RESEARCH ON SPOKEN LANGUAGE PROCESSING
Progress Report No. 28 (2007)
Indiana University

Cochlear Implant Simulations: A Tutorial on Generating Acoustic Simulations for Research[1]

Jeremy L. Loebach

Speech Research Laboratory
Department of Psychological and Brain Sciences
Indiana University
Bloomington, Indiana 47405

---

[1] This work supported by NIH-NIDCD Training Grant T32-DC00012 and NIH-NIDCD Research Grant R01-DC00111. I would like to thank Dr. Qian-Jie Fu for developing and maintaining TigerCIS, a tool that has become invaluable to my research.

Cochlear Implant Simulations: A Tutorial on Generating Acoustic Simulations for Research

Abstract. Acoustic simulations of cochlear implants have become a common tool available to researchers interested in many aspects of speech perception and cognition. Although many ways exist to create such simulations, all are derived from a common philosophical and physiological base. The primary goal of this tutorial is to instruct the reader how to make cochlear implant (CI) simulations for use in research. In order to realize this goal, the various methodologies and signal processing techniques will be reviewed, since the varied techniques used will determine the behavioral outcomes. Finally, the reader will be led through a systematic demonstration using a stand-alone simulator (TigerCIS).

Introduction

Although explicit cochlear implant simulations have been around for only about a decade, the technology and signal processing techniques that gave rise to them have existed for the better part of a century. Most modern simulations are based on the principle of the vocoder, pioneered by Homer Dudley at Bell Labs (Dudley, 1939). The vocoder was a speech synthesizer that passed the acoustic signal through a series of band pass filters, which derived the energy profiles for each band. The spectrum of each band was replaced with a synthetic source (tones and noisy hisses) which were then modulated using the energy profile appropriate for the original band. The result was a stimulus that sounded highly artificial, but was surprisingly intelligible as speech. Most modern CI simulations are based, in part, on the philosophy of Dudley's vocoder.

Although multi-channel synthesis techniques have been around since the 1930's, the development of multi-channel cochlear implants did not occur until far later. Research into the electrical stimulation of the cochlear nerves has had a long history in both animal and human models. Experiments with single electrodes inserted into the cochlear partition had been ongoing since the early 1960's, but the development of single channel cochlear implants for humans did not occur until the early 1970's, finally gaining FDA approval in the United States in 1984. While much of the research into single channel implants occurred in the U.S., the development of multi-channel implants first occurred in Australia. In fact, in 1985, the Cochlear multi-channel implant received U.S. FDA approval, barely a year after the approval of the House/3M single channel implant. This effectively put an end to single channel implants, though there are still advocates of the single channel short electrode implants (c.f. House, 1994).

Today, only 3 companies have FDA approval to produce and market cochlear implants in the United States: Med-EL, Cochlear (makers of the Nucleus series of implants) and Advanced Bionics (makers of the Clarion series of implants). Other models may be available for experimental purposes and clinical trials, and other brands do exist in European and other foreign markets. The basic design of each implant is similar, although different brands use different numbers of electrodes and different signal processing techniques to provide electrical stimulation. Most currently available implants contain six to twenty-four channels. Theoretically, increasing the number channels that are available in the electrode array will increase the amount of acoustic information that will be available to the user. Practically, however, more channels do not necessarily translate to better performance due to surgical, anatomical, and physiological constraints.

How a Cochlear Implant Works

Cochlear implants rely on the anatomy and physiology of the cochlea for their functionality. In the normal hearing ear, vibrations in the air are translated into pressure waves in the fluid of the cochlea. The basilar membrane is differentially displaced by these pressure waves depending on their frequency. The width, thickness and stiffness of the basilar membrane (BM) vary depending on longitudinal position. At the base of the cochlea (closest to the stapes) the BM is thin, narrow and tight requiring high frequency oscillations to displace it. At the apex (farthest from the stapes), the BM is wide, thick and loose, requiring low frequency oscillations to displace it. The mode of basilar membrane displacement bears some similarity to the strings of a guitar. Thin tight strings like the high E vibrate very fast to produce high-pitched sounds, whereas thick loose strings like the low E vibrate very slowly to produce low pitch sounds. The displacement patterns of the BM essentially work in reverse: high frequency sounds best displace the BM at the base, whereas low frequency sounds best displace the BM at the apex.

On top of the basilar membrane sits the organ of Corti, which contains the hair cells necessary for transduction. The cell bodies of the hair cells are embedded in the surrounding tissue so that only the stereocilia are exposed to the fluid inside the organ of Corti. These stereocilia move back and forth with the motion of the fluid (like seaweed moving back and forth with ocean waves) opening mechanically gated ion channels. The influx of potassium ions (K+) changes the resting potential of the cell, stimulating the release of neurotransmitter onto the neurons of the spiral ganglion at the base of the hair cell, effectively transducing the sound into neural impulses. The neurons of the spiral ganglion (known by many different names: auditory nerve, cochlear nerve, vestibulocochlear nerve, VIII (8th) cranial nerve) then carry the electrical impulses to the brainstem for processing.

Although the hair cells themselves do not have a frequency preference, they are situated along the length of the basilar membrane and will be displaced in a frequency dependent fashion. This means that the specific hair cells, and consequently the neurons of the spiral ganglion, will respond to specific frequencies. If the apical region of the BM is displaced, indicating a low frequency sound, the hair cells overlying this region of the BM will be stimulated, transmitting information to the brain that a low frequency sound occurred. If, on the other hand, the basal region of the BM is displaced indicating a high frequency sound, the hair cells overlying this region of the BM will be stimulated, sending information to the brain that a high frequency sound occurred. Thus, one could examine the responses of the neurons in the spiral ganglion, and based on the frequency of tones they respond best to, trace them back to the region of the cochlea and indeed the hair cells that stimulated them. This is a very important property of the spiral ganglion: it has a tonotopic organization that arises from the location of the hair cells in the organ of Corti. Due to the law of specific nerve energies, one could stimulate a single neuron in the spiral ganglion and produce the sensation of a tone of a particular pitch in the brain. This is the fundamental mechanism that cochlear implants exploit.

In sensorineural hearing loss, the normal processes of transduction are disrupted. While a myriad of etiologies can lead to specific deficits in the mechanisms of transduction, one of the most common causes of deafness occurs when the hair cells are damaged or destroyed. Since hair cells cannot regenerate, they can no longer stimulate the neurons of the spiral ganglion, so information about sound cannot get into the brain. Despite the destruction of their primary sources of input, the neurons of the spiral ganglion are far more robust, and typically do not atrophy following cochlear insult. This is a key mechanism that cochlear implants utilize: although the mechanisms of transduction have failed, the neurons of the spiral ganglion are intact, and can be electrically stimulated. The electrode array of the cochlear implant stimulates the surviving spiral ganglion neurons in order to produce the sensation of sound in the brain. More importantly, cochlear implants do not rely on extensive pre-processing; rather they simply provide an alternate form of input, allowing the ascending auditory pathway to function as it normally would.

An electrode array inserted into the round window of the cochlea brings electrical contacts close to spiral ganglion neurons. Pulses of electrical current stimulate the neurons to produce a sensation of sound in the brain. An external microphone takes in the sound, and sends it to a speech processor, which divides the acoustic information into a number of channels, changing the energy in each channel into a digital code. The coded signal is then transmitted through the skin to a receiver implanted behind the ear, which in turn sends the information to the appropriate electrodes in the cochlea as a series of current pulses. Electrodes are organized tonotopically, so that high frequency channels are located basally, and low frequency channels are located apically. Thus, the electrical impulses are delivered to the appropriate tonotopic region of the cochlea to evoke the sensation of the appropriate pitch. In the normal hearing ear, a single hair cell will stimulate a limited number of spiral ganglion neurons (typically 3-7), resulting in a very discrete encoding of frequency. Such discrete stimulation is impossible with modern implants, however, since there is some distance between the stimulation site and the target neurons. Due to spread of electrical current in the fluid and tissue of the cochlea, many more neurons are stimulated by each electrode than would be stimulated in the intact cochlea. This causes a widening of the frequency representation, which removes much of the spectral fine structure and frequency resolution of the normal auditory signal. Despite such spectral reduction, many modern cochlear implant users demonstrate very high levels of speech recognition under quiet conditions, suggesting that the limited frequency information provided by electrical stimulation is sufficient to transmit information about speech.

Acoustic Simulations of Cochlear Implants

Although no one with normal hearing can truly know what the world sounds like through a cochlear implant, simulations can approximate the experience by processing the acoustic signal in a manner similar to that of an implant's speech processor. There are three main components to such simulations: the frequency channel, the amplitude envelope and the carrier signal. Each of these can be altered to produce differential effects in perception.

The Channel. Since cochlear implant processors divide the acoustic information into a limited number of channels, simulations typically filter the acoustic signal into different frequency bands using band pass filters. The number of frequency bands that are used depends on the type of implant being modeled. For example, one-band simulations can be used to model single channel implants (Van Tasell, Soli, Kirby, & Widin, 1987), and multiple band simulations can be used to model multi-channel implants (Shannon, Zeng, Kamath, Wygonski & Ekelid, 1995; Fu & Shannon, 1999). The band pass filters are typically broad in order to limit the spectral detail in manner similar to that of a cochlear implant, and the filter bandwidth will ultimately depend on the total number of channels used (Figure 1). In single channel simulations, the band pass filter will typically cover most of the frequency range of the stimulus. In multiple channel simulations, the band pass filters are as broad as necessary to cover the range of the spectrum. In some cases, the center frequencies are selected based on known measurements such as the articulation index (AI) (French & Steinberg, 1947), or locations of critical bands in the cochlea (Greenwood, 1961; Greenwood, 1990). Channel selection is not limited by any known function, so the cutoff frequencies may vary according to the task one wishes to design. The overall selection of cutoff frequencies and width of the band pass filters should approximate reasonable values for a cochlear implant processor in order to be a maximally applicable model, however. Specific models can be made to simulate particular implant processors (CIS processor, SPEAK processor) or even specific implantees (based on their frequency sensitivity).

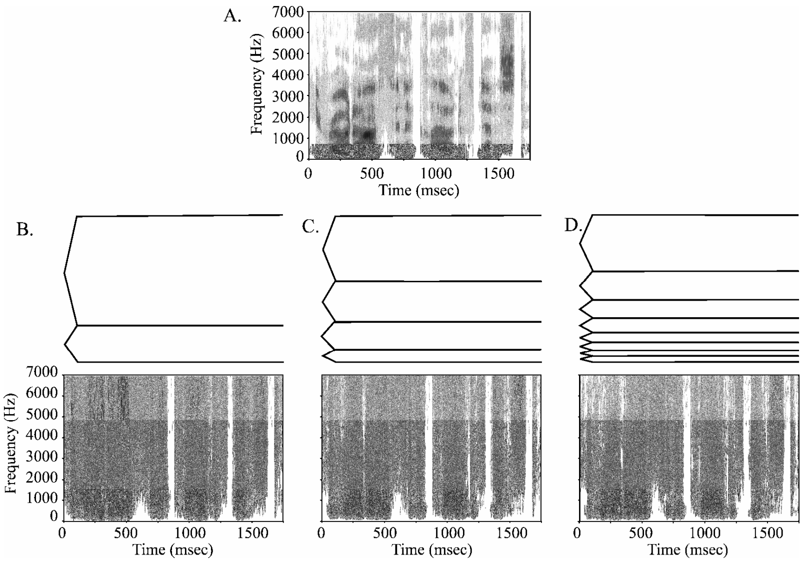

FIGURE 1: Band pass filtering allows one to simulate the limited number of channels available in a cochlear implant. The spectrogram for the naturally produced version of the sentence “He rode off in a cloud of dust” as produced by a male speaker is presented in (A). The band pass filters required to generate noise vocoded stimuli with 2 (B), 4 (C) and 8 (D) spectral channels and the spectrographic representation of the stimuli generated by them appear below. Although the majority of the spectral detail is absent from the noise vocoded stimuli, the overall spectral profiles become increasingly similar to the naturally produced stimulus as the number of bands increases from two to eight. Changes in the temporal information, however, remain unchanged across the noise vocoded stimuli.

The Envelope. Although the output of a cochlear implant is spectrally degraded, it is temporally precise. This means that the temporal modulations are not limited by any known mechanism. Typically, individuals with cochlear implants can discriminate temporal modulations upwards of 300 Hz (i.e. modulations changing at a rate of 300 times per second) (Fu & Shannon, 2000). The amplitude envelope can be thought of as a trace around the time domain waveform of a sound, and is extracted by using a low pass filter (Figure 2). How closely you trace the sound will depend on the frequency cutoff of the low pass filter. If you choose a filter with an upper frequency limit of 160 Hz, only the changes in amplitude that occur at a rate of less than 160 times per second will be preserved. Envelope derivation will also determine how well subjects perform; if you make the cutoff too low, you will undermine performance (see Rosen, 1992 for a review of the temporal cues in the envelope).

Whether modeling single or multi channel implants, the amplitude envelope must be derived from each band. One must first divide the spectrum into frequency bands using the band pass filters, and then use a low pass filter to derive the amplitude envelope from each. This means that if one were designing a 22-channel simulation, one would filter the signal into 22 channels and then extract the amplitude envelope from each. Since the energy in the spectrum will change depending on the stimulus (e.g. as is the case for consonants and vowels, where a high frequency fricative may be followed by a low frequency vowel), deriving the envelope for each band will preserve the energy in the appropriate spectral regions.

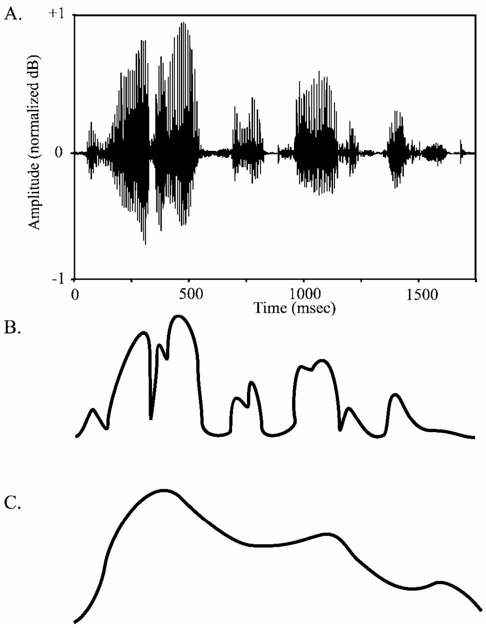

FIGURE 2: Envelope detection is based on low pass filtering the time domain waveform of the stimulus (A). The amount of temporal information preserved depends on the upper limit of the low pass filter. Higher frequency cutoffs (~300 Hz) preserve smaller changes in the time intensity information (B) than lower frequency cutoffs (~25 Hz), which only preserve course temporal changes (C). Low pass filters with sufficiently high cutoff frequencies can preserve rudimentary pitch information, if the pitch falls within the pass band.

The Carrier. Although the spectrum has been filtered into a discrete number of bands, the residual spectral information must be removed to make an effective CI simulation. This is typically done by replacing the spectral content of each band with some other signal. Two carriers have been traditionally used: white noise and sinusoids (Figure 3). White noise has proven to be a successful carrier signal for CI simulations, and the earliest pioneering work with CI simulations typically used noise based carriers to model single channel implants (Van Tasell, Soli, Kirby & Widin, 1987; Van Tasell, Greenfield, Logemann & Nelson, 1992) and multi channel implants (Blamey, Dowell, Tong, Brown, Luscombe & Clark, 1984a; Blamey, Dowell, Tong & Clark, 1984b; Shannon, Zeng, Kamath, Wygonski & Elekid, 1995). Noise is an effective way to remove the spectral detail from the frequency channels, but it may over-represent the information in the band. Recall that the electrodes in the cochlear implant do not stimulate all neurons in a given region of the cochlea evenly. If the stimulation patterns were reflective of the noise-based carrier then each electrode would be very broad, and each channel would be contiguous, with neither gaps nor overlaps between each. Physiologically, this may not be the case, since the electrodes are typically narrow, and stimulation is more likely to be focused at a given frequency, and roll off to both sides due to current spread and electrical diffusion. As such, other simulations have opted to use sinusoids as their carrier signals (Dorman & Loizou, 1997; Dorman & Loizou, 1998; Loizou Dorman, & Tu, 1999). Each sinusoid is focused at the channel center frequency, and rolls off in intensity with a given slope on either side. Indeed, some evidence for using sinewaves as carriers come from CI users themselves, who have described the sensation of electrical stimulation as being a series of beep tones rather than noise pulses (Dorman, Loizou & Rainey, 1997).

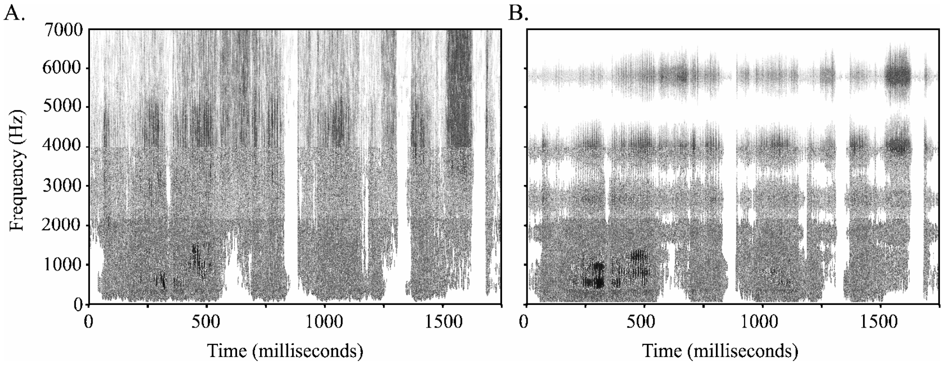

FIGURE 3: Spectrograms of the sentence "He rode off in a cloud of dust." produced using noise band carriers (A) and sinusoidal carriers (B). Although the stimuli sound qualitatively distinct, and appear to differ in overall spectral profile, studies have demonstrated that speech recognition in quiet does not appear to differ between stimuli produced with noise band and sinewave carriers.

No matter which type of carrier signal is selected, speech recognition under each type of stimulation is virtually identical (Dorman, Loizou & Rainey, 1997). Although sinusoids may offer a more realistic simulation, and indeed people with cochlear implant report that the world sounds more metallic than noisy (c.f. Chorost, 2005), the behavioral results are identical whether using noise or sinusoids, at least for speech in quiet.

Software. There are many different methods of simulating cochlear implants. Applications exist for MatLAB and stand-alone programs such as TigerCIS have also been developed for use. Any program that is capable of band and low pass filtering and can generate noise and pure tone sinusoids can be used (indeed I have done this all by hand using a Cool Edit Pro, MS Excel, and MS Notepad, see Loebach, 2005). The basic concept is simple: divide the spectrum into bands, derive the envelope from each band, replace the spectrum in each band with your carrier signal, modulate the carrier with the amplitude envelope appropriate for that band and reassemble the bands. Although similar in concept, the exact steps that you will take will depend on the specific program that you use.

One of the easiest simulators to use is TigerCIS, which was developed by Dr. Qian-Jie Fu at the House Ear Institute. It produces quality simulations using either noise band or sinewave carrier signals, and is very user friendly. It is also very malleable, allowing one to change the settings for the amplitude envelope derivation, band pass filtering, number of channels, and output filters allowing one to "shift" the electrodes in any manner. Moreover, TigerCIS is freely available on the internet at http://www.tigerspeech.com/.

How To

TigerCIS can process any Windows PCM .wav file of a standard format. Typical formats include 22,050 Hz, 16-bit stereo or mono files. If you are editing your wav files in Adobe Audition or Praat, be sure that no additional proprietary information is being placed in the header, as this will cause the file to be unreadable in TigerCIS. Occasionally, errors will be introduced if each track of the stereo .wav file contains slightly different information. This can result in a doubling of the frequency information, and an apparent upward shift in pitch. Anecdotally, I have found that mono .wav files yield the best results, and avoid processing artifacts. In addition, it is critical to level the stimuli prior to processing them with the simulator. The simulator will process whatever information is present, and if the stimuli have different

RMS amplitude values, you may under specify some portions of the signal, and over specify others. Ensuring that the stimuli are at a comparable level prior to simulation will provide you with the most accurate and appropriate simulations. The systematic demonstration below was generated using TigerCIS version 1.04.03. Although the software is frequently updated with additional features, the basic structure of the following steps will be preserved in some manner in most versions.

<u>Step 1: Select Processor</u> This step allows you to specify which type of processor you want to use. For a standard CI simulation, be sure the "Noise or sinewave vocoder" box is checked, and the "FFT-based" box is un-checked.

<u>Step 2: Select Carrier Type for Vocoder</u> This step allows you to specify the carrier that you wish to use, either sinusoids centered at the bands center frequency or noise bursts.

<u>Step 3: Select Number of Channels</u> This step allows you to set your analysis and synthesis channel filters. Although the number of channels is unlimited, the more that you use, the longer it will take to synthesize the stimuli. If you want a straightforward CI simulation the number of channels simulated should match the number of spectral channels. By varying the degree and direction of mismatch, one can make compressive simulations (the number synthesized is smaller than the number of spectral channels) or expansive (the number synthesized is larger than the number of spectral channels) simulations.

<u>Step 4: Set Analysis Filter Type</u> This step determines the method used to divide the spectrum. The default setting is based on Greenwood's function (c.f. Greenwood, 1990), which is essentially a pitch to distance map of the basilar membrane based on critical bandwidth. You can create your own mapping function as well, if you want to model a specific type of processor with specific band cutoffs. You can even program your own simulator based on patient data using their frequency channel cutoffs and electrode mappings.

<u>Step 5: Set Analysis Filter Variables</u> This step runs hand and hand with the previous. If you chose a custom filter in step 4, you should load that file here. If not, you can set your lower frequency cutoff (the lowest frequency information that you want to include), upper frequency cutoff (the highest frequency information that you want to include) and the roll off function (dB/octave) which tells the program how sharp to make the filters. For example, if you have a center frequency of 100 Hz and you roll off at 24 dB/octave, it means that as you increase the frequency by 1 octave (from 100 to 200 Hz in this case) the signal strength would drop by 24 dB. This avoids the summation of information where the band pass filters overlap (which would distort the signal, possibly causing clicks). The larger this number is the faster you roll off, and the steeper your filter slopes will be. If your roll off is too slow, channel overlap and summation can occur. If your roll off is too fast you may get transients and other distortions.

<u>Step 6: Envelope Detection</u> This step allows you to design the low pass filter to derive the amplitude envelope from each band. The first number is the cutoff frequency for the upper limit. This number should ideally be below 400 Hz, since this is the upper limit detectible by individuals with cochlear implants. If your low pass cutoff is too low (6-12 Hz) you may not adequately be representing the temporal information available to the cochlear implant user. The roll off function is similar to that used in the previous step and determines how quickly the information drops off as you increase in frequency. Select the roll off carefully, since it will affect the intelligibility (too shallow and you can include more information than is available in a CI).

<u>Step 7: Set Carrier Filter Type</u> This step is the same as in Step 4, except you are now selecting how you want to process the carrier signal. If you are using noise, this will divide the noise spectrum into bands in

a manner similar to the analysis filters in Sep 4. You can either use the Greenwood function, or model a specific implant or patient by specifying a custom filter.

Step 8: Set Carrier Filter Variables This step is the same as in Step 5, except you are now specifying the frequency range for the carrier signal. You can use values similar to or different from the filters used in Step 5, depending on what type of simulation you wish to create. If you want to make a 1:1 simulation, the frequency range should be identical to that in Step 5. If you want to simulate a basal shift of 2 millimeters based on electrode insertion depth, for instance, you would simply change the upper and lower frequency limits by a function that would equate a 2 mm shift along the basilar membrane. Since the tonotopic organization of the basilar membrane is non-uniform across frequency, the lower frequency cutoff would shift upward by about 180 Hz, and the upper frequency cutoff would shift upward by about 1800 Hz. See Zwicker for a more complete treatment (Zwicker, Flottorp & Stevens, 1957).

Step 9: Processing Now that all of the information has been specified, it is time to process the stimuli. Stimuli can be processed individually by loading a single wave file at the top, or in a batch by specifying input and output directories at the bottom. For processing single stimuli simply press the “Start Simulation” button to begin. You can then press “Play Simulation” to listen to what you just created, or “Save Simulation” to save it as a .wav file. Each time you update any of the information in Steps 1-8 you will have to reprocess the stimuli with the current settings (TigerCIS will maintain the previous settings until you exit the program).

Although there are several other useful options available in TigerCIS, the ones presented here are necessary for making standard cochlear implant simulations. The program is very malleable, allowing you to make your own creations (or even to make real time simulations of your own voice), and an online forum is available at http://www.tigerspeech.com/ for support.

Conclusions

Modern cochlear implant simulations may allow the research scientist some insight into the perceptual experiences of a person with a cochlear implant. One must never assume, however, that they are in fact reproducing those actual perceptual experiences. To that end, the closer the variables used to design the simulations are to the specifications of real cochlear implants, the more readily the results in normal hearing subjects will generalize to cochlear implant users themselves. One should be thoughtful in the selection of synthesis variables, and be able to defend their choices in publication. Whenever possible, results should be normalized by comparing findings in normal hearing subjects with those from cochlear implant users, so that a maximum benefit be can achieved for researcher, clinician and patient alike.

References

Blamey, P.J., Dowell, R.C., Tong, Y.C., Brown, A.M., Luscombe, S.M. & Clark, G.M. (1984a). Speech processing studies using an acoustic model of a multiple-channel cochlear implant, *Journal of the Acoustical Society of America, 76*, 104-110.

Blamey, P.J., Dowell, R.C., Tong, Y.C. & Clark, G.M. (1984b). An acoustic model of a multiple-channel cochlear implant, *Journal of the Acoustical Society of America, 76,* 97-103.

Chorost, M. (2005). *Rebuilt: How becoming part computer made me more human*. Houghton Mifflin: New York.

Dorman, M.F. & Loizou, P.C. (1997). Speech intelligibility as a function of the number of channels of stimulation for normal-hearing listeners and patients with cochlear implants, *American Journal of Otology, 18,* S13-S114.

Dorman, M.F. & Loizou, P.C. (1998). The identification of consonants and vowels by cochlear implant patients using a 6-channel continuous interleaved sampling processor and by normal-hearing subjects using simulations of processors with two to nine channels, *Ear and Hearing, 19,* 162-166.

Dorman, M.F., Loizou, P.C. & Rainey, D. (1997). Speech intelligibility as a function of the number of channels of stimulation for signal processors using sine-wave and noise-band outputs, *Journal of the Acoustical Society of America, 102,* 2403-2411.

Dudley, H. (1939). Remaking speech, *Journal of the Acoustical Society of America, 11,* 169-177.

French, N.R. & Steinberg, J.C. (1947). Factors governing the intelligibility of speech sounds, *Journal of the Acoustical Society of America, 19,* 90-119.

Fu, Q.J. and Shannon, R.V. (1999). Recognition of spectrally degraded and frequency-shifted vowels in acoustic and electric hearing, *Journal of the Acoustical Society of America, 105,* 1889-1900.

Fu, Q.J. & Shannon, R.V. (2000). Effect of stimulation rate on phoneme recognition by nucleus-22 cochlear implant listeners, *Journal of the Acoustical Society of America, 107,* 589-597.

Greenwood, D.D. (1990). A cochlear frequency-position function for several species—29 years later, *Journal of the Acoustical Society of America, 87,* 2592-2605.

House, W.F. (1994). *Cochlear Implants: My perspective*. All Hear Incorporated: Portland. http://www.allhear.com/monographs/m-95-htm.html

Loebach, J.L. (2005). Temporal aspects of speech: The encoding of naturally produced and spectrally reduced synthetic speech by the auditory nerve. Unpublished doctoral dissertation, University of Illinois at Urbana-Champaign.

Loizou, P.C., Dorman, M., & Tu, Z. (1999). On the number of channels needed to understand speech, *Journal of the Acoustical Society of America, 106*, 2097-2103.

Rosen, S. (1992). Temporal information in speech: Acoustic, auditory and linguistic aspects, *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences, 336*, 367-373.

Shannon, R.V., Zeng, F.G., Kamath, V., Wygonski, J. & Ekelid, M. (1995). Speech recognition with primarily temporal cues, *Science, 270,* 303-304.

Van Tasell, D.J., Greenfield, D.G., Logemann, J.J. & Nelson, D.A. (1992). Temporal cues for consonant recognition: Training, talker generalization, and use in evaluation of cochlear implants, *Journal of the Acoustical Society of America, 92,* 1247-1257.

Van Tasell, D.J., Soli, S.D., Kirby, V.M. & Widin, G.P. (1987). Speech waveform envelope cues for consonant recognition, *Journal of the Acoustical Society of America, 82,* 1152-1161.

Zwicker, E., Flottorp, G. & Stevens, S.S. (1957). Critical band width in loudness summation, *Journal of the Acoustical Society of America, 29,* 548-557.