Published in final edited form as:

Ear Hear. 2013 ; 34(6): 763–772. doi:10.1097/AUD.0b013e3182967534.

Feasibility of Real-Time Selection of Frequency Tables in an Acoustic Simulation of a Cochlear Implant

Matthew Fitzgerald¹, Elad Sagi¹, Tasnim A. Morbiwala², Chin-Tuan Tan¹, and Mario A. Svirsky¹

¹New York University School of Medicine, New York, NY

²Precor Corp

Abstract

Objectives—Perception of spectrally degraded speech is particularly difficult when the signal is also distorted along the frequency axis. This might be particularly important for post-lingually deafened recipients of cochlear implants (CI), who must adapt to a signal where there may be a mismatch between the frequencies of an input signal and the characteristic frequencies of the neurons stimulated by the CI. However, there is a lack of tools that can be used to identify whether an individual has adapted fully to a mismatch in the frequency-to-place relationship and if so, to find a frequency table that ameliorates any negative effects of an unadapted mismatch. The goal of the proposed investigation is to test the feasibility of whether real-time selection of frequency tables can be used to identify cases in which listeners have not fully adapted to a frequency mismatch. The assumption underlying this approach is that listeners who have not adapted to a frequency mismatch will select a frequency table that minimizes any such mismatches, even at the expense of reducing the information provided by this frequency table.

Design—34 normal-hearing adults listened to a noise-vocoded acoustic simulation of a cochlear implant and adjusted the frequency table in real time until they obtained a frequency table that sounded “most intelligible” to them. The use of an acoustic simulation was essential to this study because it allowed us to explicitly control the degree of frequency mismatch present in the simulation. None of the listeners had any previous experience with vocoded speech, in order to test the hypothesis that the real-time selection procedure could be used to identify cases in which a listener has not adapted to a frequency mismatch. After obtaining a self-selected table, we measured CNC word-recognition scores with that self-selected table and two other frequency tables: a “frequency-matched” table that matched the analysis filters with the noisebands of the noise-vocoder simulation, and a “right information” table that is similar to that used in most cochlear implant speech processors, but in this simulation results in a frequency shift equivalent to 6.5 mm of cochlear space.

Results—Listeners tended to select a table that was very close to, but shifted slightly lower in frequency from the frequency-matched table. The real-time selection process took on average 2–3 minutes for each trial, and the between-trial variability was comparable to that previously observed with closely-related procedures. The word-recognition scores with the self-selected table were clearly higher than with the right-information table and slightly higher than with the frequency-matched table.

Corres. Author Contact: matthew.fitzgerald@nyumc.org.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conclusions—Real-time self-selection of frequency tables may be a viable tool for identifying listeners who have not adapted to a mismatch in the frequency-to-place relationship, and to find a frequency table that is more appropriate for them. Moreover, the small but significant improvements in word-recognition ability observed with the self-selected table suggest that these listeners based their selections on intelligibility rather than some other factor. The within-subject variability in the real-time selection procedure was comparable to that of a genetic algorithm, and the speed of the real-time procedure appeared to be faster than either a genetic algorithm or a simplex procedure.

I. Introduction

When fitting a prosthetic device to a hearing-impaired individual, many aspects of the fitting process are often streamlined, with the result being that patients are often programmed with a standard set of parameters. For example, in the case of cochlear implants, the audiologist often devotes considerable time toward obtaining values at which the signal may be heard comfortably. This is a crucial step given that these values are known to vary widely across patients. In contrast, other aspects of the CI fitting, such as the frequency-allocation table, are generally left to the manufacturer-recommended settings. As a result, virtually all patients who utilize a given speech processor model are programmed with the same frequency table despite the likelihood that the frequency-to-place-of-stimulation function may be shifted, expanded or compressed with respect to that found in normal hearing listeners. This approach is predicated upon the concept that the “normal” frequency-to-place map is irrelevant to prelingually deaf patients, and that postlingually deaf patients can adapt to shifts, expansions, or compressions. This approach is supported anecdotally by the speech-perception scores exhibited by most cochlear implant users (e.g., Balkany et al., 2007), and directly from numerous reports from experimental data indicating that both normal-hearing (e.g., Rosen et al., 1999; Fu and Galvin, 2003; Faulkner et al., 2006;) and users of cochlear implants (Fu et al., 2002; Svirsky et al., 2004; Reiss et al., 2007; Reiss et al., 2012) can adapt to frequency shifts with experience. While some of these data indicate that the adaptation was incomplete, on the whole there is clear evidence that listeners can adapt at least in part to mismatches in the relationship of the input acoustic frequency to the site of stimulation within the cochlea.

While the majority of patients can likely adapt to the standard frequency table recommended by the manufacturer settings, there are also data that suggest that there are limitations in the ability of users of cochlear implants to adapt to frequency mismatches. For example, two of five users of the Med-El Combi 40+ electrode array significantly improved their vowel and consonant identification scores when the apical electrodes were deactivated and the input frequency range was reassigned to the remaining electrodes, and, 4/5 of these listeners chose to use one of the experimental tables after completing the experiment (Gani et al., 2007). The authors speculated that the initial frequency table suffered from apicalward shift (e.g., when a given electrode stimulates neurons whose characteristic frequency is lower than what is assigned by the frequency table), which was reduced by deactivating some apical electrodes, although it is also possible that the improvement could have resulted from a decrease in channel interaction in the apex of the cochlea. Further evidence that cochlear implant recipients may be unable to adapt to any frequency shift comes from data from Fu et al., 2002, in which three experienced cochlear implant recipients used an experimental frequency table that, relative to their standard frequency table, had a considerably larger “basalward shift” (this is the term used to describe the situation when a given electrode stimulates neurons whose characteristic frequency is higher than the frequency range assigned to that electrode by the frequency table that is programmed in the speech processor). Over three months of use, their word-recognition scores with the experimental table initially dropped considerably relative to their standard table, and then improved over

several weeks until reaching an asymptotic level. Despite this improvement, however, this asymptotic performance still remained worse than with their standard table, suggesting that while these listeners had adapted fully to their standard frequency table, they only adapted partially to the experimental table that imposed an additional frequency shift (Sagi et al., 2010). Unfortunately, it is impossible to establish directly whether a user of a cochlear implant has fully adapted to a frequency mismatch between the frequency range assigned by the speech processor to a given electrode and the characteristic frequency of the stimulated neurons, because the latter cannot be measured noninvasively.

One possible way to address this issue, however, would be to allow the patient to select a preferred frequency table with the intent of maximizing speech intelligibility. The assumption underlying this approach is that if a patient selects a frequency table that differs from the standard, then one could assume that the patient has not fully adapted to a potential frequency mismatch imposed by the standard frequency table. This approach, however, requires a tool that can be used to reliably obtain a frequency table that was selected with the intent of maximizing speech intelligibility. To date there have been very few efforts directly addressing this issue, although there are two procedures that have been previously investigated which could theoretically be used to do so. These procedures are the simplex procedure (e.g., Neuman et al., 1987; Preminger et al., 2000; Franck et al., 2004; Amlani and Shafer, 2009) and the genetic algorithm (e.g., Holland, 1975; Wakefield et al., 2005; Baskent and Edwards, 2007; Baskent et al., 2007).

The simplex procedure is an adaptive procedure that makes use of paired comparisons between different stimuli that differ along a pre-specified dimension, such as the amount of gain a hearing aid produces in a given frequency region. A matrix (often 5×5) is created in which each cell corresponds to a signal with given characteristics. As one moves along the x- or y-axis, the value corresponding to that axis (for example, low-frequency gain), subsequently increases or decreases. A paired-comparison forced-choice procedure is then used to compare different cells in the matrix, with successive comparisons between cells being determined by an adaptive procedure. For example, two sentences that varied along a given dimension would be presented consecutively, and the listener would report which signal was more intelligible. This decision would then determine the next set of comparisons, until a 'winning' cell with a given set of stimulus parameters has been identified. This approach has been used primarily to fit hearing-aid gain (e.g., Neuman et al., 1987; Kuk and Pape, 1992; Stelmachowicz et al., 1994; Preminger et al., 2000), although it has also been applied to other features such as noise reduction or spectral enhancement (Franck et al., 2004; Franck et al., 2007).

Genetic algorithms (Holland, 1975), are a closely-related procedure that also makes use of comparisons that can be used to determine the optimal fitting parameters for a hearing aid or a cochlear implant. In this procedure, a 'gene' is a combination of different stimulus parameters that are pre-specified by the experimenter. For example, a previous investigation that utilized an acoustic simulation of a cochlear implant used as 'genes' the number of noise-vocoder channels, the magnitude of frequency shift, and the total frequency range (Baskent et al., 2007). Different genes are then compared with one another (via a paired-comparison procedure, for example) and the outcome of that comparison is used to determine whether a given 'gene' is used in the next set of comparisons according to the genetic algorithm used. This approach is theoretically very flexible, and thus far has been utilized with acoustic simulations of a cochlear implant (Baskent and Edwards, 2007; Baskent et al., 2007) as well as with cochlear implant patients themselves in both research (Wakefield et al., 2005) and clinical environments (Govaerts et al., 2010; Varenberg et al., 2011; Holmes et al., 2012).

In the present investigation, we explored the feasibility of a new tool that could theoretically be used to determine whether a given implant user has fully adapted to the frequency table programmed into their device, and if not, to provide a recommendation to the audiologist for altering the frequency table accordingly. In this tool, listeners can adjust the frequency table in real time while listening to running speech to obtain a set of filterbanks that are perceived by the listener as optimizing speech intelligibility; such a tool is defined here as a *real-time selection procedure*. As an initial step toward establishing the feasibility of this tool, we had 34 normal-hearing individuals listen to an acoustic simulation of a cochlear implant, and then use the real-time selection procedure to choose a set of filterbanks that were perceived by the listener as maximizing speech intelligibility. These listeners subsequently performed a word-recognition task using this “*listener-selected*” table, and also using two other frequency tables. Acoustic models of a cochlear implant allow the precise selection of frequency ranges for analysis filters and vocoder noise bands, allowing us to test speech perception with known amounts of frequency mismatch. This is impossible in CI users because the characteristic frequency of the neurons that are stimulated by a given electrode cannot be determined with any degree of precision. Thus, conducting the present experiments using an acoustic model of a CI is necessary to address the questions of interest. Furthermore, our choice of listeners without any previous exposure to a CI simulation is crucial to assess the feasibility of our real-time selection tool to identify individuals who have not adapted to a frequency mismatch. Our use of this tool is similar in intent to the simplex procedure and genetic algorithms, but differs from these approaches in at least three key ways. First, those procedures rely almost exclusively on paired-comparisons, while a real-time selection procedure does not do so. In the latter approach, a listener has the option to directly compare two frequency tables, but that is not a requirement for evaluating, selecting, or discarding a given table. This may allow listeners to more rapidly discard frequency tables that are clearly unintelligible. Second, a real-time selection process allows listeners to make their selections based on running speech, rather than pre-processed sentences. This in turn allows the listener, if they choose, to listen to more than just a single sentence with a given table before making a decision as to its viability. Similarly, in the same amount of time, an individual might be able to listen to more sentences with a real-time approach than with an approach based on paired comparisons of a single sentence. Both approaches might be beneficial for some patient populations, or for individuals who need additional time or information before making a decision. The final difference is that in both the simplex procedure and genetic algorithm approach, there is a clear criterion for the end of the self-selection process, as both approaches end after the adaptive criterion has been reached. In contrast, a real-time approach as described here is more open-ended, and terminates when the participant decides he or she has determined which signal is most intelligible.

II. Experiment 1: Methods

A. Listeners

Twenty-four listeners (15 females) between the ages of 19 and 47 (mean 27.5 years, SD = 7.91) participated in this experiment. All listeners were native speakers of English and reported having normal hearing in both ears. No listener reported any history of hearing or learning disorder, nor did they have any previous experience listening to speech through a noise vocoder. All listeners were compensated for participating in this experiment. This research was approved by the NYU Medical Center IRB.

B. Acoustic simulation of a cochlear-implant

The signal processing for the acoustic simulation of the CI was adapted from the methods employed by Kaiser and Svirsky (2000). The signal was low-pass filtered with a cutoff

frequency of 20,000 Hz and digitized at a rate of 48,000 samples per second. A bank of sixth order bandpass Butterworth filters was then used to divide the auditory signal into eight frequency channels. The cutoff frequencies of these analysis filters could be changed in real time using the computer keyboard while the speech processing algorithm was running. The analysis filter cutoff frequencies were based on a filter bank similar to those used with 8-channel cochlear implants (249–6820 Hz; see Table 1), and several variations of that filter bank. Some of those variations had the same “span” as the canonical filter bank from Table 1, but a different “midpoint.” Other variations had a different span and the same midpoint, and some filter banks differed from the canonical one in both span and midpoint. Here, span is defined by first converting both the maximum and minimum frequency of the filter bank into mm of distance from the cochlear apex, using Greenwood’s equation (Greenwood, 1990) and then subtracting the minimum from the maximum. The lowest frequency of the canonical filter bank (249 Hz) corresponds to a distance of 6.7 mm from the cochlear apex, assuming a 35 mm long basilar membrane; likewise, the highest frequency of the canonical filter bank (6,820 Hz) corresponds to 27.1 mm. Therefore, the span of the canonical filter bank is 20.4 mm (that is, the difference between 27.1 mm and 6.7 mm) and the midpoint is 16.9 mm. In addition to this filter bank we created 35 other filter banks with the same approximate span of 20.4 mm but with midpoints ranging from 10.7 mm to 26.2 mm. Table 2 shows the total frequency range of these 36 filter banks. Then, for each midpoint we created seven additional filter banks with spans ranging approximately from 16 mm to 26.5 mm, in seven increments of about 1.5 mm. As an example, Table 3 shows the total frequency range for the 8 filter banks that share the approximate midpoint of 16.9 mm. Thus, a total of 288 filter banks were created and arranged in a matrix with 8 rows (representing eight different spans) and 36 columns (representing thirty-six different midpoints). By pressing different keys, listeners adjusted two parameters that allowed them to roam that matrix. When the matrix row was changed, the span changed while the midpoint of the filter bank remained largely unchanged (keys F9 and F11 on the keyboard shifted the span narrower or broader). Similarly, when the matrix column was changed, the filter bank midpoint changed while the span remained largely unchanged (keys Page Up and Page Down shifted the midpoint higher or lower). Thus, the span and midpoint of the filter bank could be adjusted independently.

For each analysis filter, the temporal envelope was extracted by half wave rectification and second order low pass filtering at 160 Hz. The temporal envelopes were then used to modulate eight noise bands that spanned the range from 851 to 16,982 Hz. When expressed in mm of displacement along the cochlea, this range represents a basalward shift of about 6.5 mm with respect to the canonical filterbank. Throughout the entire experiment, the frequency range of the noise bands remained fixed. The fixed noise bands simulate the percepts evoked by different intracochlear electrodes whose position is, obviously, also fixed.

C. Selection of listener tables

Listeners were seated in a sound-treated room, and were presented with running speech from a book on CD (Howard Zinn, *Stories Hollywood Never Tells*, 2001). This CD was sampled at 44.1 kHz, and had two different male speakers at different points in the recording. The CD was processed online in real time through our noise vocoder, and presented via loudspeakers. The listeners were instructed that the speech would sound somewhat distorted, and that they were to adjust the sound quality of that speech until it sounded “most intelligible”. All participants were given a simple visual representation of the frequency range of the analysis filters available on a computer monitor. A picture of this visual representation is available in Fig. 1: it is a dark grey rectangle embedded within a larger, light gray rectangle. The light gray rectangle was fixed and it represented the whole

frequency range that is audible to humans, and the dark grey rectangle represented the frequency range of the active filter bank. As the listener adjusted the midpoint and the span of the analysis filter bank, the dark gray rectangle would correspondingly move up, down, expand or contract. In this way, listeners obtained a self-selected table that was perceived to maximize speech intelligibility. Prior to beginning the self-selection process, listeners were presented with running speech, and the experimenter moved slowly from the lowest to the highest possible filterbank to illustrate the interface and how the sound quality would change with manipulations of the frequency table. After this brief familiarization process, the listeners were allowed to adjust the frequency table at their own pace, thereby enabling them to rapidly bypass tables with poor intelligibility, and spend more time choosing which table allowed for maximal speech understanding. This selection process was repeated three times with different starting points for the frequency table. One starting point was the previously noted canonical frequency table similar to those commonly used by CI recipients; the analysis filters in this table ranged from 249 to 6820 Hz. The second starting point ranged from 120 to 2680 Hz, and the final starting point ranged from 1728 to 20000 Hz. The order in which the three starting points were presented was randomized across listeners, and the end of a given trial was determined by the listener when they felt they had identified a frequency table that they perceived to maximize speech intelligibility, and was confirmed verbally by the listener. For an individual listener, we defined the *listener-selected* frequency table as the table whose span and midpoint are the average of the span and midpoints of the three frequency tables that they selected.

D. Speech-recognition testing

Immediately after obtaining a self-selected frequency table, listeners were tested on their ability to understand speech processed through the noise vocoder with three different frequency tables. One of these tables was the *listener-selected* frequency table described above. The second frequency table was called the *frequency-matched* table. This table employed the same frequency range used by the output noise bands (851 to 16982 Hz) and the same intermediate cutoff frequencies. Thus, with the *frequency-matched* table there was a perfect match between the analysis filters in the filter bank and the noisebands that were amplitude modulated by the output of those analysis filters. The *frequency-matched* table simulates a situation in which there is a perfect match between the frequency bands in the filter bank and the characteristic frequency of the neurons that are stimulated by the corresponding electrode in a user of a cochlear implant.

The final frequency table was named the *right-information* table. In this table, which simulates the standard clinical approach, the analysis filters ranged from 250 to 6820 Hz, while the noiseband range remained 851 to 16982 Hz. Therefore, the analysis filters in this table provide low-frequency information, but do so by mapping it to noisebands that are higher in frequency. The *right-information* table simulates a situation thought to commonly occur in cochlear implant patients, in which low-frequency information is mapped to neurons tuned to higher frequencies given the 20 mm average electrode insertion depth (Ketten et al., 1998) of a cochlear implant and the characteristic frequency of the spiral ganglion cells that are stimulated by the electrodes (Stakhovskaya et al., 2007). Thus, the *right-information* table implements a basalward shift. In the present experiment, this table is shifted by 6.5mm relative to the frequency-matched table (according to the scale of Greenwood, 1990).

For each frequency table, listeners were tested on two 50-word CNC word lists. Note that each word in the CNC lists has three phonemes, resulting in 150 phonemes per list. We used a standard recording (Peterson and Lehiste, 1962) for the CNC word lists. These lists were presented to the listeners by a custom computer program that also recorded the typewritten

responses. Listeners were instructed to repeat the target word, and their response was recorded by the experimenter prior to beginning the next trial. However, no feedback was provided to the listener. The number of words and the number of phonemes correctly identified using each frequency table were counted and subsequently converted to a percent correct value. The testing procedure was as follows: Listeners were presented with one CNC list for each frequency table, with the presentation order of the different frequency tables being determined by a pseudo-randomized Latin-square design. After completing testing with all three tables, the listeners were given a short break. Following the break, the order of the frequency tables was reversed, and listeners were re-tested on a different CNC list for each table. All word lists were presented at 70 dB SPL via loudspeakers in a sound-treated booth.

E. Statistical analyses

To determine whether the average *listener-selected* table differed significantly from either the *right-information* or the *frequency-matched* tables, we computed 99% confidence intervals on both the low- and high-frequency edges of the *listener-selected* table. We assumed in our calculations that the numbers were normally distributed, and this assumption was borne out by the data. Once making this determination, we computed confidence intervals separately for the low- and high frequency edge values by applying the equation $X = t(n-1) * STE$, where t = the values of the t distribution with $n-1$ degrees of freedom, and STE = the standard error of the mean. We then assessed whether the *right-information* or the *frequency-matched* tables fell outside of the range produced by these 99% confidence intervals. If so, then we concluded that these tables were significantly different from the *listener-selected* table.

To identify whether word-and phoneme-recognition scores obtained with the *listener-selected* table were different from those obtained with the *right-information* and *frequency-matched* tables, we performed two separate one-way analyses of variance (ANOVA) with repeated measures across the three frequency tables. One ANOVA was done on the word-recognition scores, and the other on the phoneme-recognition scores. Both ANOVAs included the combined percent-correct identifications from the two CNC word lists.

III. Experiment 1: Results

A. Frequency tables

Figure 2 shows the frequency ranges of the analysis filters for the *right-information* table (far left black column), the *frequency-matched* table (far right hashed column), the 24 individual *listener-selected* tables (narrow gray columns) and the mean *listener-selected* table (thick gray column). The present data show that, for normal-hearing adults listening to a noise vocoder with noise bands ranging from 851 to 16982 Hz, the analysis filters of the listener-selected table were on average 671 to 14527 Hz. Therefore, they were quite close to, but shifted slightly lower in frequency from the noise bands or, equivalently, from the *frequency-matched* table, whose filter passbands are identical to the noise bands. Regarding the low-frequency edge of the *listener-selected* tables, 23 of the 24 listeners selected frequency tables in which the low-frequency edge of the analysis filters (ranging from 481 to 958 Hz) fell between those of the *right-information* (250 Hz) and *frequency-matched* (851 Hz) tables. Furthermore, the 99% confidence interval of the low-frequency edge of the *listener-selected* table (ranging from 605 to 737 Hz) failed to overlap with the low-frequency edge of either the *right-information* or *frequency-matched* tables, although it was closer to the latter than the former. A similar result was observed for the high-frequency edge of the *listener-selected* tables, as 19 of the 24 listeners selected values for the analysis filters that fell between the high-frequency edge of the *right-information* (6820 Hz) and

frequency-matched (16982 Hz) tables. As with the low-frequency edge of the listener-selected tables, the 99% confidence interval of the high-frequency edge of the *listener-selected* table (12243 to 15470 Hz) fell outside the high frequency edge of the *right-information* and *frequency-matched* tables, again indicating that listeners prefer a table that fell between those two. Finally, while it should be noted that the mean *listener-selected* table was distinctly different from the other two tables, it is clear that it is considerably closer in frequency to the *frequency-matched* table than to the *right-information* table.

Listeners were able to reliably select a frequency table within a few minutes, and this selection process was quite consistent regardless of the initial starting point. Figure 3 depicts the selection process for a representative listener tested in the present experiment. The black bars represent the frequency tables selected over time when the selection process began with highest frequency table, while the gray bars depict the tables selected when the process began at its lowest frequency table. This listener rapidly progressed through many frequency tables before converging on a range that was rather close to the final table he chose. He then spent most of his time fine-tuning the table before making a final decision. For this individual, this selection process took less than five minutes. Other listeners displayed a similar pattern of behavior, and similarly fast selection times.

B. Speech understanding

The current data suggest that normal-hearing adults listening to a noise vocoder understand speech better when they are allowed to select their own frequency table. Figure 4 displays the mean word- and phoneme-recognition scores for each of the three tables. In both figures, scores for the *right-information*, *frequency-matched*, and *listener-selected* tables are shown in the left, middle, and right columns, respectively. In each instance, speech understanding is best with the *listener-selected* tables. The scores for the *frequency-matched* tables are nearly as good as those of the *listener-selected* tables, whereas the scores from the *right-information* tables are considerably worse than either of the other two frequency tables. These impressions are confirmed by statistical analysis, as ANOVAs on both the word- ($F_{2,71} = 80.553, p < 0.001$) and phoneme-recognition scores ($F_{2,71} = 183.98, p < 0.001$) revealed a significant main effect of frequency table. Post-hoc Holm-Sidak tests revealed that the word- and phoneme-recognition scores obtained with the listener-selected tables were significantly higher than those obtained with either the *right-information* ($p < 0.001$) or *frequency-matched* tables ($p = 0.032$), while the scores obtained with the *frequency-matched* table are significantly better than those observed with the *right-information* table ($p < 0.001$ in each case).

IV. Experiment 1: Discussion

The present data provide some evidence for the feasibility of using a real-time selection tool to identify cases in which individuals have not fully adapted to a frequency shift, and to provide a recommendation that may benefit speech understanding in that circumstance. Here, when presented with a sizeable frequency shift in an acoustic simulation of a cochlear implant and asked to choose a signal that was thought to maximize speech intelligibility, all 24 listeners were able to do so. Moreover, they selected filterbanks that closely approximated, but were slightly lower in frequency than those used in the *frequency-matched table* (851 to 16982 Hz). This result appears quite reasonable given that these listeners had not been given time or training that would enable them to adapt to this magnitude of frequency shift. Moreover, the observation that listeners were able to use the real-time selection procedure to obtain these tables suggests that this procedure has the potential to identify cases in which a listener has been unable to adapt to a mismatch in the frequency-to-place relationship. Finally, the small but significant improvements in speech understanding observed with the self-selected table relative to the frequency-matched table

suggests that these listeners did in fact base their selections on the intelligibility of the signal.

Here, the observation that the average low- and high-frequency edges were shifted slightly, but significantly lower in frequency relative to the *frequency-matched* table suggests that our naïve listeners were likely trying to reconcile the need for low-frequency information, which pushed down the low frequency edge of their selected table, with their own internal template for how speech should sound, which led them to select a table that is as close as possible to being frequency-matched. The observation that the self-selected tables more closely approximated the frequency-matched table almost certainly reflects the lack of experience these normal-hearing listeners had with frequency-shifted, noise-vocoded speech, as it is well documented that frequency-to-place shifts greater than approximately 3 mm in cochlear space hinder speech intelligibility with inexperienced listeners (Fu and Shannon, 1999). Conversely, if these listeners were trained to understand frequency-shifted speech, it is possible, if not likely that they would have selected different frequency tables than those observed here.

It is notable that the CNC word understanding scores were slightly higher, but significantly better for the *listener-selected* table than for the *frequency-matched* table¹. This result is consistent with previous reports suggesting that naïve listeners can not only tolerate a small degree of basalward shift (e.g., Dorman et al., 1997; Fu and Shannon, 1999), but that they can actively make use of this small shift to improve their speech-understanding abilities (see also Baskent and Shannon, 2005). We speculate that the difference in scores between the *listener-selected* and the *frequency-matched* tables most likely reflect the presence of additional information about the first formant frequency (F1) in the *listener-selected* table. As noted previously, the mean low-frequency edge of the *listener-selected* table is lower in frequency than the *frequency-matched* table (671 Hz vs. 851 Hz, respectively), resulting in additional information about the speech signal for these listeners. This additional low-frequency information increases the likelihood that F1 would be represented in the processed acoustic signal, and therefore seemed to be a reasonable explanation for the improvement in performance from the *frequency-matched* to the *listener-selected* tables. Regardless of the precise reason for the improved performance with the *listener-selected* table, the noteworthy implication here was that listeners used the real-time selection tool to obtain a frequency table that yielded improvements in speech understanding. Whether recipients of cochlear implants who have not fully adapted to the frequency table programmed in their speech processor can also use a real-time selection tool to maximize speech understanding needs to be fully explored in cochlear implant users. Suggesting this may be the case, however, are data using a different technique that reveal that self-selection of frequency tables may be useful (Jethanamest et al., 2010). In this regard, it is worth noting that the frequency table chosen by a given CI recipient might change as they adapt to the stimulation provided by their device over the first few weeks or months after implantation, potentially making this tool useful for tracking adaptation in these individuals¹.

Finally, while we believe that the present data support the feasibility of a real-time selection process, there are alternate interpretations that cannot be entirely ruled out by use of the current experimental paradigm. First, while we interpret the small but significant improvements in speech perception observed with the *listener-selected* table as evidence that the listeners here made their selections on the basis of intelligibility, we cannot rule out the

¹The poor scores in the right-information table were expected, as numerous reports have indicated that listeners requires practice to perform well in this condition, and the current experimental paradigm utilized naïve listeners precisely because these individuals would not have adapted to the 6.5 mm frequency mismatch.

possibility that listeners performed better on the listener-selected condition because of the additional exposure to this condition conferred by the self-selection process, even though this process encompassed many different frequency tables, and not just the final self-selected table used for speech testing. By this logic, any improvements seen in the self-selected condition reflect rapid learning in the listener-selected condition that occurred in the absence of explicit training or feedback, and that failed to generalize to either the frequency-matched or the right-information tables. A second concern with the experimental paradigm is that listeners may have used visual cues on the screen to aid their self-selection process, and that the results observed here are thus not based solely on the acoustic signal. By this line of thought, after completing the initial self-selection process, the listeners may have visually remembered the approximate location within the grid where their self-selected tables lay, and used that visual representation to guide their subsequent selections. To rule these possibilities out, we conducted a second experiment with an additional 10 listeners.

V. Experiment 2

VI. Experiment 2: Methods

A. Listeners—Ten listeners (5 females) between the ages of 18 and 39 (mean 25.7 years, $SD = 7.36$) participated in this experiment. As with the first experiment, all listeners were native speakers of English and reported normal hearing in both ears with no history of learning disability or any prior experience listening to noise-vocoded speech.

B. Selection of frequency tables and speech testing—All frequency-table selections were made in the same way as in Experiment 1 with two small differences. First, all selections were timed by a stopwatch so as to determine exactly how much time was required to make each selection. Second, after making an initial selection with visual input, the monitor was turned off, and the listener then repeated the self-selection process without any visual input. All speech testing was conducted using the same protocol described in Experiment 1, except that we spoke conversationally for 5 minutes prior to testing with each frequency table. This allowed the listener to have the same amount of experience with each frequency table prior to testing, with the goal of ensuring that any differences in speech perception between different frequency tables did not reflect rapid, incidental learning.

VII. Experiment 2: Results

There are three notable results from Experiment 2. First, the frequency-table selections obtained without visual input closely resembled those obtained in Experiment 1. Figure 5 shows the individual and mean self-selected tables obtained from the 10 listeners tested here. As with Experiment 1, these selections were on average shifted slightly lower in frequency than the frequency-matched table, but were clearly higher than the right-information table. Specifically, for 8 of the 10 listeners tested here, the low-frequency edge of the self-selected table was lower than that of the frequency-matched table. Thus, the lack of visual input did not appear to affect the average selection made by these listeners.

The second noteworthy result from Experiment 2 is that the time required for these ten listeners to select a preferred frequency table ranged between 30 seconds and 5 minutes. Table 4 indicates the amount of time it took for listeners to select preferred frequency table for each of the four runs. Eight of 10 listeners required more time to perform the self-selection task the first time they performed the task, even though that selection was made with the visual input as a guide. Subsequent selections made without visual input generally took less time, presumably due to familiarity with the interface and the self-selection procedure itself. As a general rule, all selections made after the first required a similar amount of time.

Finally, the pattern of word- and phoneme-recognition scores obtained here (Figure 6) closely resembled that seen with Experiment 1 (see Fig. 4). As before, the best scores were observed with the self-selected table. Slightly lower scores were seen with the frequency-matched table, and the lowest scores were seen with the right-information table. Results of a one-way RMANOVA were consistent with this interpretation. As with Experiment 1, a significant effect was observed across the three frequency tables for the word- ($F_{2,18} = 96.317$, $p < 0.001$) and phoneme-recognition scores ($F_{2,18} = 209.372$, $p < 0.001$). Post-hoc Holm-Sidak tests indicated that the word- and phoneme-recognition scores obtained with the self-selected and frequency-matched table were both significantly higher than with the right-information table ($p < 0.001$ in each case). The only difference between these data and Experiment 1 is that in Experiment 2, the difference between the self-selected and frequency-matched table was not significant ($p = 0.092$).

VIII. Experiment 2: Discussion

The results of this experiment are notable in two regards. First, even when listeners had no visual input, they were able to select a frequency table that closely resembled those obtained when visual input was available. This result, when coupled with the improved word-recognition scores seen in Experiment 1 with the self-selected table, strongly suggests that listeners based their selections on the intelligibility of the auditory speech signal, with little, if any input from other factors.

Second, the pattern of word-recognition scores was virtually identical across Experiments 1 and 2: the best scores were observed in the listener-selected condition, slightly worse scores in the frequency-matched condition, and poor scores in the right-information condition. The only difference between the two studies was that in Experiment 1, significantly higher word-recognition scores were observed with the self-selected table than with the frequency-matched table, while in Experiment 2, this difference was not significant. We think the most likely explanation for minor discrepancy between experiments was the number of listeners participating in each experiment (24 in Experiment 1 vs. 10 in Experiment 2). Thus, we interpret these speech-perception data as evidence that the higher word-recognition scores in the listener-selected condition reflect the additional low-frequency information provided by this table, rather than a rapid, incidental learning associated with slightly increased exposure to this condition via the self-selection process. Here, in Experiment 2, listeners received an equal amount of experience with each condition prior to testing (e.g., a few minutes of conversational speech), yet we observed the same general pattern of results as in Experiment 1.

IX. General discussion

Taken together, the data from Experiments 1 and 2 suggest that self-selection of frequency tables may be a useful tool to determine when individuals have not fully adapted to mismatches in the frequency-to-place relationship. Moreover, these data indicate that judgments made with the self-selection process as well as subsequent speech-understanding abilities appear to be based on the acoustic characteristics of the signal and not on other non-auditory factors. For such a tool to be feasible to utilize in patient populations, it must be both fast and reliable, and the present data with normal-hearing listeners selecting frequency tables in an acoustic simulation of a cochlear implant suggests this may be the case.

While it is difficult to directly compare the time required to complete the self-selection procedure with that utilized by other optimization procedures, the speed at which the present data were obtained suggest that it should at minimum be no slower than other optimization procedures. Here, the average amount of time needed to converge on a preferred frequency table was approximately 3 minutes when listeners had little experience with the real-time

procedure. In contrast, the durations previously reported with other optimization procedures range from a few minutes to several hours of testing (Neuman et al., 1987; Preminger et al., 2000; Saunders and Cienkowski, 2002; Wakefield et al., 2005; Baskent et al., 2007; Holmes et al., 2012). On the surface, this suggests that the real-time procedure may be faster than other optimization procedures. However, previous utilizations of other optimization procedures have generally employed large search spaces with multiple parameters; in the present experiments, the listener only manipulated the frequency table. Moreover, in some instances the optimization procedure may not have been optimized for practical application. As an example, when a genetic algorithm was used to optimize speech perception in a CI simulation, it was noted by the authors that the time needed to converge on a preferred setting (25 minutes in this instance) could have been reduced by changing some variables of the algorithm (Baskent et al., 2007). For these reasons, it is difficult to directly compare the speed of our real-time procedure with that of other optimization procedures. However, given that the real-time procedure only required a few minutes per run, we speculate that it would be at minimum no slower than other optimization procedures.

The present data also suggest that while there is some within-subject variability in the selections made by each individual, the magnitude of that variability is similar to that previously observed with a genetic algorithm. Figure 5 depicts all selections made from the additional 10 listeners who completed Experiment 2. From this figure it is evident that within a given listener, there can sometimes be a sizeable amount of variability with regard to the self-selected frequency table. For example, on average the largest deviation between the highest and lowest frequency table was 242 Hz for the low-frequency edge of the self-selected table. A visual inspection of the data reveals, however, that while there may be considerable variability in the individual table selections, the general trend within a given individual was similar. For example, in eight of ten listeners, the low-frequency edge of the self-selected table either matched or was lower than the low-frequency edge of the frequency-matched table across all four individual selections. Moreover, in the remaining two listeners, only a single selection of the four had a low-frequency value higher than that of the frequency-matched table. Perhaps more important, the between-run variability for a given individual was similar to that previously observed for a genetic algorithm that was used to select parameters with the intent of maximizing speech intelligibility in a cochlear implant simulation that also had frequency shifted speech (Baskent et al., 2007). In that study, a ‘best theoretical solution’ was known (e.g., one that maximized the number of available channels, and minimized the amount of frequency shift or expansion/compression of the signal), but listeners did not always select this ‘best’ solution. More important, in the individual examples provided in that manuscript, frequency shifts of 3–4 mm in cochlear space were selected by an individual listener between different runs of the algorithm. To compare the variability observed with their genetic algorithm and the real-time approach used here, we took our listener (L26) with the largest between-run difference in the low-frequency edge of the table, and converted the low-frequency edge of the runs with the highest and lowest frequency to mm in cochlear space according to the Greenwood equation (Greenwood, 1990). The results of that analysis indicated a 3 mm shift between the selections with the highest and lowest low-frequency edge, suggesting that at worst, the variability observed in the self-selection procedure is similar to that previously seen with a genetic algorithm (Baskent et al., 2007).

It is interesting to note that Baskent et al. (2007) observed some runs of the algorithm that yielded significantly worse speech-understanding abilities than their ‘optimal’ solution, but these occurrences were generally rare. Here, we also observed sporadic instances in which the real-time procedure yielded frequency tables that would be expected to result in poorer word-recognition scores than either a frequency-matched table, or the mean self-selected table. For example, L26 and L33 both had single runs in which the listener selected a table

with a low-frequency edge of 964 Hz, while L30 had a single run in which the low-frequency edge was 441 Hz. In these listeners, these runs appeared to be outliers in that they differed from the rest of the results obtained from that listener. While we did not measure word-recognition abilities with each run, given the additional loss of low-frequency information (L26 and L33), or the increased amount of basalward shift relative to their other selections (L30), it is reasonable to assume that these runs might have yielded poorer word-recognition scores than the other selections made by these individuals, as was previously reported by Baskent et al (2007).

In summary, the present results support the hypothesis that listeners exposed to noise vocoded speech can enhance speech intelligibility by manipulating the frequency table. Given the noise bands used in this study, the speech intelligibility enhancement was achieved by a compromise between minimizing frequency mismatch and adding first formant frequency information. Lastly, the frequency table selection process was completed within five minutes. Taken together, these results support the feasibility of applying the frequency table selection approach to clinical populations.

Acknowledgments

This work was supported by NIH/NIDCD grants DC09459 (PI: Fitzgerald), and DC03937 (PI: Svirsky).

References

- Amlani AM, Shafer EC. Application of paired-comparison methods to hearing aids. *Trends Amplification*. 2009; 13:241–259.
- Balkany T, Hodges A, Menapace C, Hazard L, Driscoll C, Gantz B, Kelsall D, Luxford W, McMenomy S, Neely JG, Peters B, Pillsbury H, Roberson J, Schramm D, Telian S, Waltzman S, Westerberg B, Payne S. Nucleus Freedom North American Clinical Trial. *Otolaryngol Head Neck Surg*. 2007 May; 136(5):757–762. [PubMed: 17478211]
- Baskent D, Shannon RV. Speech recognition under conditions of frequency compression and expansion. *J. Acoust. Soc. Am*. 2005; 117:1405–1416. [PubMed: 15807028]
- Baskent D, Edwards B. Simulating listener errors in using genetic algorithms for perceptual optimization. *J. Acoust. Soc. Am*. 2007; 121:EL238–EL243. [PubMed: 17552575]
- Baskent D, Eiler CL, Edwards B. Using genetic algorithms with subjective input from human subjects: implications for fitting hearing aids and cochlear implants. *Ear Hear*. 2007; 28:370–380. [PubMed: 17485986]
- Dorman MF, Loizu PC, Rainey D. Simulating the effect of cochlear implant insertion depth on speech understanding. *J. Acoust. Soc. Am*. 1997; 102:2993–2996. [PubMed: 9373986]
- Faulkner A, Rosen S, Norman C. The right information may matter more than frequency-place alignment: Simulations of frequency-aligned and upward shifting cochlear implant processors for a shallow array insertion. *Ear. Hear*. 2006; 27:139–152. [PubMed: 16518142]
- Franck BA, Dreschler WA, Lyzenga J. Methodological aspects of an adaptive multidirectional pattern search to optimize speech perception using three hearing-aid algorithms. *J. Acoust. Soc. Am*. 2004; 116:3620–3628. [PubMed: 15658712]
- Franck BA, Boymans M, Dreschler WA. Interactive fitting of multiple algorithms implemented in the same hearing aid. *Int. J. Audiol*. 2007; 46:388–397. [PubMed: 17680471]
- Fu Q-J, Shannon RV. Recognition of spectrally degraded and frequency-shifted vowels in acoustic and electric hearing. *J. Acoust. Soc. Am*. 1999; 105:1889–1900. [PubMed: 10089611]
- Fu Q-J, Galvin JJ 3rd. The effects of short-term training for spectrally mismatched noise-band speech. *J. Acoust. Soc. Am*. 2003; 113:1065–1072. [PubMed: 12597199]
- Fu QJ, Shannon RV, Galvin JJ 3rd. Perceptual learning following changes in the frequency-to-electrode assignment in the Nucleus 22 cochlear implant. *J. Acoust. Soc. Am*. 2002; 112:1664–1674. [PubMed: 12398471]

- Gani M, Valentini G, Sigrist A, Kos MI, Boex C. Implications of deep electrode insertion on cochlear implant fitting. *J. Assoc. Res. Otolaryngol.* 2007; 8:69–83. [PubMed: 17216585]
- Govaerts PJ, Varenberg B, De Ceulaer G, Daemers K, De Beukelaer C, Schauwers K. Development of a software tool using deterministic logic for the optimization of cochlear implant processor programming. *Otol. Neurotol.* 2010; 31:908–918. [PubMed: 20418791]
- Greenwood DD. A cochlear frequency-position function for several species – 29 years later. *J. Acoust. Soc. Am.* 1990; 87:2592–2605. [PubMed: 2373794]
- Holland, JH. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence.* Arbor, MI: University of Michigan Press; 1975.
- Holmes AE, Shrivastav R, Krause L, Siburt HW, Schwartz E. Speech based optimization of cochlear implants. *Int. J. Audiol.* 2012; 51:806–816. [PubMed: 22978753]
- Jethanamest D, Tan C-T, Fitzgerald MB, Svirsky MA. A new software tool to optimize frequency table selection for users of cochlear implants. *Otol. Neurotol.* 2010; 31:1242–1247. [PubMed: 20729774]
- Kaiser AR, Svirsky MA. Using a personal computer to perform real-time signal processing in cochlear implant research. *Proceedings of the IXth IEEE-DSP Workshop.* 2000 Oct.:15–18.
- Ketten DR, Skinner MW, Wang G, Vannier MW, Gates GA, Neely JG. In vivo measures of cochlear length and insertion depth of Nucleus cochlear implant electrode arrays. *Ann. Otol. Rhinol. Laryngol.* 1998; 175(Suppl):1–16.
- Kuk FK, Pape NM. The reliability of a modified simplex procedure in hearing-aid frequency-response selection. *J. Speech. Hear. Res.* 1992; 35:418–429. [PubMed: 1573880]
- Neuman AC, Levitt H, Mills R, Schwander T. An evaluation of three adaptive hearing aid selection strategies. *J. Acoust. Soc. Am.* 1987; 82:1967–1976. [PubMed: 3429734]
- Peterson GE, Lehiste I. Revised CNC lists for auditory tests,” *J. Speech Hear. Disord.* 1962; 27:62–70.
- Preminger JE, Neuman AC, Bakke MH, Walters D, Levitt H. An examination of the practicality of the simplex procedure. *Ear Hearing.* 2000; 21:177–193. [PubMed: 10890726]
- Reiss LA, Turner CW, Erenburg SR, Gantz BJ. Changes in pitch with a cochlear implant over time. *J. Assoc. Res. Otolaryngol.* 2007; 8:241–257. [PubMed: 17347777]
- Reiss LA, Perreau AE, Turner CW. Effects of lower frequency-to electrode allocations on speech and pitch perception with the Hybrid short-electrode cochlear implant. *Audiol. Neurotol.* 2012; 17:357–372.
- Rosen S, Faulkner A, Wilkinson L. Adaptation by normal listeners to upward spectral shifts of speech: implications for cochlear implants. *J. Acoust. Soc. Am.* 1999; 106:3629–3636. [PubMed: 10615701]
- Sagi E, Fu Q-J, Galvin JJ, Svirsky MA. A model of incomplete adaptation to a severely shifted frequency-to-electrode table mapping by cochlear implant users. *J. Assoc. Res. Otolaryngol.* 2010; 11:69–78. [PubMed: 19774412]
- Saunders GH, Cienkowski KM. A test to measure subjective and objective speech intelligibility. *J. Am. Acad. Audiol.* 2002; 13:38–49. [PubMed: 11833991]
- Stakhovskaya O, Sridhar D, Bonham BH, Leake PA. Frequency table for the human cochlear spiral ganglion: Implications for cochlear implants. *J. Assoc. Res. Otolaryngol.* 2007; 8:220–233. [PubMed: 17318276]
- Stelmachowicz PG, Lewis DE, Carney E. Preferred hearing-aid responses in simulated listening environments. *J. Speech Hear. Res.* 1994; 37:712–719. [PubMed: 8084201]
- Svirsky MA, Silviera A, Neuberger H, Teoh SW, Suarez H. Long-term auditory adaptation to a modified peripheral frequency table. *Acta. Otolaryngol.* 2004; 124:381–386. [PubMed: 15224858]
- Varenberg B, Govaerts PJ, De Ceulaer G, Daemers K, Schauwers K. Experiences of the use of FOX, an intelligent agent, for programming cochlear implant sound processors in new users. *Int. J. Audiol.* 2011; 50:50–58. [PubMed: 21091083]
- Wakefield GH, van den Honert C, Parkinson W, Lineaweaver S. Genetic algorithms for adaptive psychophysical procedures: recipient-directed design of speech-processor maps. *Ear Hear.* 2005; 26(4 Suppl):57S–72S. [PubMed: 16082268]

Zinn H. Stories Hollywood never tells. Alternative Tentacles. 2001 Audio CD.

Figure 1. GUI for self-selection of frequency tables

This figure depicts the graphical user interface utilized by our listeners for real-time selection of frequency tables. The light grey rectangle represents the total audible range of human hearing, while the dark gray rectangle depicts the frequency range of the analysis filters that are active. By pressing keys on a keyboard, listeners can change which analysis filters are active, and the grey rectangle will move accordingly. In this way listeners adjusted the speech signal with the intent of maximizing speech intelligibility.

Figure 2. Frequency tables

This figure depicts the frequency range of some important tables employed in Experiments 1 and 2, including the *right-information* (shaded rectangle on far left), *frequency-matched* (hashed rectangle on far right), *mean listener-selected* (thick gray rectangle) and *individual listener-selected tables* (thin gray rectangles). Dashed lines indicate the edges of the *frequency-matched* table.

Figure 3. Individual frequency-table selection, Experiment 1

The listener-selection table process is shown for two different starting points, the lowest-frequency point (gray rectangles) and the high-frequency point (black rectangles). For each starting point, the geometric means of the frequency tables traversed by this listener are connected by solid lines (lowest-frequency starting point) and short dashed lines (highest-frequency point). Thick dashed lines indicate the edges of the *frequency-matched* table.

Figure 4. Word- and phoneme-recognition scores

CNC word- and phoneme-recognition scores are indicated for the right-information (shaded bar), frequency-matched (hashed bar), and listener-selected table (open bar). Error bars indicate $\pm$ one standard error of the mean across listeners.

Figure 5. Individual frequency-table selection, Experiment 2

Same as Figure 2 except that data are shown only for the 10 listeners participating in Experiment 2.

Figure 6. Word- and phoneme-recognition scores, Experiment 2

Same as Figure 4 except that data are shown only for the 10 listeners participating in Experiment 2.

Figure 7. Frequency-table selections for each individual trial

Same as Figures 2 and 5, except that, for a given listener, the results of each individual trial are depicted.

Table 1

Canonical filter bank: a frequency table similar to that used with an 8-channel device

249 – 489	489 – 747	747 – 996	996 – 1463	1463 – 2150	2150 – 3159	3159 – 4642	4642 – 6820
-----------	-----------	-----------	------------	-------------	-------------	-------------	-------------

Table 2

Frequency range for 36 filter banks with a similar span: the range of the canonical filter bank from Table 1 is indicated with bold type

20 – 2680
20 – 2884
28 – 3102
42 – 3336
57 – 3586
73 – 3854
90 – 4142
108 – 4450
128 – 4780
149 – 5134
171 – 5513
195 – 5919
221 – 6354
249 – 6820
278 – 7320
310 – 7855
344 – 8429
381 – 9043
420 – 9702
462 – 10408
506 – 11164
554 – 11974
606 – 12842
661 – 13773
720 – 14770
783 – 15838
851 – 16982
924 – 18209
1002 – 19523
1085 – 20000
1175 – 20000
1271 – 20000
1373 – 20000
1484 – 20000
1601 – 20000
1728 – 20000

Table 3

Frequency range for filter banks with an approximate midpoint of 16.9 mm, and different spans: the range of the canonical filter bank from Table 1 is indicated with bold type

398 – 4972	342 – 5532	292 – 6154	249 – 6820	207 – 7610	170 – 8459	137 – 9400	107 – 10444
------------	------------	------------	-------------------	------------	------------	------------	-------------

Amount of time needed to complete the real-time self-selection procedure for each measurement made with each listener. A run was terminated when the listener felt she had identified the frequency table that was perceived to maximize speech intelligibility.

Table 4

Run #	L25	L26	L27	L28	L29	L30	L31	L32	L33	L34
Run 1	3:45	2:15	4:38	3:42	4:53	1:52	1:19	1:15	2:54	2:21
Run 2	1:40	1:45	2:06	1:19	3:34	1:02	2:19	0:31	1:46	1:59
Run 3	1:45	3:45	1:51	2:32	2:53	1:22	0:58	0:32	1:44	1:39
Run 4		3:30	2:26	2:04	3:28	1:18	1:48	0:36	0:48	1:54