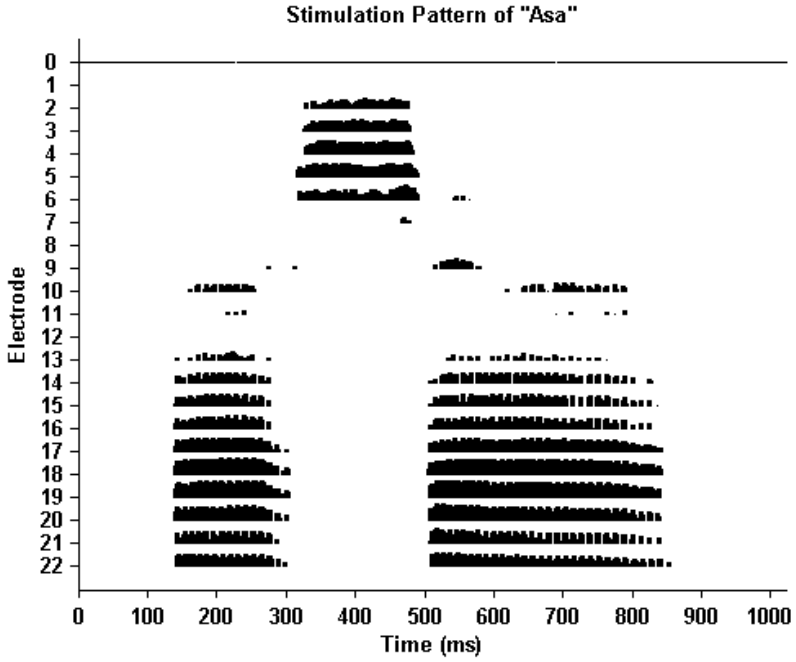

FIGURE 2.3: Stimulation pattern of the speech token "asa." This stimulation pattern, referred to as an electrodogram, demonstrates the frequency and temporal information that is presented to a cochlear implant user during the speech token "asa." Time is plotted on the x axis, and the y axis designates electrode number. If an electrode is to be stimulated at a given time, a "tick" mark, with amplitude corresponding to the stimulation current level, will appear at the corresponding location.

2.2 Channel Interactions

Channel interactions result from the lack of electrode independence within a cochlear implant, and they can occur in the form of current summation and neural interactions. Cochlear implants were designed such that each electrode is assumed to be independent of the others, and the success of cochlear implant strategies requires some level of independence between the channels (e.g. [3]).

The type of channel interaction that concerns the study outlined in this document is that of neural interactions, which can occur because each electrode most likely does not stimulate a unique set of nerve fibers. Forward masking and gap detection are two examples of experimental tasks that can be used to quantify the amount of neural overlap that exists between channels (e.g. [1], [2], [3], [6]). This study focuses on the forward masking approach because physiological approaches to estimating forward masking data suggest that these measurements may be completed in a clinically relevant time frame.

2.3 Forward Masking

Forward masking is a phenomenon experienced by both normal hearing and cochlear implant listeners in which the presence of one stimulus impedes a person's ability to perceive a subsequent stimulus (e.g. [16], [17], [18], [19], [20]). This phenomenon has been measured in single fibers as the reduction of a response to one stimulus (probe) when following a separate stimulus (masker) (e.g. [18]). In normal hearing listeners, this results in a shift of the probe threshold, the level at which the probe stimulus would first become audible (e.g. [10]). Because the cochlea is arranged tonotopically and forward masking results from stimulating the same nerve fibers in succession (e.g. [16], [17], [18]), masker stimuli that are close in frequency to the following probe stimuli result in a greater masking effect, or a greater threshold shift

(e.g. [21]). As the time between the masker and the probe stimuli increases and nerves are able to recover, this threshold shift decreases (e.g. [10], [18]).

Forward masking has been measured to quantify the recovery of nerves from adaptation (e.g. [17], [18]), which results in increased thresholds of nerves fibers following a masker stimulus (e.g. [20]). Shannon (1983) suggests that the adaptation recovery process associated with electrical stimulation is longer than that for normal hearing. The study explains this finding by noting that the amount of adaptation in normal hearing may be reduced due to the stochastic firing of nerves [20]. Additionally, electrical stimulation causes nerve fibers to fire more rapidly [22], and an increased firing rate is reported to result in longer recovery processes [18].

Cochlear implant users are also presented with less frequency and temporal information than normal hearing listeners, further limiting the amount of information they receive. Studies of forward masking in electric hearing could help researchers better understand the connection between signals that are being sent to the implant and signals that are being perceived by the implanted listener. With this in mind, previous studies have measured forward masking in order to study temporal interactions (e.g. [1], [2], [3], [4], [5], [6], [7]) and spatial interactions (e.g. [1], [2], [3], [6]) that occur in cochlear implants.

In an implanted cochlea, forward masking results from stimulating the same underlying neural population, and it can occur when the masker (initial stimulus) and the probe (subsequent stimulus) are located on the same or different electrodes. The degree of forward masking may be a measure of the degree of neural overlap between two electrodes (e.g. [3]). Because neural overlap generally decreases as distance between electrodes increases, the threshold shift of a probe preceded by a masker tends to decrease with increased distance between the channels (e.g. [3]).

Figure 2.4, which is a plot of current level as a function of electrode number, demonstrates a simplified example of forward masking data. The blue stem located

at electrode 11 represents the stimulating (masking) pulse, and the green curve represents the elevated thresholds of neighboring electrodes following the masker stimulation pulse. That is, at each electrode location on the x-axis, the presence of the masker pulse elevates the threshold by the amount of current indicated by the green curve. The threshold shift that results at each electrode quantifies the degree of forward masking present in that channel due to stimulation on electrode 11.

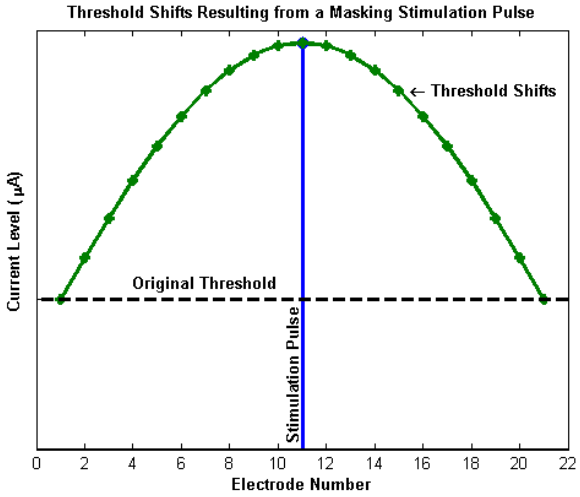

FIGURE 2.4: Simulated forward masking pattern resulting from a stimulation pulse presented to electrode 11. The green curve represents the elevated thresholds of neighboring electrodes following the masker stimulation pulse presented to electrode 11.

When considering how changes in stimulus parameters would affect a forward masking pattern, there are two main stimulus parameters to consider. The first parameter is concerned with how the pattern scales with masker amplitude. That is, the parameter is associated with how the amplitude of the threshold shifts change with masker pulses of different amplitudes. Nelson and Donaldson (2001) found that the amplitude of single-pulse forward masking patterns scale linearly with masker amplitude (relative to the masker's threshold) [23]. The next parameter to consider is the how the masking pattern decays with the time delay between the masker and

the probe. Nelson and Donaldson (2001) found that the rapid recovery process from single-pulse forward masking completes by about 10 ms for most subjects [23].

As previously mentioned, forward masking may be used as a measure of channel interactions. In an ideal cochlear implant with no channel interactions, the green curve in Figure 2.4 would be an impulse at electrode 11, suggesting that thresholds of neighboring electrodes are not increased as a result of stimulation on electrode 11. Since this is the assumption under which cochlear implants operate, forward masking may be tied to a subject's speech recognition performance. Chatterjee and Shannon (1998) discussed a possible correlation between speech recognition and sensitivity to forward masking parameter alterations [1]. Additionally, Throckmorton and Collins (1999) found a correlation between average levels of forward masking and speech recognition, which led them to hypothesize a link between forward masking and imperceptible information. They further hypothesized that this correlation may be due to both spectral and temporal interactions, as quantified by forward masking [2]. Although forward masking appears to be an accurate measure of neural interactions, psychophysically gathering the necessary data requires experimental time that is too long for clinical relevance.

2.4 Electrically Evoked Compound Action Potential (ECAP)

The lack of clinical relevance of psychophysical forward masking techniques was one reason alternate measures of forward masking were sought by the research community. One alternate measure, the ECAP, is a measure of a neural population's response to a stimulus (e.g. [24], [25], [26], [27], [28], [29]). In cochlear implant applications, one electrode stimulates a nerve population with a pulse, and a neighboring electrode records the resulting neural response (e.g. [30]). Because ECAPs require high amplitude pulses, they cannot be measured at all current levels within a subject's dynamic range.

The ability to measure ECAPs in implanted subjects has many possible applications. One such application is the potential to confirm the proper functioning of a cochlear implant and to confirm that the nerve fibers are responding to stimulation [31]. Some studies also found that ECAP thresholds may help clinicians establish an appropriate dynamic range for individual subjects [32], [33], [34]. This application would be extremely helpful when working with young children, who may be unable to provide the feedback required to set the necessary parameters. Additionally, ECAPs measured in different subjects may be indicative of the amount of nerve survival surrounding a given electrode [35], [36], [37]. Others have investigated using ECAPs to measure channel interactions [11], [12], [38] and to measure refractory periods of nerve fibers [39].

For a biphasic current pulse, the ECAP response is characterized by a sharp negative peak followed by a gradual positive peak, and a simplified conceptual rendering is shown following a biphasic stimulation pulse in Figure 2.5. This figure, which plots voltage as a function of time, is not drawn to scale. In reality the ECAP amplitude is very small compared to the stimulus, and it is therefore necessary to remove the artifact created by the stimulus pulse to accurately measure the ECAP response (e.g. [26], [30], [34], [40]).

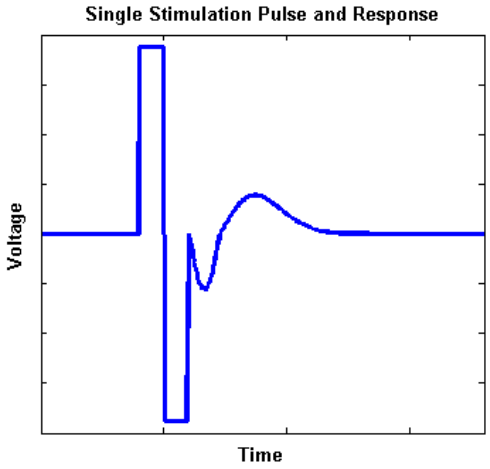

FIGURE 2.5: Biphasic stimulation pulse and resulting ECAP response. The ECAP is characterized by a sharp negative peak followed by a gradual positive peak for a biphasic stimulation pulse. Note: figure is not drawn to scale.

Many artifact reduction techniques have been developed and studied. One such technique, alternating polarity, involves stimulating the nerves with a pulse of one polarity, and repeating the measurement with a pulse of opposite polarity (e.g. [41]). The mean of both measurements is computed such that the two stimulation pulses average to zero, and only the ECAP response remains. This technique suffers inaccuracies, as responses to stimuli of different polarity do not have identical amplitudes and response delays [42]. Another artifact reduction technique subtracts an artifact template from the stimulus artifact and ECAP response. The template is measured using a subthreshold stimulus and its amplitude is scaled to match that of the stimulus from which the ECAP is measured [42]. However, scaling the template to match the stimulating pulse may reduce the signal to noise ratio of the recording [41].

A more robust and widely used artifact removal method, the subtraction method, was used in this study [26], [30]. The subtraction method was developed to utilize the refractoriness of nerve fibers. This algorithm, illustrated in Figure 2.6, subtracts the neural response to a masker followed by a probe (third row) from the sum of the

neural response to a masker (top row) and the neural response to a probe (second row). If the masker and and the probe pulses are presented in rapid succession, the nerves that responded to the masker pulse will be in their absolute refractory period during the presentation of the probe pulse. During the absolute refractory period, nerve fibers are unable to respond to a stimulus. Therefore, if the probe is fully masked in the masker-followed-by-probe response (third row), the measurement recorded during this panel will simply be the two stimulation pulses and the ECAP response to the masker pulse. Once this measurement is subtracted from the sum of the masker-alone and probe-alone measurements (first and second rows of Figure 2.6), the subtraction method will result in the neural response to the probe alone (as seen in the bottom row of the figure) [26], [30].

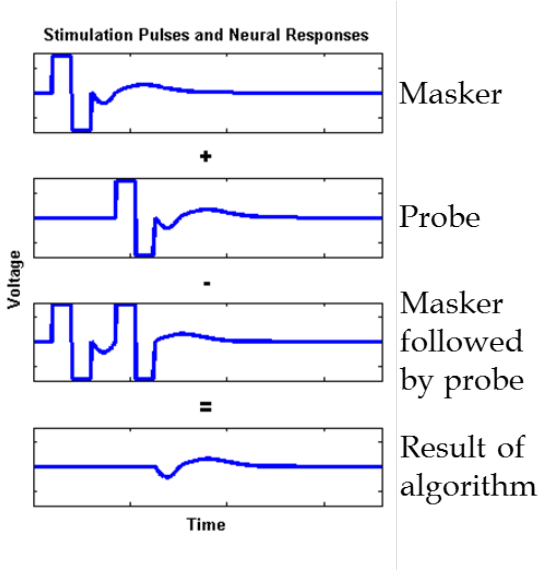

FIGURE 2.6: The subtraction method for ECAP artifact removal. The signals displayed in the top two panels are summed, and the signal in the third panel is subtracted from the summation. If the probe is fully masked in the third panel, the resulting signal (shown in the bottom panel) will be equivalent to the response resulting from the probe alone.

Although the subtraction method was created to measure a neural response to

a probe pulse [26], [30], it has been suggested that the subtraction method can be extended to estimate a forward masking pattern for a given electrode (e.g. [11], [12]). As previously mentioned, if the probe pulse is fully masked when following the masker pulse, the ECAP that results from the subtraction method will simply be the response to the probe pulse. However, as the time between the masker and the probe pulses is increased in the third panel of Figure 2.6, the nerve fibers may begin to recover from the masker stimulation pulse. If the probe response is only partially masked when preceded by a masker pulse, the ECAP resulting from the subtraction method will be the neural response to the probe-alone pulse minus the (partially masked) neural response to the probe when following the masker. Different degrees of masking in the third panel of Figure 2.6 can result in ECAPs of different sizes when the subtraction method is performed. Therefore, ECAP measurements can approximate the amount of forward masking occurring. Specifically, a large amount of masking results in a large ECAP response measured using the subtraction algorithm. Alternatively, if no masking is present in the masker-followed-by-probe scenario, no ECAP will result from this method.

To further illustrate this concept, simulations of three different forward masking scenarios are depicted in Figure 2.7, which again plots voltage as a function of time. The plots in the top row represent the masker-followed-by-probe stimulation pulses and responses, and the corresponding bottom plots display the results obtained using the subtraction method. In the first vertical column, the masker and probe pulses are separated by enough time that the probe experiences no masking, and no ECAP results from the subtraction method. The second column demonstrates a scenario in which the masker and probe pulses are separated by enough time that the probe pulse is only partially masked. The result of the subtraction method is a small ECAP, the response to the probe-alone minus the response to a partially masked probe. The third vertical pair of plots represents the scenario in which the probe pulse is fully

masked by the masking pulse. Therefore a full ECAP results from the subtraction method, as seen in the rightmost bottom plot.

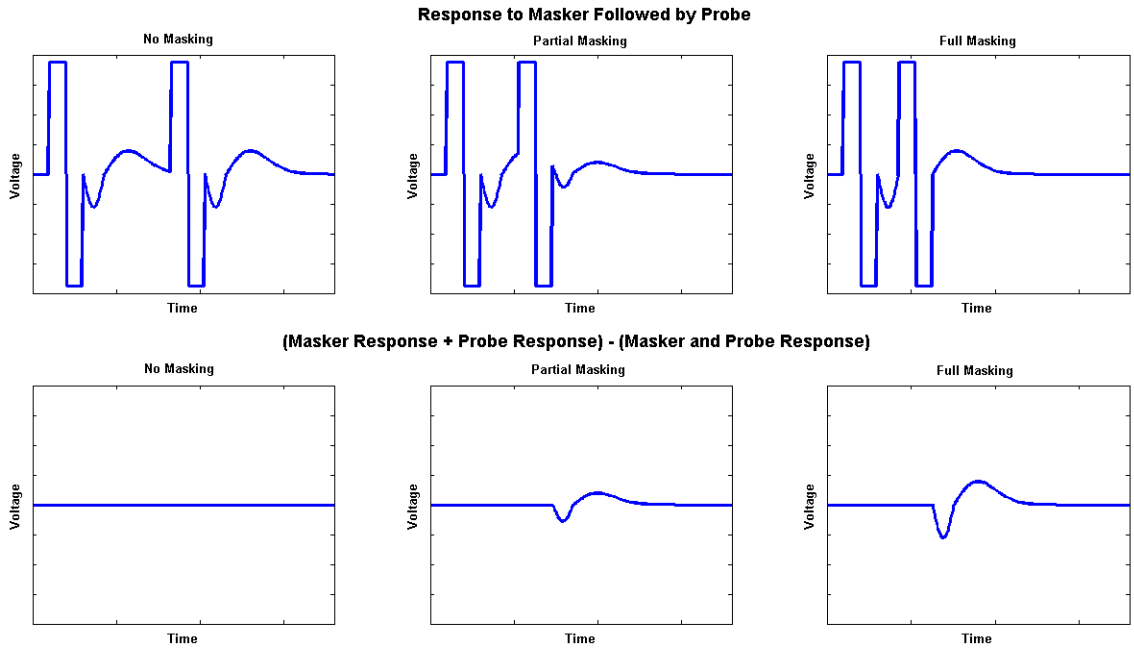

FIGURE 2.7: ECAPs resulting from various masker-probe delays. In the first vertical pair of plots, the masker and probe pulses are separated by enough time that the probe pulse is not masked. The second vertical pair of plots demonstrates a scenario in which the masker and probe pulses are separated by enough time that the probe pulse is only partially masked. The third vertical pair of plots represents the scenario in which the probe pulse is fully masked by the masker pulse.

2.4.1 ECAPs as a Forward Masking Approximation

As mentioned previously, this study aimed to use physiological ECAP measurements to assess forward masking patterns on a subject by subject and electrode by electrode basis. Hughes and Stille (2008) investigated whether psychophysical forward masking patterns and physiological masking patterns are correlated. Although 10 subjects showed strong correlation, 2 subjects showed only moderately strong correlation and 6 subjects showed weak correlation. It was hypothesized that the lack of correlation may have been due to discrepancies in the ECAP and psychophysical forward

masking experimental parameters. One such discrepancy was the use of different masker-probe delays, the ECAP delay being 0.5 ms and the psychophysical delay being 2 ms. Additionally, single-pulse stimuli were used to measure ECAPs, while pulse train stimuli were used to measure psychophysical forward masking. Finally, different current levels were used when recording the two types of measurements. Both psychophysical and physiological data were collected at the same percentage of their respective dynamic ranges. However, since the two stimuli were of different pulse rates, the current levels resulting from the dynamic range percentage differed [43].

In a subsequent study, Hughes and Stille (2009) found a strong correlation between ECAP threshold shifts and psychophysical threshold shifts. In this study, an effort was made to ensure that the ECAP and psychophysical stimuli were as similar as possible, including parameters such as pulse width, interphase gap, pulse rate, and amplitude. Simultaneous stimulation was used for the ECAP and psychophysical measurements, and threshold shift was determined by measuring the threshold of a probe stimulus, followed by measuring the threshold of the probe stimulus in the presence of a (simultaneous) subthreshold "interaction stimulus" [9].

2.4.2 ECAPs as a Measure of Channel Interactions

Although the design of cochlear implants assumes independence between electrodes, channel interactions occur. It has been suggested that these channel interactions could affect speech recognition [1], [2]. Although forward masking is one method of measuring channel interactions (e.g. [1], [2], [3], [4], [5], [6], [7]), masking patterns cannot be measured psychophysically in a clinically relevant time frame. However, a correlation has been found between ECAP threshold shifts and psychophysical threshold shifts [9], suggesting that psychophysical measurements may be approximated physiologically.

3

ECAP Detection

In order to complete this and future studies, hundreds of measurements must be collected and identified as either containing an ECAP or not containing an ECAP. Completing this process visually for every subject would be very time consuming, and determining whether a signal contains an ECAP becomes difficult near threshold, possibly resulting in inconsistencies. In order to address these two concerns, several automatic ECAP classifiers were considered as part of this study.

3.1 Subjects

Three postlingually deaf subjects, all users of Cochlear Corporation's CI24 family of devices, participated in this study. The stimulation mode for each subject was monopolar 1+2 (MP1+2), in which both available extra-cochlear electrodes (labeled as 1 and 2) were used as ground electrodes. Table 3.1 contains the subjects' demographic information.

This study was completed in six to eight sessions, each of which was three to four hours in length. Subject S2 volunteered his time, while the remaining participants were compensated. The Institutional Review Board of Duke University approved the

use of human subjects in the experiments associated with this study.

Table 3.1: Demographic information for implanted subjects

Subject ID	Gender	Age (years)	Age at onset of deafness (years)	Age at implantation (years)	Implant type	Mode of stimulation
S1	F	47	10	41	CI24RE	MP1+2
S2	M	57	48	49	CI24R	MP1+2
S3	M	58	15	53	CI24RE	MP1+2

3.2 Stimuli and Equipment

The ECAP data was collected using a Nucleus Freedom Sound Processor and the Neural Response Telemetry (NRT) system available in the Cochlear Custom Sound EP software [34], [40]. Prior to testing, the subjects' thresholds (Ts) and maximum comfortable levels (MCLs) were measured. These values were used to determine the masker and probe current levels that would be appropriate for testing. The stimuli presented in this study consisted of biphasic pulses with pulse widths of 25 microseconds. A delay of 122 microseconds was included between the stimulation pulses and the onset of recording. Data was recorded using NRT via a recording electrode located two electrodes from the probe in the apical direction. For Subject S2, the measured data was amplified with a 60 dB gain. In order to prevent saturation, a 50 dB gain was used for Subjects S1 and S3. An interpulse interval of 400 microseconds was inserted between the masker and probe pulses. Additionally, each ECAP measurement was the result of averaging 100 recordings

3.3 ECAP Classifiers

As previously mentioned, the large quantity of signals that needed to be evaluated for this study motivated the evaluation of automatic ECAP detectors. Initially,

linear regression was considered as a method to determine ECAP threshold, since that is the method used by the Custom Sound software. In addition other ECAP detectors, namely K nearest neighbors, Fisher's linear discriminant, the generalized likelihood ratio test, an energy detector, and template correlation, were implemented and applied to the ECAP classification problem.

3.3.1 Linear Regression

When determining ECAP thresholds automatically, the Cochlear Custom Sound EP software measures the peak-to-peak amplitudes of ECAPs at various current levels. (A peak-to-peak amplitude is measured from the ECAP's characteristic sharp negative peak to the gradual positive peak. See Figure 3.1 for a reference.) The plot of ECAP amplitude as a function of probe current level results in a sigmoidal curve, to which the software applies linear regression to approximate the ECAP threshold. Figure 3.2 illustrates ECAP data and a linear regression fit, as an illustration of a successful application of this method. The circles represent ECAP amplitudes measured at the given probe current levels, and the straight line illustrates the result of linear regression. (Current level (CL) is defined on Cochlear's 1-255 scale, which relates to current (I) for CI24R implants by the equation $I(\mu A) = 10 \times e^{\frac{CL \times ln(175)}{255}}$, and for CI24RE implants by the equation $I(\mu A) = 17.5 \times 100^{\frac{CL}{255}}$.) The approximate ECAP threshold is determined to be the probe current level at which the ECAP amplitude would equal zero, according to linear regression. In Figure 3.2, the ECAP threshold occurs at a current level of approximately 133.6 current steps.

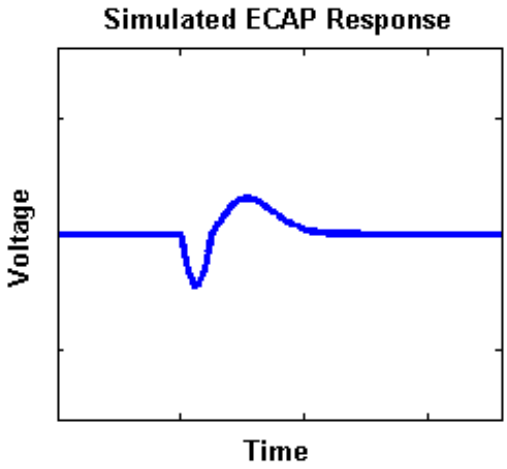

FIGURE 3.1: A simulated ECAP response. The Cochlear Custom Sound EP software measures the peak-to-peak amplitudes of ECAPs at various current levels to determine ECAP thresholds. Note the characteristic negative and positive peaks in this figure.

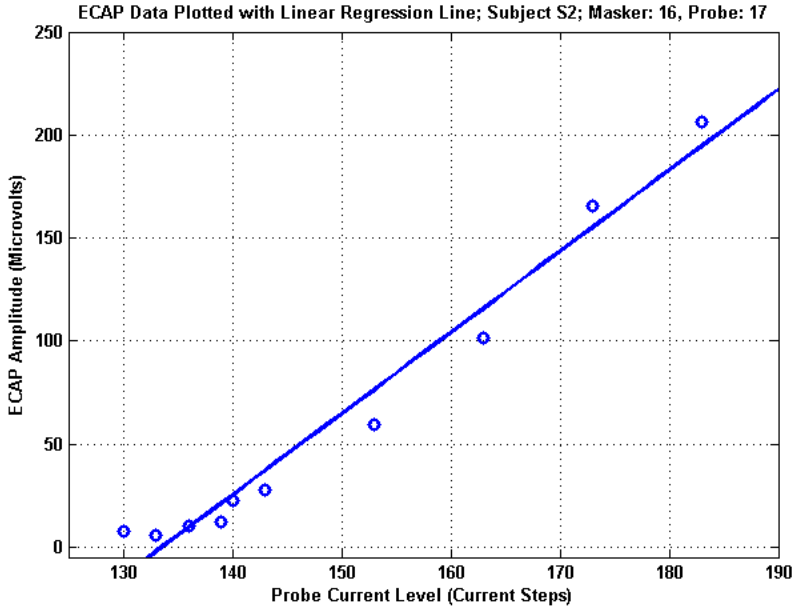

FIGURE 3.2: An approximation of ECAP threshold by means of linear regression. The circles represent ECAP amplitudes measured at the given probe current levels, and the straight line is the result of linear regression. This plot suggests an ECAP threshold at a current level of approximately 133.6.

Although Figure 3.2 demonstrates promising results, linear regression alone can

be unreliable in some cases. Figure 3.3 shows a series of measurements in which ECAPs were not measurable until current levels near the subject's maximum comfortable level were delivered. Because ECAPs were not measurable at lower probe current levels, and higher current levels would have resulted in an uncomfortable sensation, peak-to-peak ECAP amplitudes could only be obtained for a limited range of probe current levels, as shown in Figure 3.3. The expected sigmoidal shape of ECAP amplitude as a function of probe current level could not be obtained, as only measurements at the flat "base" of the anticipated sigmoidal curve could be made. When applied to this incomplete sigmoid, linear regression determined an incorrect ECAP threshold at a current level of approximately 101. The ECAP threshold for this masker-probe pair was found to be at a current level of 188 when determined visually.

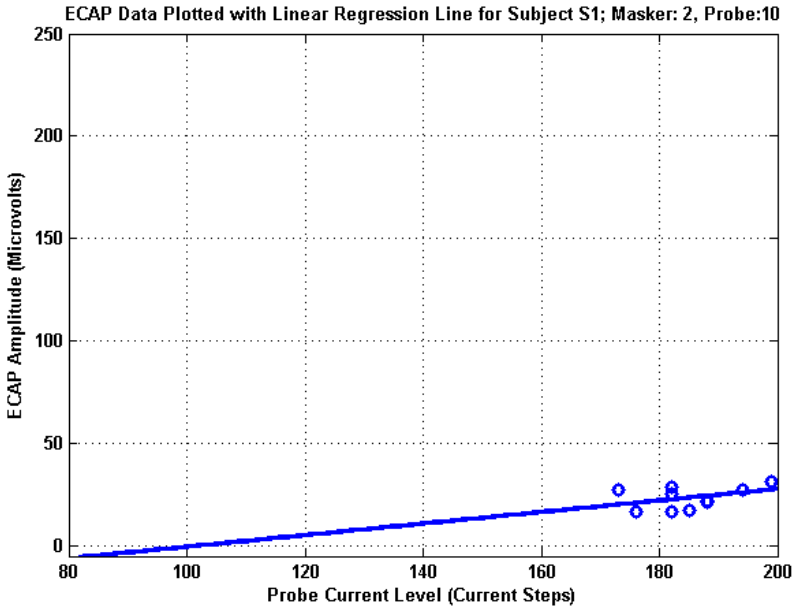

FIGURE 3.3: An inaccurate approximation of ECAP threshold by means of linear regression. The circles represent ECAP amplitudes measured at the given probe current levels, and the straight line shows the result of linear regression.

Even when the linear regression method performs accurately, as in Figure 3.2, it

consistently results in lower ECAP threshold estimates than does visual detection. The sigmoidal shape of the ECAP-amplitude versus probe-current-level curve skews the x-intercept in favor of lower thresholds. This skew of the data can be seen in Figure 3.4, which plots the kernel density estimates of ECAP thresholds determined visually and those determined using linear regression for each subject. The means of two distributions are statistically significantly different for all three subjects (student t-test, $p < 0.05$).

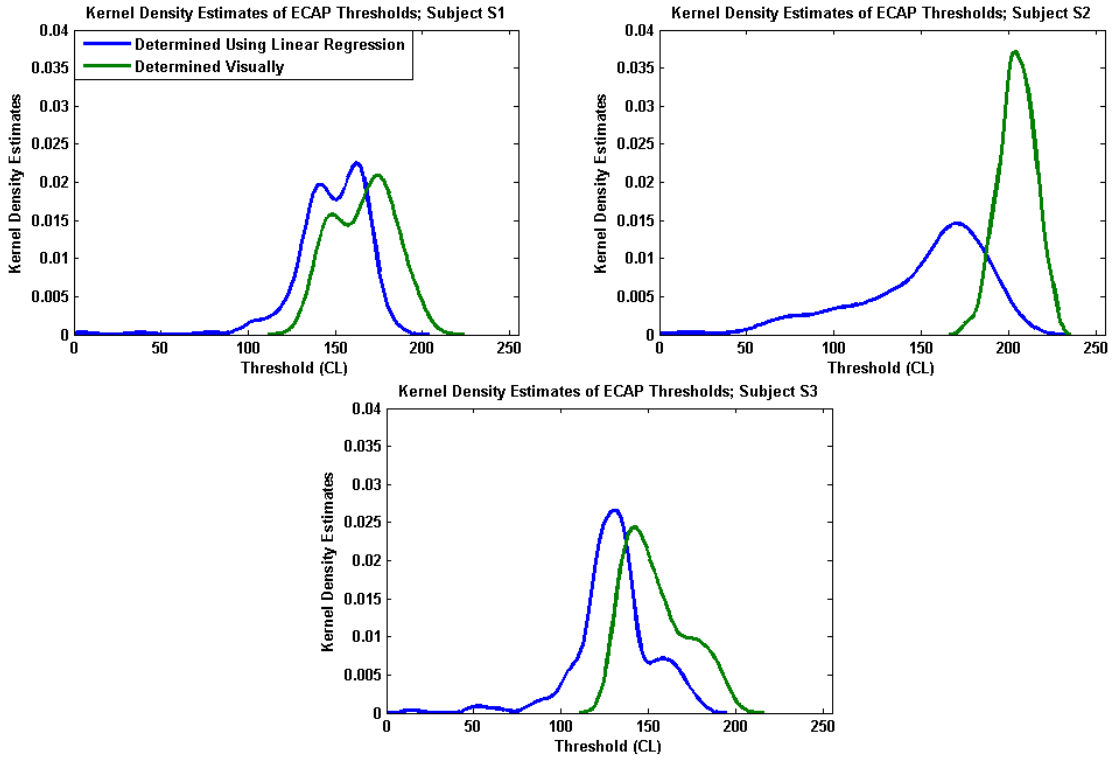

FIGURE 3.4: Kernel density estimates of ECAP thresholds determined using the linear regression method and determined by visual inspection for three subjects. The top left plot demonstrates the results for Subject S1, the top right plot displays the results for Subject S2, and the results for Subject S3 are shown in the bottom plot. Clearly, ECAP thresholds determined by the linear regression method are lower on average than those determined visually.

Figure 3.5 displays the kernel density estimates of ECAP thresholds for all sub-

jects determined both visually and using linear regression. This plot and Figure 3.4 clearly show that the skew of ECAP thresholds determined by linear regression occurs both when considering subject-specific thresholds and when considering thresholds gathered from multiple participants. The means of the distributions containing data across subjects, similarly to those containing subject-specific data, are statistically significantly different (student t-test, $p < 0.05$). If relied upon for approximate forward masking patterns, this skew of ECAP thresholds would result in inaccuracies which, in turn, would result in erroneous predictions of the "masked" pulses that this study is attempting to identify.

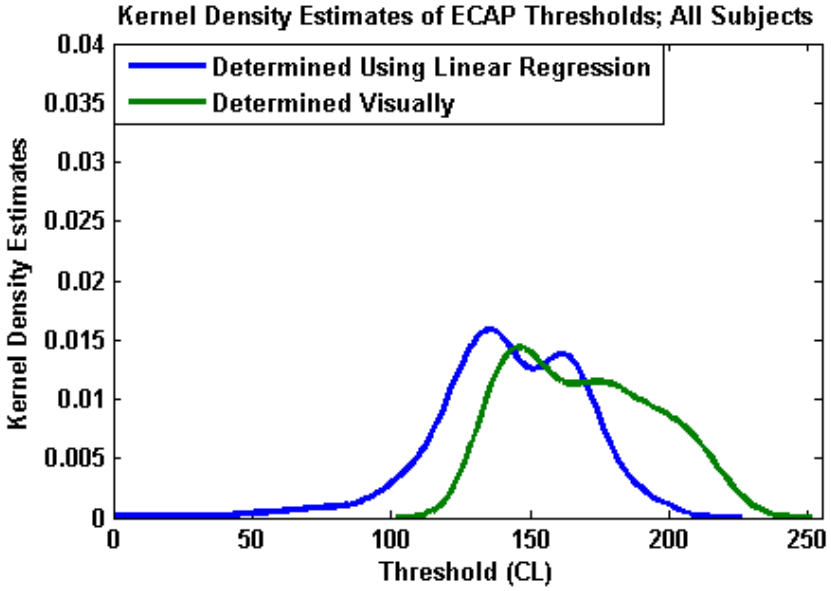

FIGURE 3.5: Kernel density estimates of ECAP thresholds determined using the linear regression method and determined by visual inspection across subjects. Clearly the skew of ECAP thresholds resulting from linear regression estimates is not simply a phenomenon of subject-specific data. Rather, ECAP thresholds determined by linear regression are lower than those determined visually when considering each subject's data separately and when combining data across subjects.

Although linear regression was a logical method to consider, it is sometimes an unreliable ECAP threshold detection method. Both this unreliability and the skew of ECAP thresholds determined using this method motivated investigation

into alternate ECAP detection algorithms.

3.3.2 Feature Extraction

Additional ECAP detection algorithms that were considered utilized specific characteristics, or features, of the ECAP signals. One such feature was the correlation between the input signals and a template. The template was the average of ten signals known to contain strong ECAPs. Other features included the "local minimum" and "local maximum" of the signals. These features were extracted from the time segments in which they are most likely to appear, to account for the fact that the characteristic negative peak of an ECAP typically occurs between 0.2 and 0.4 ms after the onset of the stimulus and the characteristic positive peak usually occurs between 0.6 and 0.8 ms after the stimulus onset [33], [34], [44]. Features in this study also included the global minimum and global maximum values, in an effort to incorporate extrema that deviated from the expected range. The latencies of the local and global minimums and maximums with respect to the onset of recording were also considered. Finally, the energy of each signal was included as a feature, as signals that contain ECAPs will most likely have greater energy than those that do not.

Figure 3.6 demonstrates these features, after being z-scored, for Subjects S1, S2, and S3 (as noted in the individual plot titles). The blue symbols represent features extracted from data that contained an ECAP, and the red symbols represent features gathered from data that did not contain an ECAP (as determined by visual inspection). A great deal of overlap occurs between the features of the two classes of signals, with some features being more informative than others. However, detection algorithms that consider the information from all features combined may be able to more strongly separate signals that contain an ECAP from those that do not.

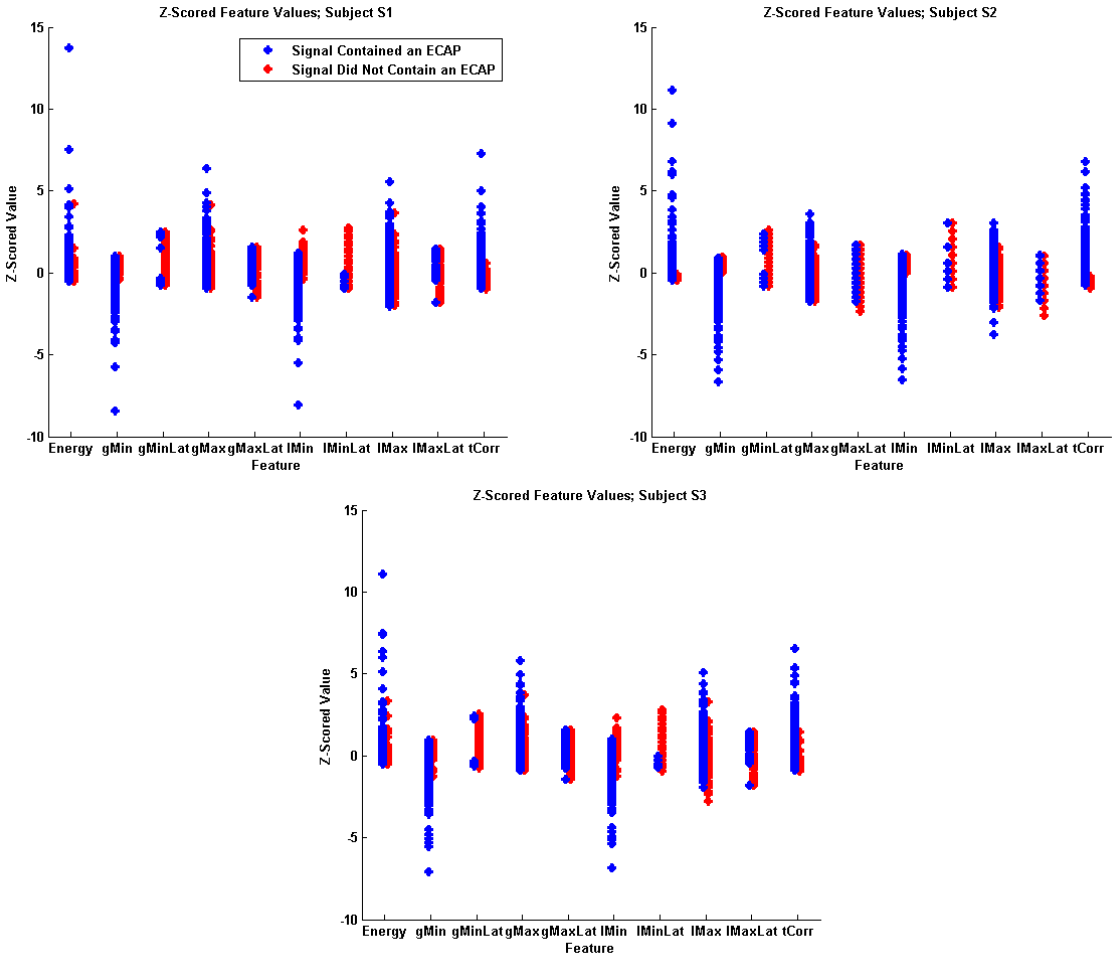

FIGURE 3.6: Z-scored ECAP features for Subjects S1, (top left), S2 (top right), and S3 (bottom). The blue symbols represent features extracted from data that contained an ECAP, and the red symbols represent features gathered from data that did not contain an ECAP (as determined by visual inspection).

In an effort to determine whether the ECAP detectors to be discussed could be applied universally, features extracted across subjects were also examined and can be seen in Figure 3.7. The distributions of these features are comparable with the subject-specific features, suggesting that the ECAP detectors may be generalizable across subjects.

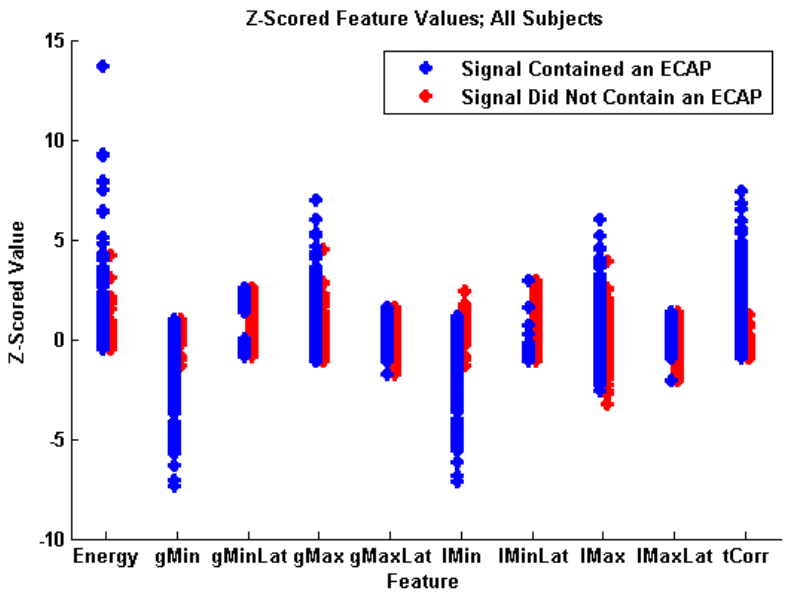

FIGURE 3.7: Z-scored ECAP features combined across subjects. The blue symbols represent features extracted from data that contained an ECAP, and the red symbols represent features gathered from data that did not contain an ECAP (as determined by visual inspection).

3.3.3 Cross-Validation

The ECAP detectors considered in this study were trained and tested using three-fold cross-validation. This technique divides the available data into three groups, or folds. During each iteration, two folds are used to train the classifiers, and the remaining fold is used for testing purposes. This process is completed three times, with each fold acting as the "testing" fold for one iteration. The results from each iteration were combined when evaluating the classifiers. Cross-validation allows the classifiers to have access to more training and testing data than they would otherwise be exposed to (e.g. [45]).

3.3.4 K-Nearest Neighbors

A K nearest neighbor (KNN) classification algorithm was implemented, using cross-validation and the features discussed above, to detect the presence of an ECAP

in a signal. For this method, the 'K' neighbors that are closest to the data point in question are considered. K was assigned a value of six, as this was empirically determined to provide the best results for the ECAP data. The threshold used to generate the receiver operating characteristic (ROC) curve (see Figures 3.9 and 3.10) is the number of neighbors, out of the six being considered, that must contain an ECAP in order for the data under consideration to be labeled as containing an ECAP (e.g. [45]).

3.3.5 Fisher's Linear Discriminant

A Fisher's linear discriminant was also used to detect ECAPs in this study. This classification method projects the feature space onto a single dimension with weight vector $\mathbf{w} : y = \mathbf{w}^T\mathbf{x}$. $\mathbf{X}$ represents the input vector (of dimension D), and $\mathbf{w}$ is selected such that the separation between the classes is maximized. That is, Fisher's linear discriminant aims to maximize the distance between the classes' means, while maintaining a small variance within each class. If $\mathbf{S}_W$ is defined as the complete within-class covariance matrix and $\mathbf{m}_1$ and $\mathbf{m}_2$ are the mean vectors of classes C_1 and C_2, respectively,

$$\mathbf{S}_W = \sum_{n \in C_1} (\mathbf{x}_n - \mathbf{m}_1)(\mathbf{x}_n - \mathbf{m}_1)^T + \sum_{n \in C_2} (\mathbf{x}_n - \mathbf{m}_2)(\mathbf{x}_n - \mathbf{m}_2)^T. \tag{3.1}$$

$\mathbf{W}$ is calculated such that $\mathbf{w} \propto \mathbf{S}_W^{-1}(\mathbf{m}_2 - \mathbf{m}_1)$. Since Fisher's Linear Discriminant aims to project the data onto a single dimension, the direction rather than the magnitude of $\mathbf{w}$ is important. The result, $y = \mathbf{w}^T\mathbf{x}$, can then be subjected to a threshold in order to classify the data (e.g. [45]).

3.3.6 Generalized Likelihood Ratio Test

Another classifier considered for the ECAP detection problem was the generalized likelihood ratio test (GLRT). This detector, as implemented in this study, assumes

that the features are normally distributed. Figure 3.8 displays the kernel density estimates of the features for each subject and for data gathered across subjects. Although not ideal for all features, the Gaussian distribution captures the general trends of the different densities.

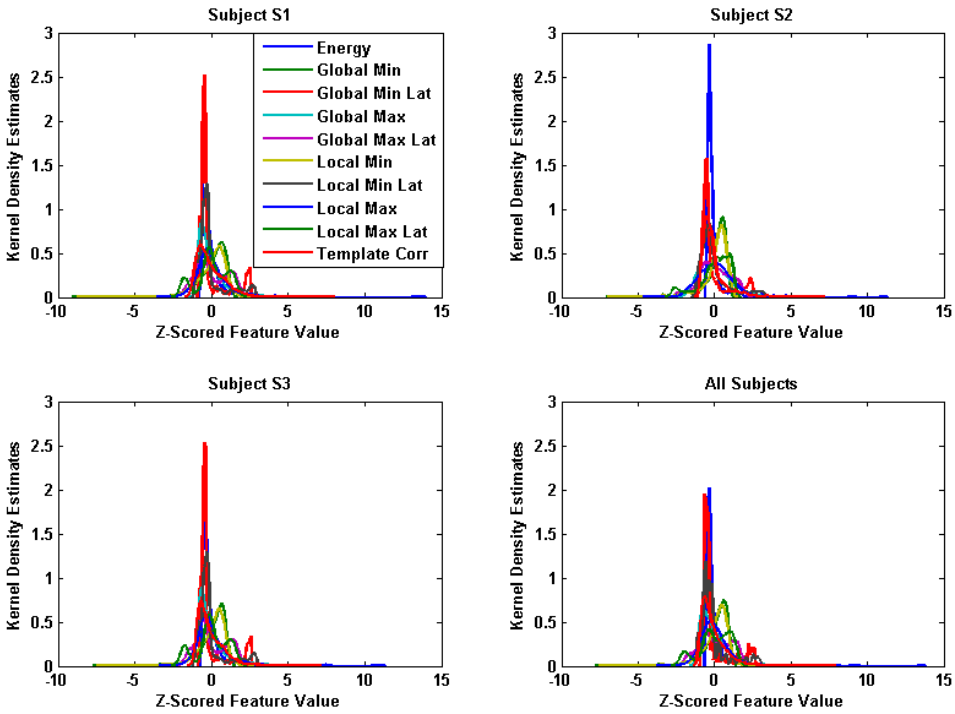

FIGURE 3.8: Kernel density estimates of the features used for ECAP detection. The GLRT assumes that the features are distributed normally. Although this is not an ideal fit for each density displayed here, the Gaussian distribution appears to be acceptable for the group of features.

For this detector, the likelihood ratio

$$\lambda(\mathbf{x}) = \frac{p(\mathbf{x}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)}{p(\mathbf{x}|\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)} \tag{3.2}$$

is computed, where

$$p(\mathbf{x}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) = \frac{\exp(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_1)^T\boldsymbol{\Sigma}_1^{-1}(\mathbf{x}-\boldsymbol{\mu}_1))}{(2\pi)^{\frac{D}{2}}|\boldsymbol{\Sigma}_1|^{\frac{1}{2}}} \tag{3.3}$$

and

$$p(\mathbf{x}|\boldsymbol{\mu_0}, \boldsymbol{\Sigma_0}) = \frac{\exp(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_0)^T\boldsymbol{\Sigma}_0^{-1}(\mathbf{x}-\boldsymbol{\mu}_0))}{(2\pi)^{\frac{D}{2}}|\boldsymbol{\Sigma}_0|^{\frac{1}{2}}} \tag{3.4}$$

$\mathbf{X}$ is the D-component input vector, $\boldsymbol{\mu}$ is the D-component mean vector of the features, and $\boldsymbol{\Sigma}$ is the D-by-D covariance matrix of the features. The subscripts 1 and 0 refer to the cases in which an ECAP is present and in which there is no ECAP present, respectively. To classify the data, $\lambda(\mathbf{x})$ is compared to a threshold.

3.3.7 Energy Detector

A simple energy detector was also considered as a baseline classifier for the ECAP detection segment of this study. This detector requires only that the signal energy be calculated and compared to a threshold for classification. If a signal contains energy above this threshold, it is classified as containing an ECAP.

3.3.8 Template Correlation

Finally, as another baseline ECAP detector, template correlation was used to classify the signals. A template was created by averaging ten signals known to contain strong ECAPs, and decision statistics were obtained by correlating an input signal with this template. The decision statistic for each signal was the maximum correlation that resulted when shifting the template to all possible locations relative to the input signal.

3.4 ECAP Detection Results

The ROCs for subject-specific ECAP detectors are shown in Figure 3.9. The plot on the top left demonstrates results obtained using Subject S1's data, the plot on the top right contains results found using Subject S2's data, and the bottom plot was created using Subject S3's data. These algorithms outperform the baseline detectors (the energy detector and template correlation), with the KNN and GLRT methods outperforming the others overall. However, further investigation is required to obtain a higher probability of detection, while maintaining a low probability of false alarm.

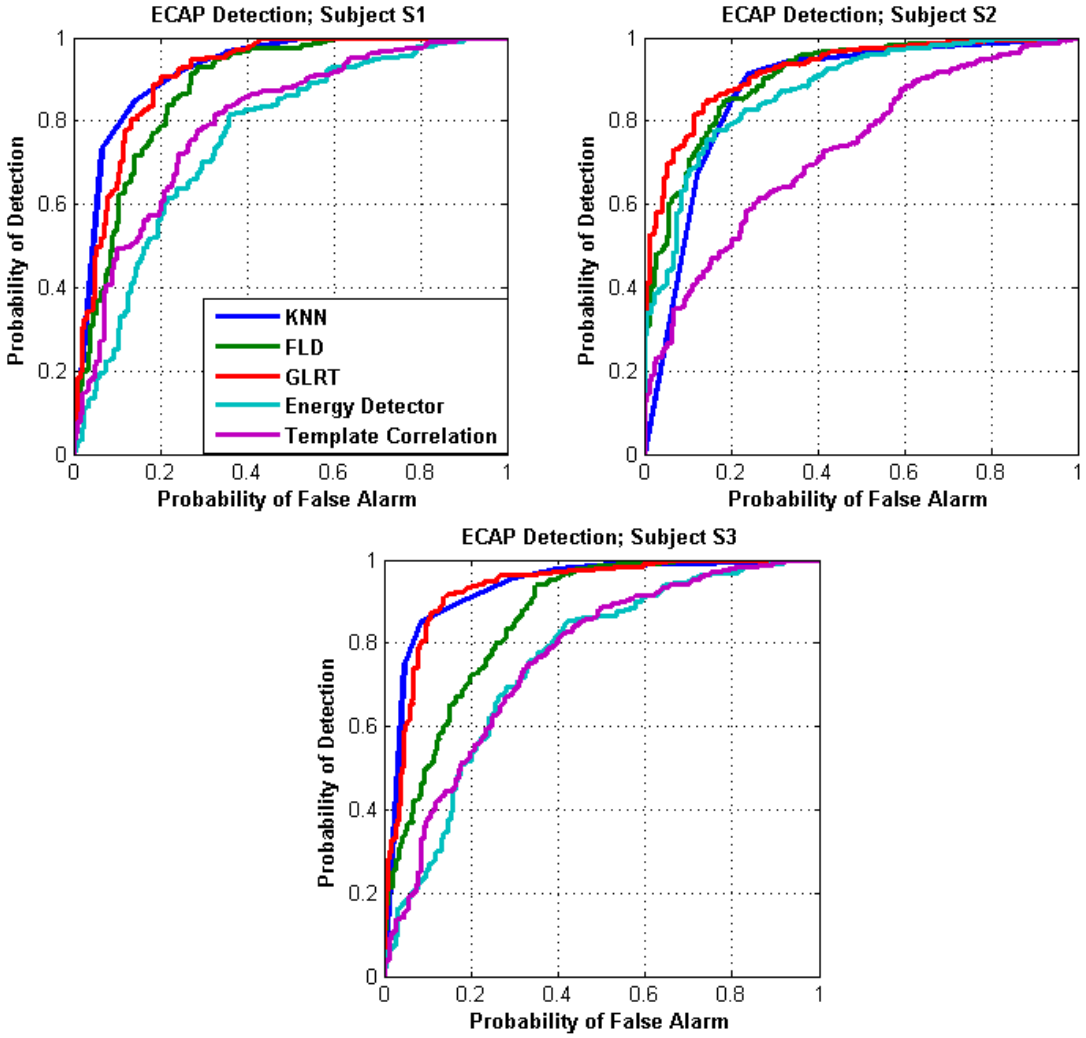

FIGURE 3.9: Subject-specific ROC curves for the ECAP detectors. The plot on the top left demonstrates the results obtained using Subject S1's data, the results in the top right plot were found using Subject S2's data, and the bottom plot was created using Subject S3's data.

The classifier that will be used in a future study will be trained on data collected from various subjects, in order to maintain generality. Figure 3.10 demonstrates the performance results of the ECAP detectors using data collected from all three subjects. These results suggest that the ECAP detectors are generalizable across subjects but, similarly to the subject-specific detectors, they will require future improvements. These improvements could be in the form of selecting more relevant

features or implementing more appropriate detection algorithms. Similarly to the subject-specific results, the KNN and GLRT methods appear to be the strongest ECAP classifiers overall.

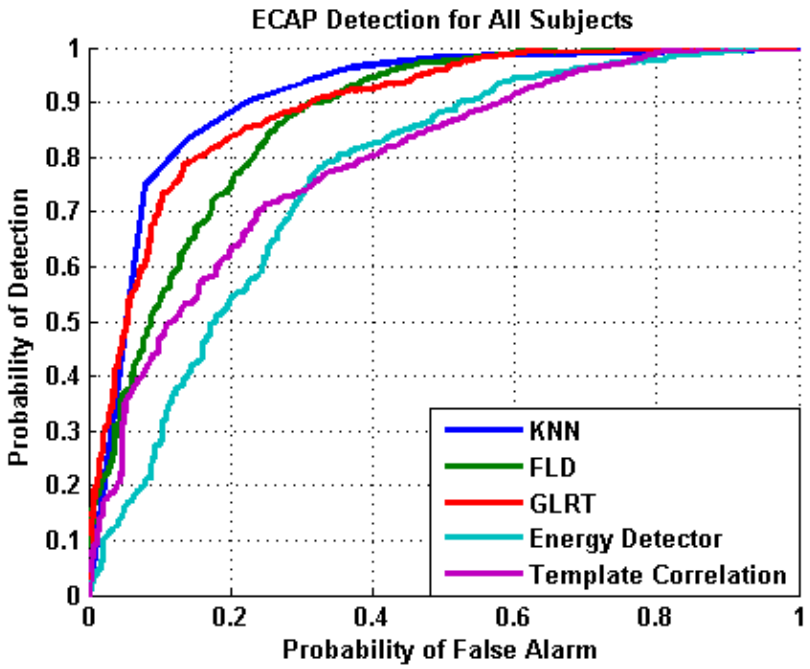

FIGURE 3.10: ROC curves for ECAP detectors trained and tested using data gathered from all subjects. The success of these detectors suggests that, although future work is required, the ECAP detectors are generalizable across subjects.

3.5 ECAP Detection Conclusions

Although the linear regression method is used clinically, this study determined that the skew of ECAP thresholds resulting from linear regression would introduce undesirable inaccuracies into this study. The other ECAP detectors that were considered in this study outperformed the two baseline detectors, energy detection and template correlation. Although these detectors could prove useful for future larger scale investigations, the ECAP thresholds for the remainder of this study were determined visually.

4

Experimental Methods

This study utilized forward masking measurements to estimate information in the stimulation patterns of speech that, since it is masked, may not be received by the user. These estimates could potentially guide new speech processing algorithms to mitigate lost information or substitute alternate information. Because forward masking patterns cannot be measured psychophysically in a clinically relevant time frame [2], this study obtained forward masking patterns with ECAP measurements. Several experiments were required to achieve the goals of approximating forward masking patterns with ECAP measurements. First, ECAP threshold patterns were measured for each electrode in the array so that each electrode could act as a masker with all other electrodes acting as probes. Data was then collected from a psychophysical forward masking task, and this data was used to scale the ECAP threshold patterns into forward masking patterns.

4.1 ECAP Data Collection

4.1.1 Experimental Plan

This study aimed to determine the shift in the ECAP threshold of a probe electrode that resulted from stimulation on a preceding masker electrode. (ECAP threshold is defined to be the threshold at which an ECAP can first be measured). Since a greater degree of forward masking results in a strong ECAP, threshold decreases as masking increases.

Each electrode was designated as the "masker" when measuring its ECAP threshold pattern. The masker electrode was held constant and was presented at 90% of the subject's dynamic range throughout the measurements. The amplitude of the probe pulses varied between 50-90% of the subject's dynamic range. The current level at which an ECAP was first visualized was recorded manually as the ECAP threshold for that masker-probe pair, and ECAP thresholds were determined for all masker-probe combinations. As an example, the ECAP threshold profile for Subject S3, masker electrode 8 is shown in Figure 4.1. Unlike forward masking threshold shifts, greater amounts of masking are reflected by smaller ECAP thresholds.

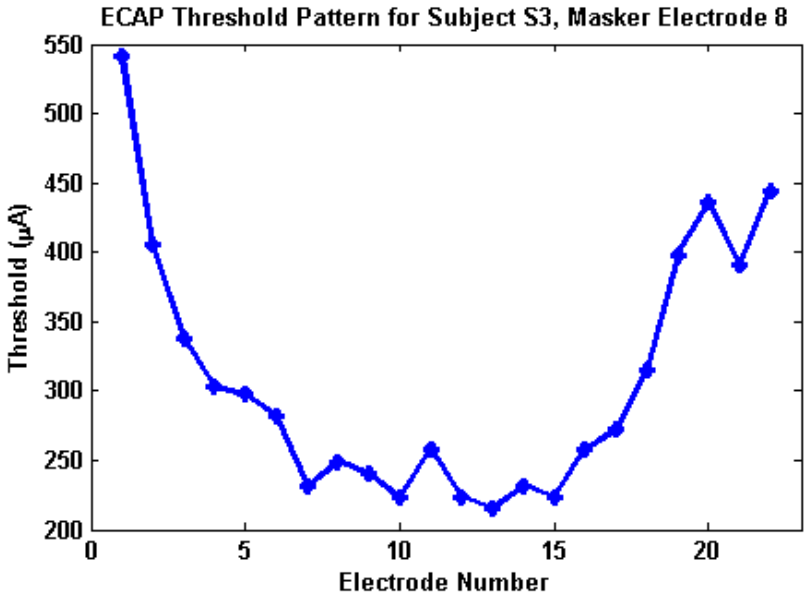

FIGURE 4.1: ECAP threshold pattern for Subject S3, masker electrode 8. The current level required to achieve an ECAP threshold is plotted for all probe electrodes. Note that smaller threshold values correspond to greater amounts of masking.

Once ECAP thresholds were determined for all masker-probe combinations, ECAP threshold patterns were to be scaled into forward masking patterns. A previous study suggests that psychophysical threshold shifts and ECAP threshold shifts are correlated [9]. Therefore, psychophysical forward masking measurements were recorded to scale the ECAP threshold patterns into forward masking threshold shift patterns. The specifics describing how the ECAP threshold patterns were scaled are outlined in a future section.

4.2 Psychophysical Forward Masking Data Collection

Threshold shifts measured using traditional forward masking approaches, in which a masker pulse train precedes a probe pulse train, were found to not correlate strongly with ECAP threshold shifts in all subjects [43]. However, Hughes and Stille (2009) used simultaneous stimulation when measuring ECAP threshold shifts and psychophysical masking patterns, and they found a correlation between the two mea-

surements. Their success was thought, in part, to be due to the similarity between the ECAP measurements and the psychophysical measurements [9].

The NRT system available in the Cochlear Custom Sound EP software does not allow simultaneous stimulation while measuring ECAPs. Rather, all ECAPs in this study resulted from measuring the response to a masker followed by a probe, and ECAP threshold shifts could not be measured directly. Alternatively, the ECAP threshold patterns that were measured were scaled into forward masking patterns using select psychophysical measurements to approximate ECAP "probe-alone" values. (The methods of approximation will be explained in a future section.) The psychophysical measurements were in the form of interleaved masking, which consists of a pulse train in which each masker pulse is followed by a probe pulse in an "interleaved" pattern. By ensuring that the interleaved masking parameters and the ECAP measurement parameters were as similar as possible, this study aimed to maintain the correlation between ECAP and psychophysical threshold shift measurements found by Hughes and Stille (2009).

Although some psychophysical measurements were required to approximate forward masking patterns with ECAP threshold patterns, this study measured a select few psychophysical threshold shifts. Therefore, the method outlined in this document is still more time efficient than traditional psychophysical methods.

4.2.1 Stimuli and Equipment

The psychophysical forward masking experiment consisted of 300 ms biphasic pulse trains, each pulse 25 μs wide. The interpulse gap for Subjects S1 and S3 was 7 μs, and the gap for Subject S2 was 8 μs due to hardware limitations associated with his implant. Stimuli were presented using the Nucleus Implant Communicator (NIC), which was driven by a PC.

Prior to testing, subjects' thresholds and maximum comfortable levels were mea-

sured using 80 pps pulse trains of length 300 ms. These measurements assisted in determining the range of stimulation levels for the following experiment. Parameters throughout this experiment were selected to match those used to measure ECAP threshold patterns.

Two different pulse trains were presented to the subjects (see Figure 4.2 for a zoomed-in illustration of these stimuli). The first pulse train contained pulses presented at a rate of 80 pps (a). The second pulse train contained the same 80 pps pulse train, which became the "masker" pulses, with the addition of "probe" pulses inserted 400 μs after each of these masker pulses (b). The masker pulses were presented at 90% of the subject's dynamic range (in both scenarios (a) and (b)), and the amplitude of the probe pulses varied depending on the accuracy of the subjects' responses during the task. This was completed for all electrodes, with both masker and probe presented to the same electrode. Similar forward masking experiments, referred to as "interleaved masking" have been implemented in [4], [46], [47], [48].

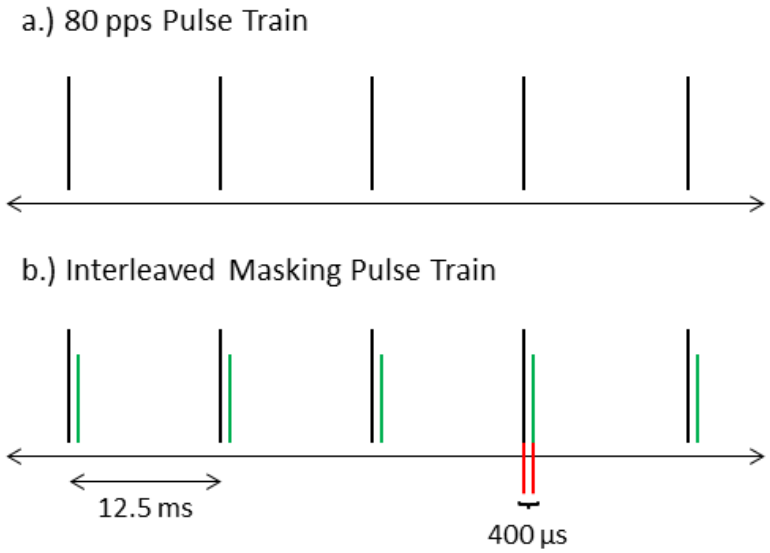

FIGURE 4.2: One section of the stimuli used in the psychophysical forward masking task. Both pulse trains were 300 ms in length (containing more pulses than shown here). Pulse train (a) was simply a pulse train presented at 80 pps. Pulse train (b) contained the same base 80 pps pulse train, with additional "probe" pulses inserted 400 μs after each "masker" pulse. Both masker and probe stimuli were presented to the same electrode.

4.2.2 Experimental Plan

Subjects were presented with three pulse trains, separated by 500 ms of silence. Pulse train (a) (see Figure 4.2) was presented in interval 1 as well as in interval 2 or 3 (determined randomly). Pulse train (b) was presented in the remaining interval.

A graphical user interface (GUI), shown in Figure 4.3, was presented to the subjects. After pressing the "Ready" button they were presented with the three stimuli, while corresponding buttons on the GUI were illuminated. Subjects were instructed to select the interval that differed from the others (the interval containing the interleaved masking stimulus). After a response was submitted, the button associated with the correct response was highlighted in green to provide feedback.

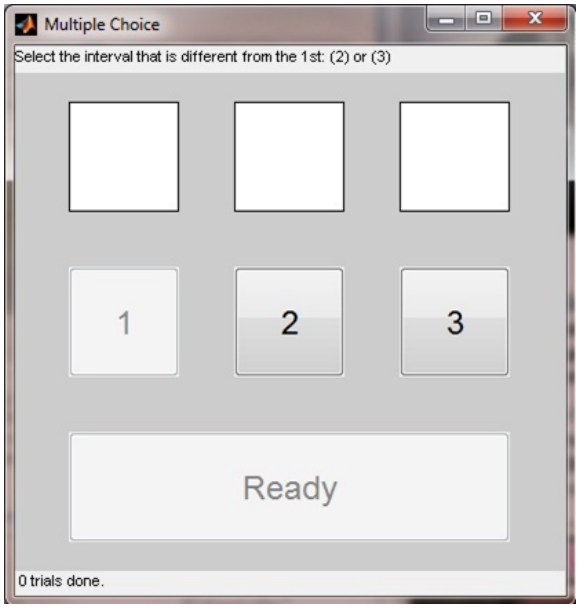

FIGURE 4.3: The graphical user interface (GUI) presented to subjects during the forward masking task. Subjects were instructed to press the "Ready" button, listen to the stimuli (while corresponding segments of the GUI were illuminated), and select which interval (2 or 3) contained the stimulus that differed from the others. After each response, the button associated with the correct response was highlighted in green to provide feedback.

The masker and the probe pulses were initially presented at 90% of the subject's dynamic range. This task was an adaptive "two-down, one-up" procedure (converging on the 70.7% location on the psychometric function [49]), which required two correct responses for the probe current level to be decreased, but only one incorrect response for the probe current level to be increased. The step size was two current steps (on Cochlear's 1-255 scale) for the first four reversals. A "reversal" occurred when an incorrect response was made after two correct responses, or when two correct responses occurred after an incorrect response. The step size was then decreased to one current step for the final four reversals. An average of the probe current levels presented during the final four reversals was recorded for each trial, and five trials were averaged for each electrode.

4.3 Forward Masking Patterns

As discussed in Section 4.2, psychophysical measurements were necessary to determine the scaling factors that converted ECAP threshold curves to forward masking curves. The psychophysical shifts were calculated for each electrode when the masker and the probe shared a common location. This was completed by subtracting each electrode's psychophysical threshold measurement from the corresponding psychophysical forward masking threshold measurement. Since this psychophysical threshold shift corresponds to the ECAP threshold shift, subtracting the psychophysical threshold shifts from the ECAP thresholds measured when the masker and the probe were on the same electrode, as seen in Figure 4.4, results in the ECAP "probe-alone" threshold at that electrode location. The collection of probe-alone values are the scale factors needed to convert the ECAP threshold curves to psychophysical threshold shift curves.

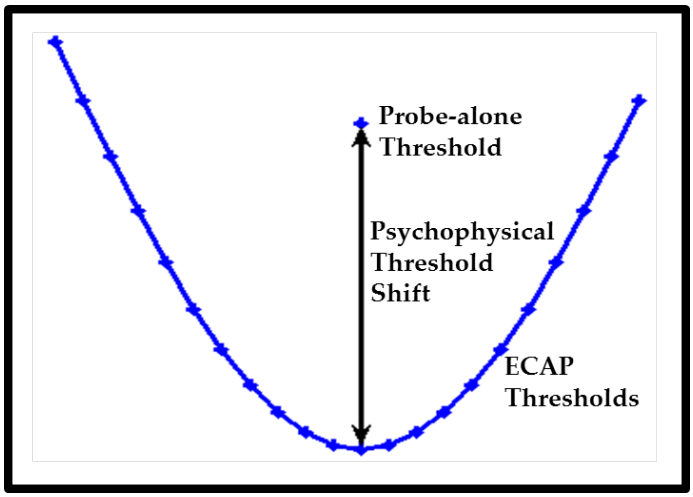

FIGURE 4.4: Model of calculating ECAP probe-alone threshold. The ECAP probe-alone threshold was calculated by subtracting the psychophysical threshold shift (the difference between the psychophysically-measured masked threshold and the psychophysical probe-alone threshold) from the ECAP threshold measured when the masker and the probe were located on the same electrode in the array.

Once the ECAP probe-alone threshold was calculated for every electrode in the array, ECAP thresholds could be converted into forward masking threshold shifts. This was completed by subtracting the probe-alone values from the ECAP thresholds that were associated with those probe locations [9]. The threshold shift pattern that resulted for Subject S3, masker electrode 8 is shown in Figure 4.5. Nonmonotonicities similar to those observed in this figure have also been reported in the literature [1], [2], [3], [5].

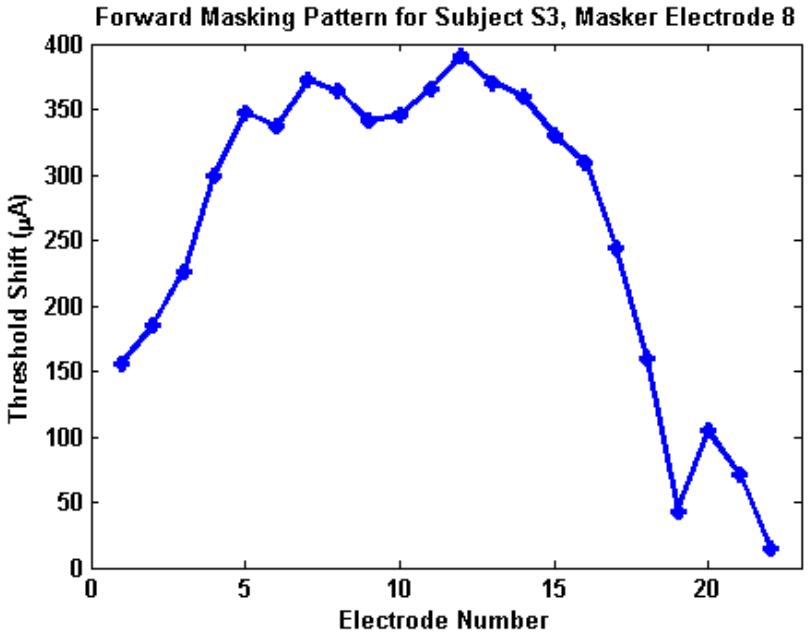

FIGURE 4.5: Threshold shift pattern for Subject S3, masker electrode 8. Note that larger threshold shift values now correspond to greater amounts of masking.

Threshold shift patterns are extremely variable in width and shape, and ECAPs are not always measurable for every masker-probe combination at current levels comfortable for the subject. Oftentimes, as in Figure 4.6, ECAPs are only measurable when the masker and probe are separated by a minimal distance on the electrode array.

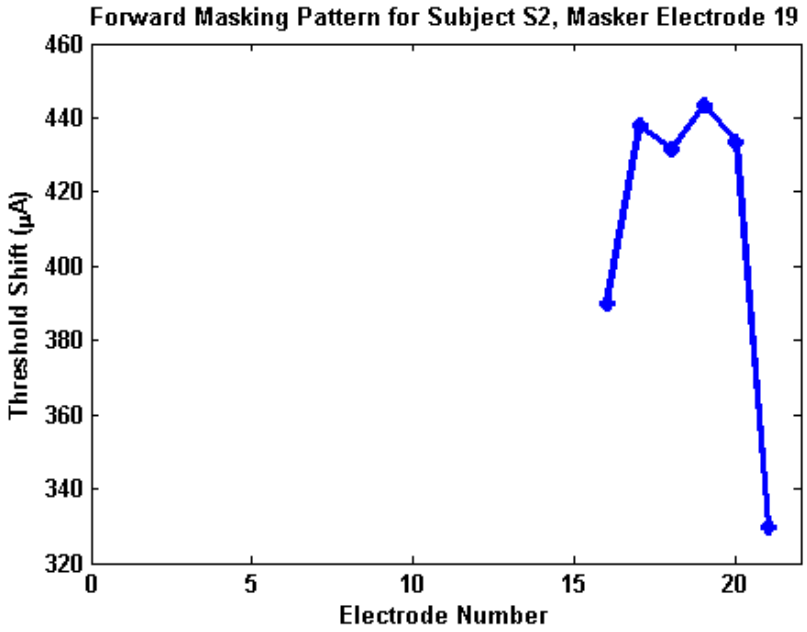

FIGURE 4.6: Threshold shift pattern for Subject S2, masker electrode 19. This threshold shift pattern is notably narrower than that shown in Figure 4.5. Such a narrow pattern is not uncommon, as ECAPs are not measurable for all masker-probe combinations at current levels comfortable for the subject.

Additionally, the peaks of forward masking patterns do not always occur when the masker and the probe are on the same electrode, as shown in Figure 4.7. This figure displays the forward masking pattern that resulted when electrode 3 was designated as the masker electrode, but the peak of the masking pattern occurred when the probe was presented to electrode 5. Such results are not uncommon and have been reported in the literature (e.g. [2]).

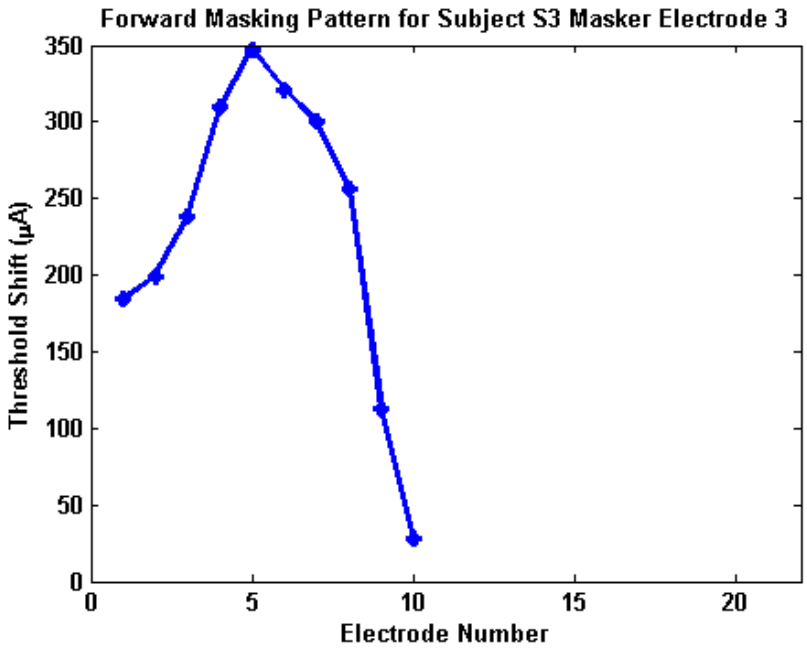

FIGURE 4.7: Threshold shift pattern for Subject S3, masker electrode 3. Although the masker was located at electrode 3, the peak of this forward masking pattern occurs at electrode 5. Peaks that occur when the masker and the probe are located at different electrodes have been reported in the literature (e.g. [2]).

Many forward masking pattern approximations resulted in deviations from the "ideal" case in which a maximum occurs when the masker and the probe are presented to the same electrode and the curve monotonically decreases as the distance between the masker and probe is increased. However, the discrepancies seen in this study are not uncommon and have been reported in the literature.

5

Analysis of Forward Masking Patterns

Compared to normal hearing listeners, cochlear implant users are presented with limited frequency and temporal information. Some of this limited information is potentially imperceptible due to channel interactions, and forward masking has been used to identify some of these interactions (e.g. [1], [2], [3], [4], [5], [6], [7]).

This study approximated forward masking patterns using ECAP measurements. In order to determine whether the approximations were accurate and whether forward masking can predict imperceptible information, this study applied the forward masking patterns to speech stimuli.

5.1 Forward Masking Patterns Applied to Speech Stimuli

5.1.1 Stimulation Patterns

To validate the accuracy of the forward masking patterns, and to investigate whether they can be used to predict imperceptible information, the patterns were used to remove "masked" pulses from speech stimulation patterns. Figure 5.1 shows the stimulation pattern for the original speech token "asa." Although this is similar to the electrodogram displayed in Section 2.1, Figure 5.1 displays the stimulation pat-

tern generated using Subject S3's clinical parameters. Recall, these electrodograms consist of pulses that occur on a given electrode (frequency band) at a given time. The amplitude of these pulses corresponds to the current level that is delivered.

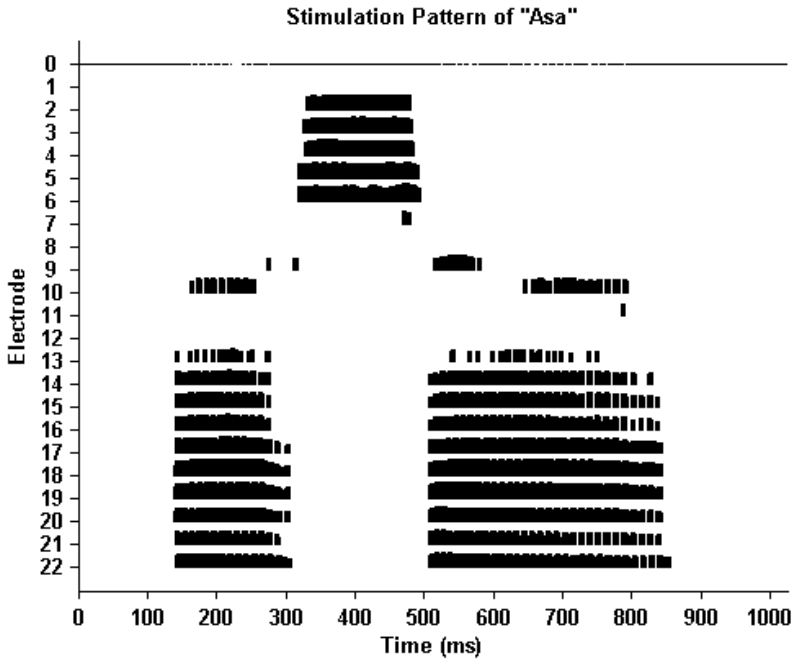

FIGURE 5.1: Stimulation pattern of the speech token "asa" generated using Subject S3's clinical parameters. This stimulation pattern, referred to as an electrodogram, demonstrates the frequency and temporal information that is presented to Subject S3 during the speech token "asa." Time is plotted on the x axis, and the y axis designates the electrode. If a given electrode is to be stimulated at a given time, a "tick" mark, with amplitude corresponding to the stimulation current level, will appear at the corresponding location

5.1.2 Applying Forward Masking Patterns to Stimulation Patterns

When deciding which pulses in a given stimulus were likely masked, each pulse was considered to be a "masker" pulse that could mask future "probe" pulses. The previously measured forward masking patterns were linearly scaled according to the amplitude of each masker pulse [23], and the amplitudes of the patterns were decayed according to the delays between the maskers and the probes. Probe pulses with

amplitudes less than the sum of their original thresholds and the threshold shift imposed by the masker pulse were removed from the stimulation pattern before being presented to the subject. Figure 5.2 demonstrates the stimulation pattern for the speech token "asa" according to Subject S3's clinical parameters after "masked" information has been removed from the stimulus. In this example, 35.10% of pulses were removed from the original stimulation pattern. When applying the subject-specific forward masking patterns to speech stimuli for Subjects S1, S2, and S3, an average 37.41%, 59.05%, and 28.22% of pulses, respectively, were assumed to be masked and were therefore removed from the stimulation patterns.

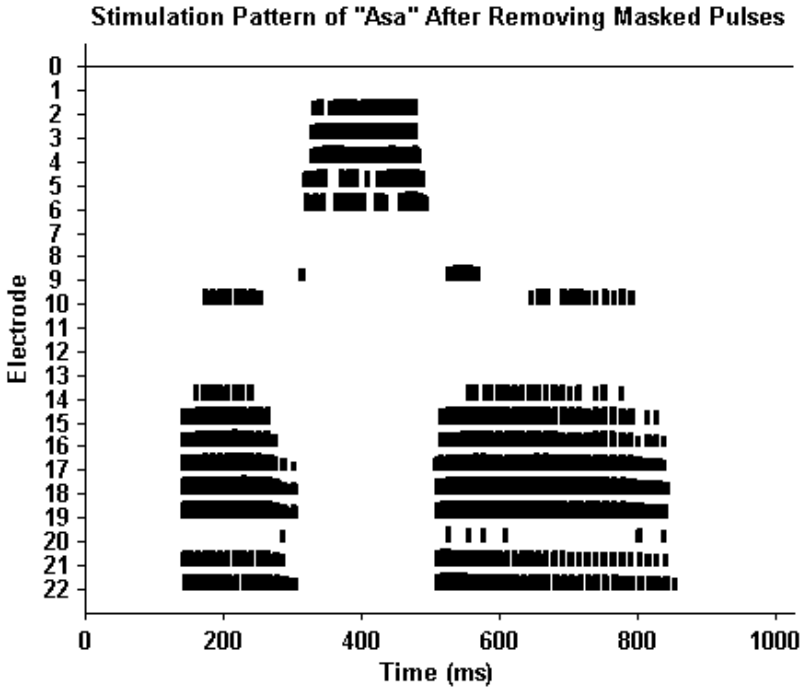

FIGURE 5.2: Stimulation pattern of the speech token "asa" after removing masked pulses. This stimulation pattern demonstrates the unmasked frequency and temporal information, as estimated using Subject S3's forward masking patterns. 35.10% of pulses were removed from the original stimulation pattern shown in Figure 5.1.

5.1.3 Forward Masking Parameters

Forward masking patterns must be scaled based on the masker amplitude and the time between the masker and the probe. As previously mentioned, forward masking patterns in this study were linearly scaled according to the amplitude of each masking pulse [23]. Additionally, it has been found that forward masking patterns decay exponentially with time (e.g. [18], [23], [32], [50]). A previous study investigated the temporal decay resulting from single-pulse forward masking, finding that the rapid recovery process completes by about 10 ms for most subjects [23].

Although the study outlined in this document approximated subject-specific forward masking patterns, the amplitude and temporal scaling parameters associated with the patterns reflected averages of those reported in the literature. As a result this study anticipated that when testing speech recognition, parameters would most likely require alterations to fit each subject personally.

5.2 Speech Reception Threshold

For this portion of the study, a speech reception threshold (SRT) task was completed. Subjects were presented with sentence stimuli and were asked to reproduce what they heard. A correct response would cause the signal to noise ratio (SNR) of the subsequent sentence to decrease, while an incorrect response would result in an increased SNR. This one-down, one-up procedure converges to the 50% location on the psychometric function [49], and therefore determines the noise level at which subjects can correctly identify sentences 50% of the time.

The SRT task was selected for this experiment because it focuses around the noise level at which subjects struggle to understand speech. A more traditional task in which stimuli are presented at a fixed SNR is potentially less robust and may allow subjects to simply compensate for missing information in the stimuli from which

"masked" pulses were removed.

5.2.1 Stimuli and Equipment

This portion of the study presented stimuli utilizing the parameters (pulse rate, pulse width, interphase gap, etc.) that the subjects used most frequently in every-day situations. Thresholds and maximum comfortable levels, however, were measured to account for differences in hardware.

5.2.2 Experimental Plan

The SRT task utilized sentences from the Hearing in Noise Test (HINT) [51]. The sentences are divided into 25 "lists," each containing ten sentences. The first sentence of each list was presented to the subject at an SNR of 16 dB. If the subject was able to correctly reproduce the sentence (with the exception of plurals and articles), the sentence was considered to be correct, the SNR was decreased by 2 dB, and the next sentence was presented. If the subject could not correctly reproduce the sentence, the SNR was increased by 2 dB before the subsequent sentence was presented. This one-down, one-up procedure continued until all ten sentences in the given list had been presented. The final SRT value was determined to be the average of the SNRs of sentences 5-10.

To test if removing masked pulses affected speech recognition, the SRTs were measured for two conditions in all subjects. The first condition consisted of speech stimuli presented in their original form, with all pulses retained. In the second condition, "masked" pulses were removed from the original stimuli according to the subject-specific forward masking patterns. Five repeats were conducted for each condition, and the distributions resulting from these repeats were compared to determine the effects of removing masked pulses. If this study successfully predicted imperceptible information, the distributions resulting from the two conditions should

be comparable.

5.2.3 Results

For Subjects S1 and S2, the initial results reflected differences in the SRT distributions when all pulses were retained and when masked pulses were removed from speech stimuli. These differences suggested that the two subjects required forward masking pattern scaling parameters that deviated from averages reported in the literature. In order to decrease the amount of masked pulses in the stimulation patterns of these two subjects, the decay scaling parameters were altered. Specifically, this study decreased the recovery process completion time, or the time after which the exponentially decaying forward masking pattern would have approximately no effect on future probe stimuli.

Figure 5.3 displays the speech reception thresholds that were found for the three subjects and two conditions. No statistically significant differences were found between the distributions of the two conditions for all three subjects ($p = 0.54$, $p = 0.48$, and $p = 0.32$, respectively). These distributions suggest that forward masking can potentially predict imperceptible information that can be dropped from a stimulation pattern without negative effects on speech recognition.

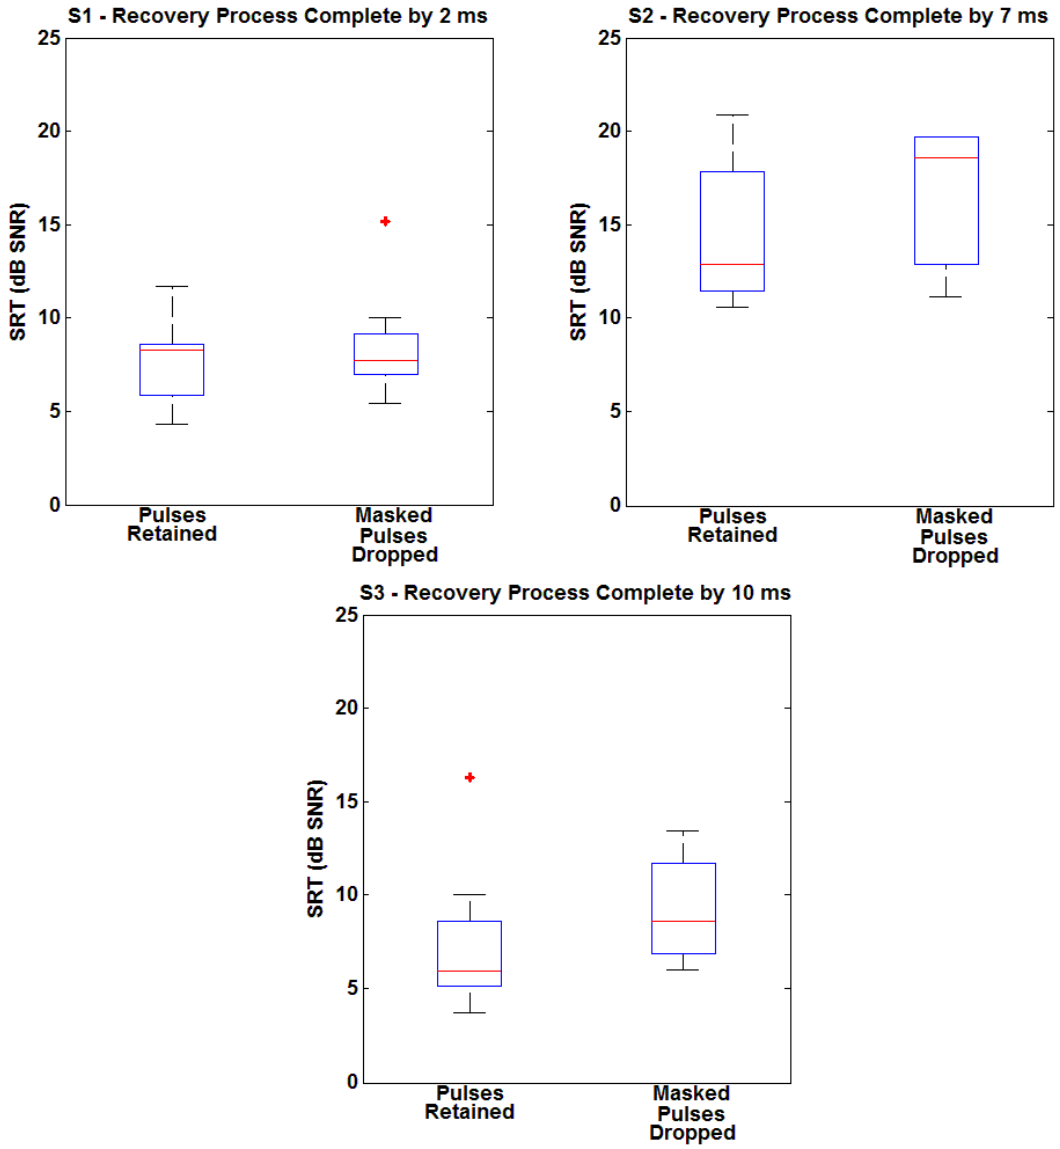

FIGURE 5.3: Speech reception thresholds measured for three subjects and two conditions: original sentence stimuli, and sentence stimuli from which masked pulses were removed. For all three subjects, the differences between the "pulses retained" and the "masked pulses dropped" distributions were not found to be statistically significant.

As noted in the title of the top left plot of Figure 5.3, a recovery process completion time of 2 ms was used for Subject S1. Because this parameter is unrealistic according to the literature (e.g. [23]), a post hoc analysis, comparing alterations to

the amplitude and temporal scaling factors, was completed.

5.2.4 Post Hoc Comparison of Parameter Effects

Although a future study will investigate the relationship between altering the scaling parameter (by decreasing the amplitudes of the masking patterns) and altering the temporal decay (by shortening the recovery process completion time) an initial comparison suggests that the results of the two alterations are similar. Figure 5.4 displays the estimated average percentages of pulses masked from sentence stimuli as a function of electrode number for the three subjects. Results for the two modifications, altering the temporal decay parameter (blue) and altering the scaling parameter (red), were calculated by applying the subject-specific forward masking patterns to sentence stimuli in a MATLAB simulation. This simulation recorded the percentages of pulses masked per channel for each subject, averaging the results over 100 sentences. The two alterations resulted in similar masking percentages per channel when the overall percentages of masked pulses were equal. For Subject S3 the results are identical for both cases because, for this subject, the recovery process completion time was set to 10 ms. Therefore, the amount of masking did not need to be reduced by either shortening the recovery process completion time or by decreasing the masking pattern amplitude.

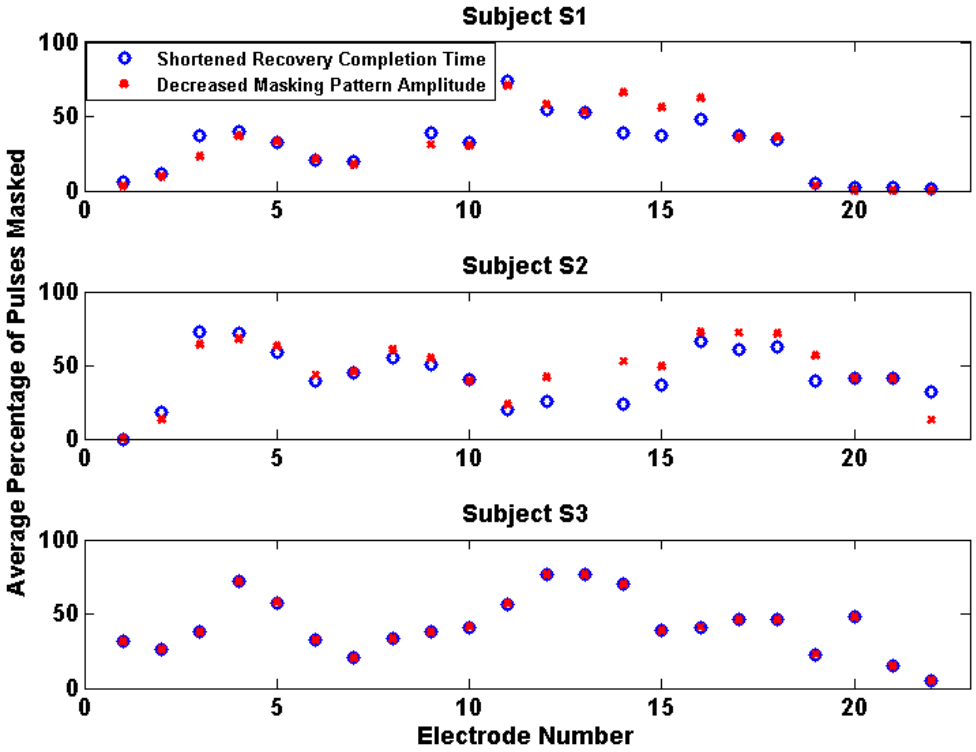

FIGURE 5.4: The percentages of pulses masked per channel for three subjects, when altering the recovery process completion time versus altering the masking pattern amplitude. This figure suggests that altering either parameter results in similar percentages of masked pulses per channel. Note: the recovery process completion time for Subject S3 was set equal to 10 ms. Therefore, the amount of masking did not need to be reduced by either shortening the recovery process completion time or by decreasing the masking pattern amplitude. As a result, the plots are identical for the two methods.

5.3 Analysis of Masked Speech

This study also investigated the amount of masking that occurs in different parts of speech. Understanding which phonemes are most vulnerable to masking could aid in understanding the perceptual effects of the phenomenon. Figure 5.5 displays the number of pulses masked from different parts of speech as a function of the total number of pulses originally present. The slopes for each condition, displayed in the

upper left corner of each plot, suggest subtle differences in the amount of masking occurring per phoneme. These differences, if statistically significantly different, could aid in a better understanding of forward masking's effect on different parts of speech.

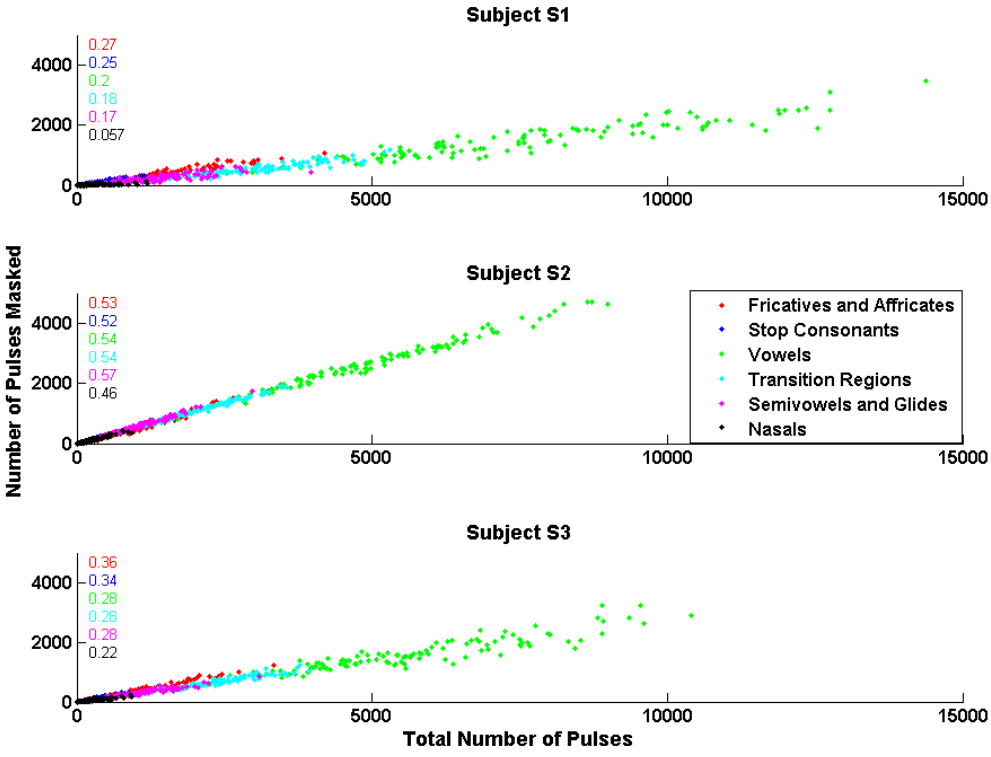

FIGURE 5.5: The number of pulses masked from different parts of speech as a function of the total number of pulses originally present. Data for Subjects S1, S2, and S3 are displayed in the top, middle, and bottom plots, respectively. The corresponding slopes are displayed in the upper left corner of each plot. This figure implies subtle differences in the amount of masking occurring per part of speech.

Figure 5.6 displays the average amount of masking that occurs per phoneme across subjects. Using this data and comparing all phonemes to vowels, only nasals were found to be statistically different (student t-test, $p < 0.05$). However, the subjects who participated in this experiment had low levels of masking on the low frequency channels associated with nasals, suggesting that this finding may be specific to those

who participated in this experiment.

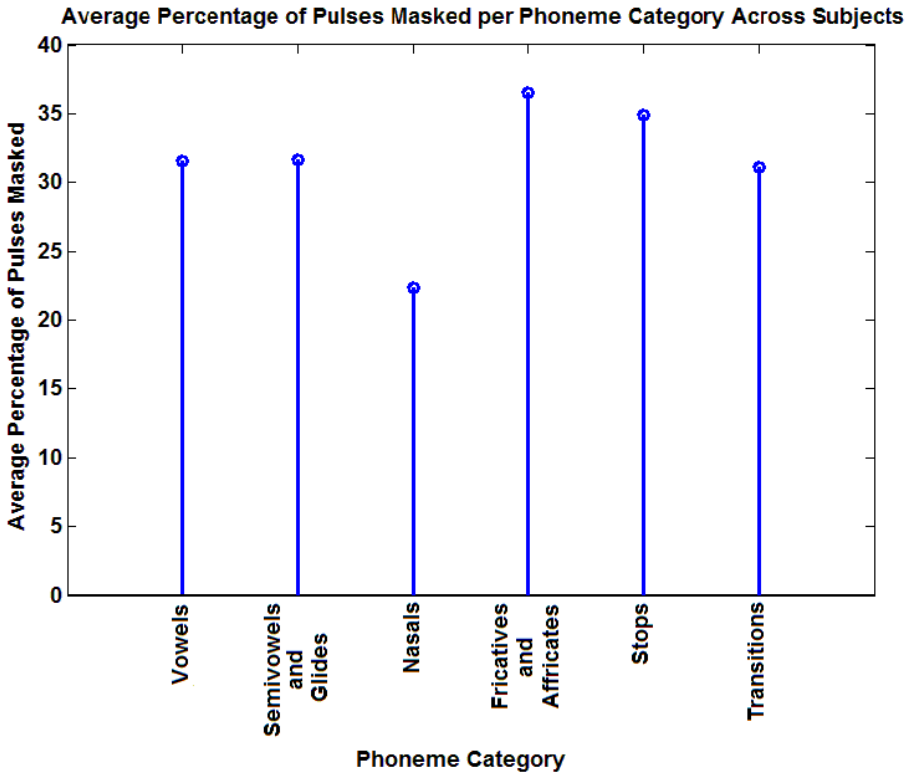

FIGURE 5.6: Stem plot of the average percentage of pulses masked per phoneme across subjects. Although nasals were statistically different than vowels (student t-test, $p < 0.05$), this may be due to subject-specific low levels of masking on low frequency channels.

Figure 5.7 demonstrates the low levels of masking on the low frequency channels (high electrode numbers) associated with nasals for Subject S1. This figure, which is a plot of the forward masking patterns for electrodes 19-22, displays low levels of masking for electrodes 19-22 when the probe is on or near the masker location. As these should be the instances when masking is greatest, the levels of masking seen here are lower than expected.

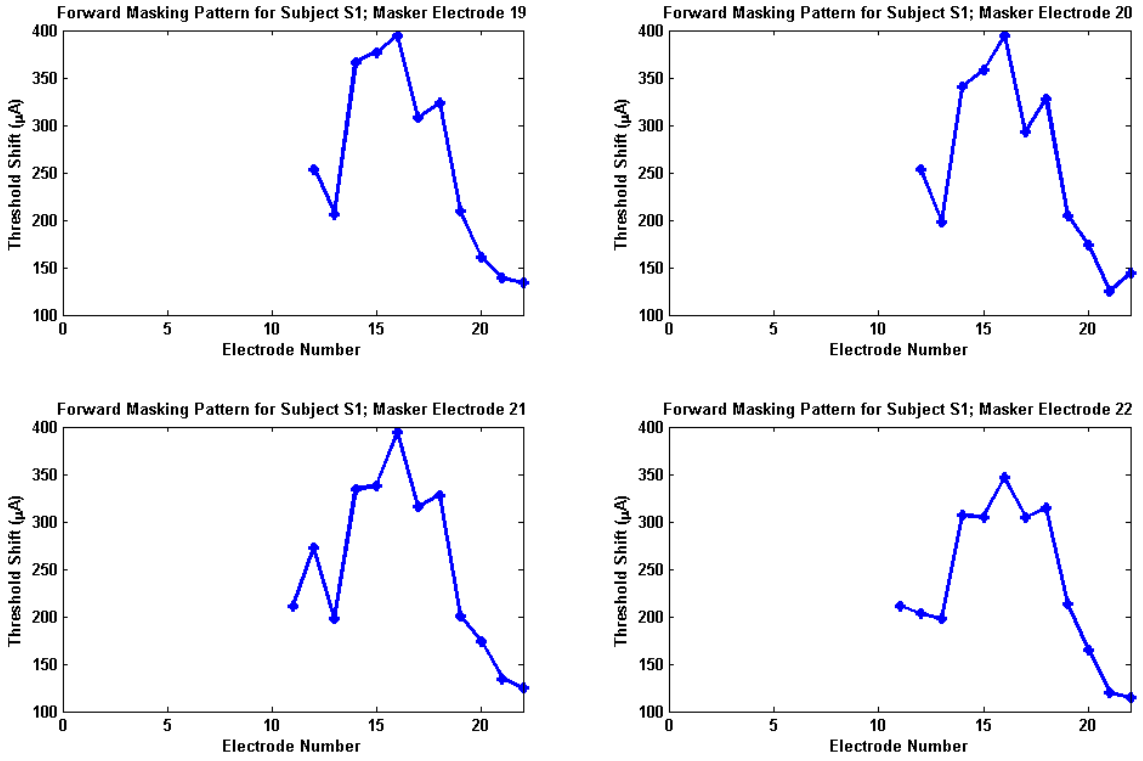

FIGURE 5.7: Forward masking patterns for Subject S1, masker electrodes 19-22. The greatest amount of masking is expected when the masker and probe are on or near the same electrode, and these plots demonstrate low levels of masking on low frequency channels (high electrode numbers) in these scenarios.

5.4 Discussion

This study successfully applied forward masking pattern approximations to speech stimuli, removing masked pulses from sentence stimuli. The SRT task evaluated each subject's speech recognition abilities for two conditions: sentence stimuli from which all pulses were retained and sentence stimuli from which masked pulses were removed. No statistically significant differences were found between the two distributions of SRTs for all subjects, suggesting that forward masking, with some level of adjustment, can predict imperceptible pulses that can be dropped from stimuli without negatively effecting speech recognition. The adjustment that this study made was in

the form of altering the temporal decay parameter. A post hoc analysis suggested that altering this parameter has similar effects to altering the scaling parameter of the forward masking patterns' amplitudes.

Finally, an analysis of the amount of masking occurring per speech phoneme resulted in inconclusive evidence for the wider cochlear implant population. Rather, it is hypothesized that the statistically significant difference found between the amount of information lost from nasals and vowels was a subject-specific result tied to characteristics of the forward masking patterns approximated in this study.

6

Conclusions

The results of this study suggest that forward masking patterns approximated using ECAP threshold measurements could predict information that is imperceptible for cochlear implant users. This study utilized subject- and electrode-specific forward masking patterns to estimate threshold shifts that resulted from each stimulating pulse. These threshold shifts were used to determine stimuli that may be "masked" as a result of a prior stimulus. To assess the hypothesis that these pulses were in fact imperceptible, an SRT task was completed for sentence stimuli from which "masked" pulses were retained and those from which "masked" pulses were removed. For all three subjects, no statistically significant differences were found between the distributions of the two SRTs. These results suggest that masked pulses may be removed from sentence stimuli without altering a cochlear implant listener's speech perception abilities, which could guide speech processing algorithms that either mitigate lost information or substitute alternate information.

The SRT task was used in an attempt to discourage subjects from simply compensating for missing information when masked pulses were removed from sentence stimuli. This task focuses around the noise levels at which subjects struggle to un-

derstand speech, potentially resulting in a more robust task than one in which speech is presented at a fixed SNR. However, to further argue that subjects are not simply compensating for missing information when masked pulses are removed, a future study will compare the SRTs measured in this document with those measured when pulses are removed at random. If the "masked" pulses are in fact not being perceived by the subjects, the scenario in which random pulses are removed should result in higher (worse) SRTs than the case in which masked pulses are removed.

Although the initial results are promising, further investigation is required to better fit the masking patterns to each subject. In this study, the amount of masking occurring in the sentence stimuli was decreased by altering the recovery process of the threshold shift patterns. Although adjusting the recovery process and decreasing the amplitude of the threshold shift patterns resulted in similar per channel masking, future work will investigate the accuracy of both approaches.

While approximating the forward masking patterns, this study highlighted the need for ECAP detection algorithms. Although these detectors performed relatively well, future work will require developing a method that will increase the probability of detection while maintaining a low probability of false alarm. Once improvements have been made, the results of this study suggest that the ECAP detectors will be generalizable across subjects.

This study also investigated the amount of masking that occurs in different speech phonemes. Although the amount of masking that occurs during the presentation of nasals was statistically significantly different than that occurring for vowels, the subjects who participated in this study experienced low levels of masking on the low frequency channels associated with nasals. Therefore this finding may be specific to the participants of this study.

In order to obtain more robust results, additional subjects will also be recruited. Besides providing support to the claim that removing masked pulses does not effect

speech recognition, more participants could help this study further investigate the effects of masking on different phonemes. This study observed low levels of masking on the low frequency channels of the current participants. Additional subjects could result in varied forward masking patterns, and thus different amounts of masking per phoneme.

Bibliography

[1] M. Chatterjee and R. Shannon, "Forward masked excitation patterns in multielectrode electrical stimulation," *J. Acoust. Soc. Am.*, vol. 103, pp. 2565–2572, 1998.

[2] C. Throckmorton and L. Collins, "Investigation of the effects of temporal and spatial interactions on speech-recognition skills in cochlear implant subjects," *J. Acoust. Soc. Am.*, vol. 105, pp. 861–873, 1999.

[3] R. Shannon, "Multichannel electrical stimulation of the auditory nerve in man. ii. channel interaction," *Hear. Res.*, vol. 12, pp. 1–16, 1983.

[4] Y. Tong and G. Clark, "Loudness summation, masking, and temopral interaction for sensations produced by electric stimulation of two sites in the human cochlea," *J. Acoust. Soc. Am.*, vol. 79(6), pp. 1958–1966, 1986.

[5] H. Lim, Y. Tong, and G. Clark, "Forward masking patterns produced by interacochlear electrical stimulation of one and two electrode pairs in the human cochlea," *J. Acoust. Soc. Am.*, vol. 86(3), pp. 971–980, 1989.

[6] R. Shannon, "Forward masking in patients with cochlear implants," *J. Acoust. Soc. Am.*, vol. 88, pp. 741–744, 1990.

[7] P. Blamey and G. Dooley, "Pattern recognition and masking in cochlear implant patients," *J. Acoust. Soc. Am.*, vol. 97, pp. 271–278, 1993.

[8] W. Nogueira, A. Buchner, T. Lenarz, and B. Elder, "A psychoacoustic NofM-type speech coding strategy for cochlear implants," *J. Appl. Sig. Process*, vol. 18, pp. 3044–3059, 2005.

[9] M. Hughes and L. Stille, "Psychophysical and physiological measures of electrical-field interaction in cochlear implants," *J. Acoust. Soc. Am.*, vol. 125, pp. 247–260, 2009.

[10] H. Duifhuis, "Consequences of peripheral frequency selectivity for nonsimultaneous masking," *J. Acoust. Soc. Am.*, vol. 54(6), pp. 1471–1973, 1973.

[11] L. Cohen, L. Richardson, E. Saunders, and R. Cowan, "Spatial spread of neural excitation in cochlear implant recipients: comparison of improved ecap method and psychophysical forward masking," *Hear. Res.*, vol. 179, pp. 72–87, 2003.

[12] P. Abbas, M. Hughes, C. Brown, C. Miller, and H. South, "Channel interaction in cochlear implant users evaluated using the electrically evoked compound action potential," *Audiol. Neurootol.*, vol. 9(4), pp. 203–213, 2004.

[13] H. Schuknecht, *Pathology of the Ear.* Lea and Febiger, 1993.

[14] P. Loizou, "Introduction to cochlear implants," *IEEE Signal Process. Mag.*, vol. 18(1), pp. 32–42, 1999.

[15] F. Zeng, S. Rebscher, W. Harrison, X. Sun, and H. Feng, "Cochlear implants: system design, integration, and evaluation," *IEEE Rev. Biomed. Eng.*, vol. 1, pp. 115–142, 2008.

[16] R. Plomp, "Rate of decay of auditory sensation," *J. Acoust. Soc. Am.*, vol. 36(2), pp. 277–282, 1964.

[17] R. Shannon, "Two-tone unmasking and suppression in a forward-masking situation," *J. Acoust. Soc. Am.*, vol. 59(6), pp. 1460–1470, 1976.

[18] D. Harris and P. Dallos, "Forward masking of auditory nerve fiber responses," *J. Neurophysiol.*, vol. 42(4), pp. 1083–1107, 1979.

[19] W. Jesteadt, S. Bacon, and J. Lehman, "Forward masking as a function of frequency, masker level, and signal decay," *J. Acoust. Soc. Am.*, vol. 71, pp. 950–962, 1982.

[20] R. Shannon, "Multichannel electrical stimulation of the auditory nerve in man. i. basic psychophysics," *Hear. Res.*, vol. 11, pp. 157–189, 1983.

[21] R. Wegel and C. Lane, "The auditory masking of one sound by another and its probable relation to the dynamics of the inner ear," *Phys. Rev.*, vol. 23, pp. 266–285, 1924.

[22] N. Kiang and E. Moxon, "Physiological considerations in artificial stimulation of the inner ear," *Ann. Otol. Rhinol. Laryngol.*, vol. 81, pp. 714–730, 1972.

[23] D. Nelson and G. Donaldson, "Psychophysical recovery from single-pulse forward masking in electric hearing," *J. Acoust. Soc. Am.*, vol. 109(6), pp. 2921–2933, 2001.

[24] D. Nagal, "Compound action potential of the cochlear nerve evoked electrically," *Arch. Otol. Rhinol. Laryngol.*, vol. 206, pp. 293–298, 1974.

[25] V. Prijs, "On peripheral auditory adaptation. ii. comparison of electrically and acoustically evoked action potentials in the guinea pig," *Acustica*, vol. 45, pp. 1–13, 1980.

[26] R. Charlet de Sauvage, Y. Cazals, J. Erre, and J. Aran, "Acoustically derived auditory nerve action potentials evoked by electrical stimulation: An estimation of the waveform of single unit contribution," *J. Acoust. Soc. Am.*, vol. 73(2), pp. 616–627, 1983.

[27] P. Stypulkowski and C. van den Honert, "Physiological properties of the electrically stimulated auditory nerve. i. compound action potential recordings," *Hear. Res.*, vol. 14, pp. 205–223, 1984.

[28] J. Aran, J. Erre, and R. Charlet de Sauvage, "Derived evoked potentials for continuous tones using a hybrid electrical-acoustical stimulation," *Hear. Res.*, vol. 20, pp. 289–293, 1985.

[29] J. Aran, J. Erre, H. Hiel, and P. Goeury, "Distribution of viii nerve excitation by pure tones derived by electrical stimulation and acoustic masking," *Acta. Otolaryngol.*, vol. 103, pp. 593–601, 1987.

[30] C. Brown, P. Abbas, and B. Gantz, "Electrically evoked whole-nerve action potentials: Data from human cochlear implant users," *J. Acoust. Soc. Am.*, vol. 88(3), pp. 1385–1391, 1990.

[31] B. Charasse, H. Thai-Van, C. Berger-Vachon, and L. Collet, "Assessing auditory nerve refractory function with a modified subtraction method: results and mathematical modeling," *Clin. Neurophysiol.*, vol. 114, pp. 1344–1450, 2003.

[32] C. Brown, P. Abbas, J. Borland, and M. Bertschy, "Electrically evoked whole nerve action potentials in ineraid cochlear implant users: Responses to different stimulating electrode configurations and comparison to psychophysical responses," *J. Speech Hear. Res.*, vol. 39(3), pp. 453–467, 1996.

[33] C. Brown, P. Abbas, and B. Gantz, "Preliminary experience with neural response telemetry in the Nucleus CI24M cochlear implant," *Am. J. Otol.*, vol. 19(3), pp. 320–327, 1998.

[34] P. Abbas, C. Brown, J. Shallop, J. Firszt, M. Hughes, S. Hong, and S. Staller, "Summary of results using the Nucleus CI24M implant to record the electrically evoked compound action potential," *Ear Hear.*, vol. 20(1), pp. 45–59, 1999.

[35] L. Smith and F. Simmons, "Estimating eighth nerve survival by electrical stimulation," *Ann. Otol. Rhinol. Laryngol.*, vol. 92, pp. 19–23, 1983.

[36] R. Jyung, J. Miller, and S. CAnnon, "Evaluation of eighth nerve integrity by the electrically evoked middle latency response," *Otolaryngol. Head Neck Surg.*, vol. 101, pp. 670–682, 1989.

[37] R. Hall, "Estimation of surviving spiral ganglion cells in the deaf rat using the electrically evoked auditory brainstem response," *Hear. Res.*, vol. 45, pp. 123–136, 1990.

[38] L. Cohen, E. Saunders, and L. Richardson, "Spatial spread of neural excitation: comparison of compound action potential and forward-masking data in cochlear implant recipients," *Int. J. Audiol.*, vol. 43, pp. 1–10, 2004.

[39] C. Miller, P. Abbas, and B. Robinson, "Response properties of the refractory auditory nerve fiber," *J. Assoc. Res. Otolaryngol.*, vol. 2(3), pp. 216–232, 2001.

[40] W. Lai and N. Dillier, "A simple two-component model of the electrically evoked compound action potential in the human cochlea," *Audiol. Neurootol.*, vol. 5, pp. 333–345, 2000.

[41] C. Miller, P. Abbas, and C. Brown, "An improved method of reducing stimulus artifact in the electrically evoked whole-nerve potential," *Ear Hear.*, vol. 21(4), pp. 280–290, 2000.

[42] C. Miller, P. Abbas, J. Rubinstein, B. Robinson, A. Matsuoka, and G. Woodworth, "Electrically evoked compound action potentials of guinea pig and cat: Responses to monopolar, monophasic stimulation," *Hear. Res.*, vol. 119, pp. 142–154, 1998.

[43] M. Hughes and L. Stille, "Psychophysical versus physiological spatial forward masking and the relation to speech perception in cochlear implants," *Ear Hear.*, vol. 29, pp. 435–452, 2008.

[44] H. Cullington, "Preliminary neural response telemetry results," *Br. J. Audiol.*, vol. 34(3), pp. 131–140, 2000.

[45] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[46] F. Zeng, H. Chen, and S. Han, "Temporal masking in electric hearing," *J. Assoc. Res. Otolaryngol.*, vol. 6, pp. 390–400, 2005.

[47] O. Macherey, R. Carlyon, A. van Wieringen, J. Deeks, and J. Wouters, "Higher sensitivity of human auditory nerve fibers to positive electrical currents," *J. Assoc. Res. Otolaryngol.*, vol. 9(2), pp. 241–251, 2008.

[48] B. Kwon and C. von den Honert, "Spatial and temporal effects of interleaved masking in cochlear implants," *J. Assoc. Res. Otolaryngol.*, vol. 10, pp. 447–457, 2009.

[49] H. Levitt, "Transformed up-down methods in psychoacoustics," *J. Acoust. Soc. Am.*, vol. 49, pp. 467–477, 1971.

[50] M. Chatterjee, "Effects of stimulation mode on threshold and loudness growth in multielectrode cochlear implants," *J. Acoust. Soc. Am.*, vol. 105, pp. 850–860, 1999.

[51] M. Nilsson, S. D. Soli, and J. A. Sullivan, "Development of the hearing in noise test for the measurement of speech reception thresholds in quiet and in noise," *J. Acoust. Soc. Am.*, vol. 95(2), pp. 1085–1099, 1994.