Speech Processors for Auditory Prostheses

NIH Contract N01-DC-92100

QPR #2: Apr-May-June 1999

Robert V. Shannon
Qian-Jie Fu
Lendra Friesen
Monita Chatterjee
John Wygonski
John Galvin III
Fan-Gang Zeng *
Mark Robert
Xiaosong Wang

House Ear Institute
2100 W. Third St.
Los Angeles, CA 90057

*** University of Maryland**
College Park, MD 20742

30 July 1999

ABSTRACT

In this quarter we continued hardware and software development on research interfaces for the Clarion and Nucleus-24 cochlear implant systems. The Clarion research interface (CRI) is fully debugged and is now ready for experimental use. We have received hardware and software from Cochlear Corp in the last quarter that will allow the development of two types of interface to the Nucleus-24 device: a laboratory based software-only interface using the existing clinical interface, and a hardware interface that will provide a platform for development of wearable processors.

In this report we provide an update on a previously reported experiment on stimulation rate, and a brief update is presented on “holes in hearing” experiment in progress. We also report in detail on the results of two studies: a study of long-term learning, and a comparison of the Nucleus-22 and Clarion implant systems in noise as a function of the number of electrodes. In the learning study three Nucleus-22 implant patients were given an experimental processor that shifted the frequency assignments to electrodes down in frequency by one octave. Measurements were made weekly over a period of three months. With only a few exceptions, performance dropped substantially, and then recovered somewhat over the three month interval, with most of the improvement coming in the first few weeks. However, for most test materials the performance after three months of experience was considerably poorer than the initial performance level.

In the noise study performance was measured on a variety of speech materials in 10 Nucleus-22 subjects and 10 Clarion subjects as a function of the number of electrodes used in the processor. Performance with both implants increased as the number of electrodes was increased from 2 to 8, and little improvement in performance was observed in Nucleus-22 patients as the number of electrodes was further increased to 20. Performance with the two implant devices was not significantly different at all numbers of electrodes and at all noise levels, in spite of the considerable differences between the two systems in electrode design and speech processing strategies.

IMPLANT INTERFACE DEVELOPMENT

Clarion Research Interface

The Clarion Research Interface is now fully debugged and ready for use in experiments. Some additional applications programming is necessary to integrate this interface into the laboratory experimental programs. We anticipate that we will initiate experiments with this interface in the next quarter.

Nucleus-24 Research Interface

We have received hardware and software from Cochlear Corp in the last quarter that will allow the development of two types of interface to the Nucleus-24 device: a laboratory based software-only interface using the existing clinical interface, and a hardware interface that will provide a platform for development of wearable processors. The software interface will utilize the Clinicians' Programming Interface (CPI), which is the standard clinical interface for the Nucleus-24 device. The software interface will allow preprocessed stimuli to be presented to patients with the Nucleus-24 system in the laboratory. This will allow psychophysical experiments and speech processor manipulations in acute studies. Primary development of this software will take place in the next quarter.

The hardware interface will utilize a Motorola 56300 series DSP platform that will process speech in real time and code the signals in an appropriate fashion for transmission over the radio-frequency coils to the implanted receiver/stimulator. This interface will allow construction of wearable processors that will be capable of running novel experimental signal processing strategies to patients with the Nucleus-24 system. We anticipate that this hardware and software development will take place over the next two quarters.

OVERVIEW OF EXPERIMENTS IN PROGRESS

Stimulation Rate

In response to reviewer's comments on our previous study we ran an additional three implant subjects in the experiment to assess the effects of stimulation rate on speech recognition. In this experiment, which was included as a draft manuscript in the previous quarterly report, we measured speech recognition as a function of the stimulation rate in four-channel CIS processors. Six Nucleus-22 subjects participated in the study. Custom, four-channel CIS processors were constructed for each subject and stimuli were pre-processed off line for later presentation. Speech recognition was measured for stimulation rates of 50, 100, 150, 200, 300, 400 and 500 pps/electrode. The envelope filter was fixed at 40% of the stimulation rate. The results, presented in Figure 1, show an increase in performance with increasing stimulation rate up to 150 pps/electrode, and no change in performance from 150 to 500 pps/electrode.

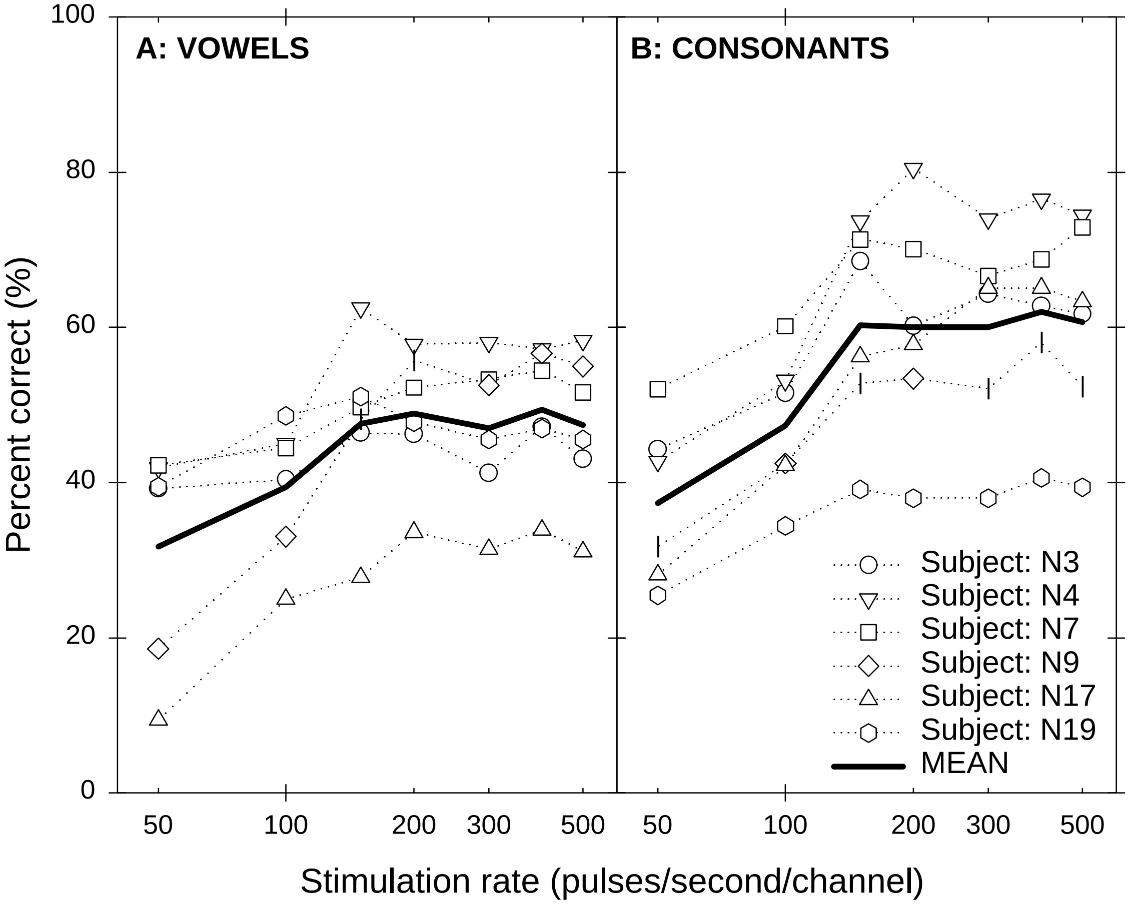

FIGURE 1: Vowel and consonant recognition as a function of stimulation rate for six subjects with 4-channel CIS processors.

To ensure that the limitation on performance was not the low-pass envelope filter the stimulation rate was fixed at 500 pps/electrode and the envelope filter was changed from 2 Hz to 160 Hz. Measures were also collected from five normal-hearing subjects for comparison. The normal-hearing subjects listened to a 4-noise-band processor with parameters that were identical to the 4-channel CIS processors except that the carrier in each channel was a band of noise rather than a 500 Hz pulse train. The results, presented in Figure 2, show no significant change in performance for low-pass filter settings above 20 Hz. These results indicate that temporal information in speech does not increase with increases in pulse rate above 150 Hz, at least up to the 500 Hz limit of this experiment. The data from the additional three implant listeners did not change the results or conclusions of the original study.

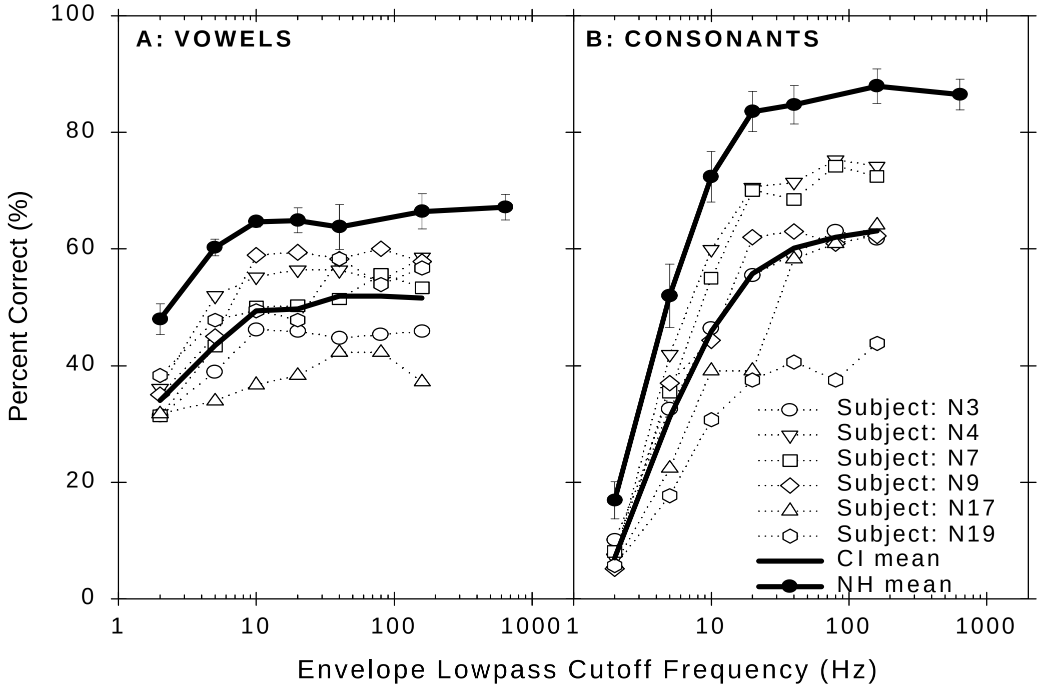

FIGURE 2: Vowel and consonant recognition for six implant listeners as a function of the low-pass envelope filter cutoff frequency in 4-channel CIS processors. Average results from 5 normal-hearing listeners with similar processing are presented for comparison.

Holes in Hearing

One of the important issues in electrical stimulation of the cochlea is the uniformity of nerve survival in the deaf cochlea. The density and uniformity of nerve survival may depend on the type of pathology that caused the deafness of the individual patient. If the nerve supply is uneven, then evenly spaced electrodes will not produce the intended pattern of excitation along the auditory nerve population. To assess the importance of the uniformity of nerve survival, we designed an experiment to place “holes” along the tonotopic dimension of the cochlea. We have now collected data from several patients with cochlear implants and will collect similar data in normal-hearing listeners for comparison.

“Holes” were created in Nucleus-22 implant patients by the following method. Two or more sequential electrodes were selected to define the hole. The hole was created in either the apical, middle, or basal region of the electrode array. Holes were created that were 2, 4, 6, or 8 electrodes in width. The electrodes selected for the hole were turned off. The full 20-band frequency analysis was still used, but the filter bands that would have normally been assigned to the electrodes in the hole were re-routed. Four conditions were assessed: the filter bands normally assigned to the electrodes in the hole were (1) reassigned to the electrode on the apical edge of the hole, (2) reassigned to the electrode at the basal edge of the hole, (3) evenly split between the electrodes at the apical and basal edge of the hole, or (4) dropped. Performance was measured in all conditions on 16 medial consonants, 12 medial vowels, and TIMIT sentences. Figure 3 presents preliminary results from one Nucleus-22 subject (N4) for these conditions.

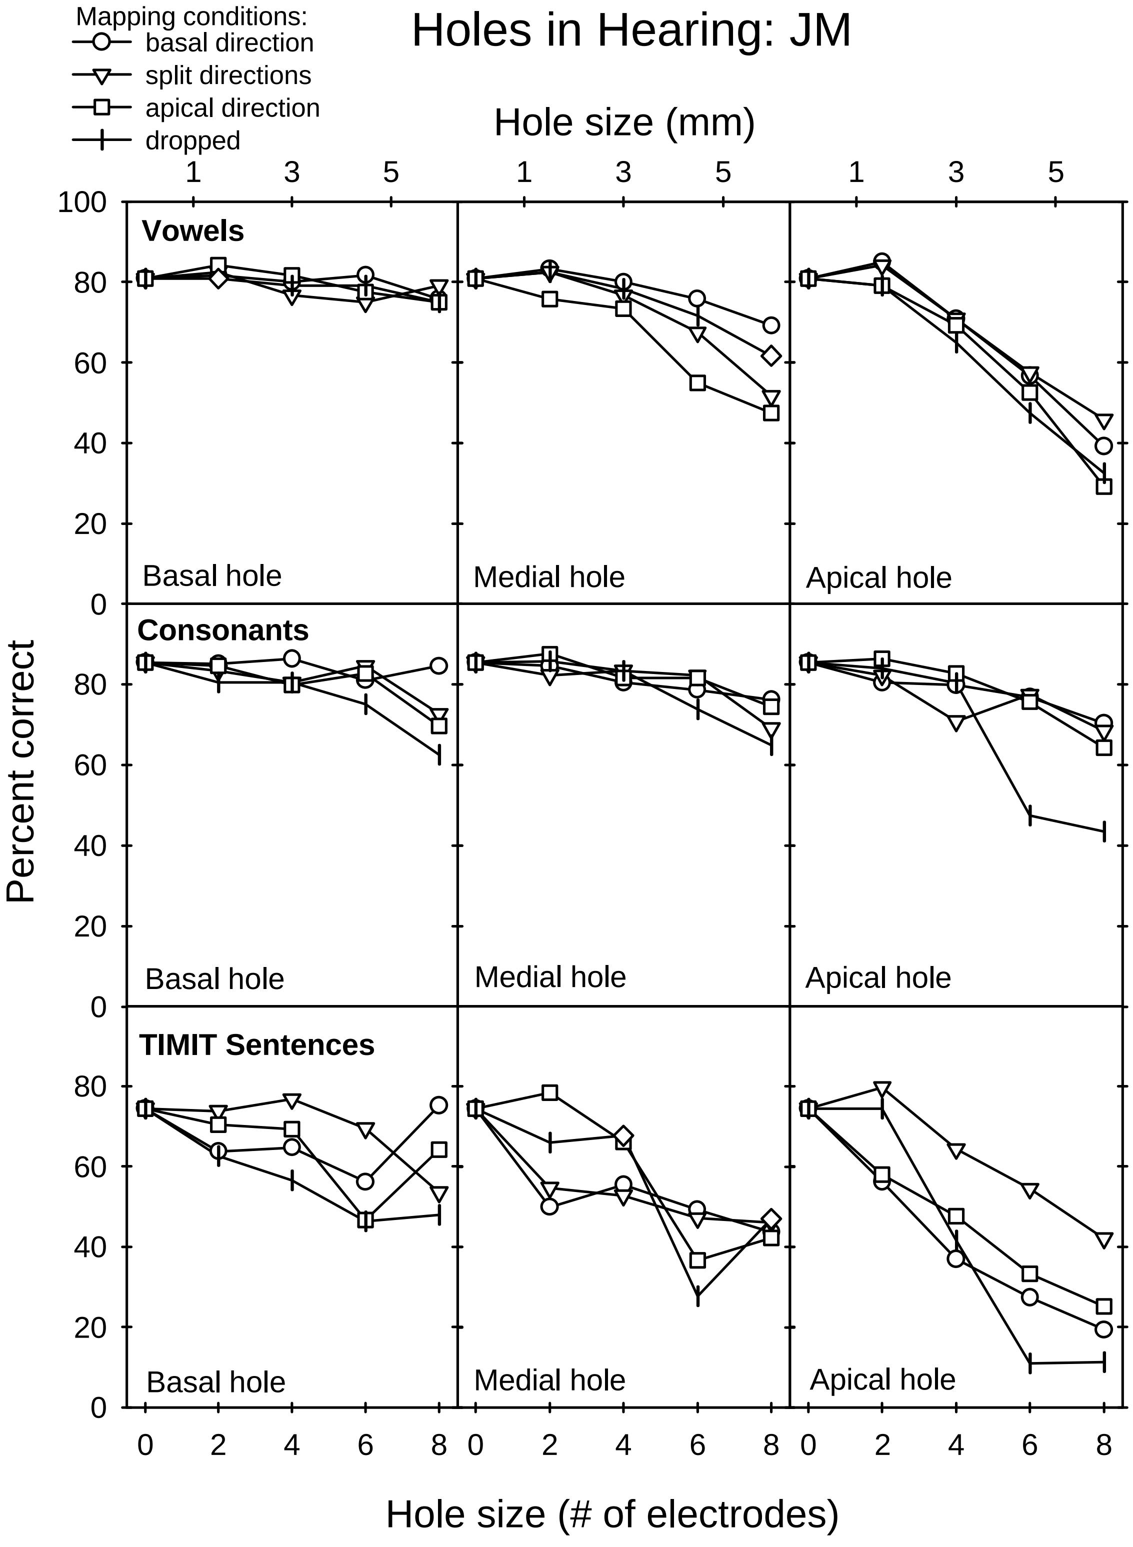

FIGURE 3: Recognition of consonants, vowels, and TIMIT sentences by Nucleus-22 subject N4 under conditions in which electrodes were removed from their map to create a “hole” in the pattern of stimulation.

These results indicate that:

1. an apical hole is more damaging to intelligibility than a basal hole
2. the disruption of intelligibility increases with the size of the hole
3. dropping the speech information was slightly more harmful than reassigning the information to the edges of the hole.

Points 1 and 2 are to be expected from a consideration of speech information and may even be quantitatively predictable from the articulation index, an analysis we will perform when the experiment is complete. However, based on our earlier results on the misrepresentation or warping of spectral information in speech (Shannon et al., 1998; Fu and Shannon, 1998) we anticipated that dropping the information from the hole would be less harmful than assigning it to the edges of the hole, which creates a warping of the spectral information around the hole. The fact that this warping was not worse than dropping the information, even for large holes, was a surprising result. As expected, the least disruption was observed when the filters were split to both sides of the hole.

Number of Channels in Noise

Several studies have shown that speech recognition improves with the number of spectral channels (e.g., Shannon et al., 1995, Fishman et al., 1997). However, these studies were done in quiet listening conditions. We are replicating the Fishman et al. study in quiet listening conditions as well as in four levels of added noise. Speech-shaped noise was added to all stimuli prior to processing to achieve signal-to-noise levels of +15, +10, +5, and 0 dB. The number of electrodes was varied from 2 to 8 for the Clarion device (2 to 7 for “enhanced bipolar” devices) and from 2 to 20 for the Nucleus-22 device. For 10 patients with the Nucleus-22 device the electrode reduction was accomplished in the same manner as in the Fishman et al. study, i.e., the output of multiple analysis filters were assigned to selected electrodes, while other electrodes were turned off. For the 9 subjects with the Clarion device the normal clinical fitting system was used to reduce the number of electrodes. In this case, the full speech spectral range was divided into the same number of segments as there were electrodes. So as the number of electrodes was reduced, the full spectral range was simply divided into fewer bands.

Preliminary results are presented in Figure 4 for recognition of vowels, consonants, phonemes in CNC words, and HINT sentences. In all cases performance increased as the number of electrodes increased from 2 to 7 or 8. However, there was little or no increase in recognition with the Nucleus device when the number of electrodes was increased from 7 to 10 to 20, consistent with the earlier results in quiet by Fishman et al. (1997). Results from normal-hearing (NH) listeners under similar conditions (from Fu et al., 1998) are presented for vowel and consonant recognition for comparison. Note that performance for NH listeners continued to increase as the number of spectral channels was increased, unlike the CI listeners whose performance reached asymptote at 7 or 8 electrodes. It is not clear at the present time why the performance of CI patients does not continue to increase with the number of electrodes.

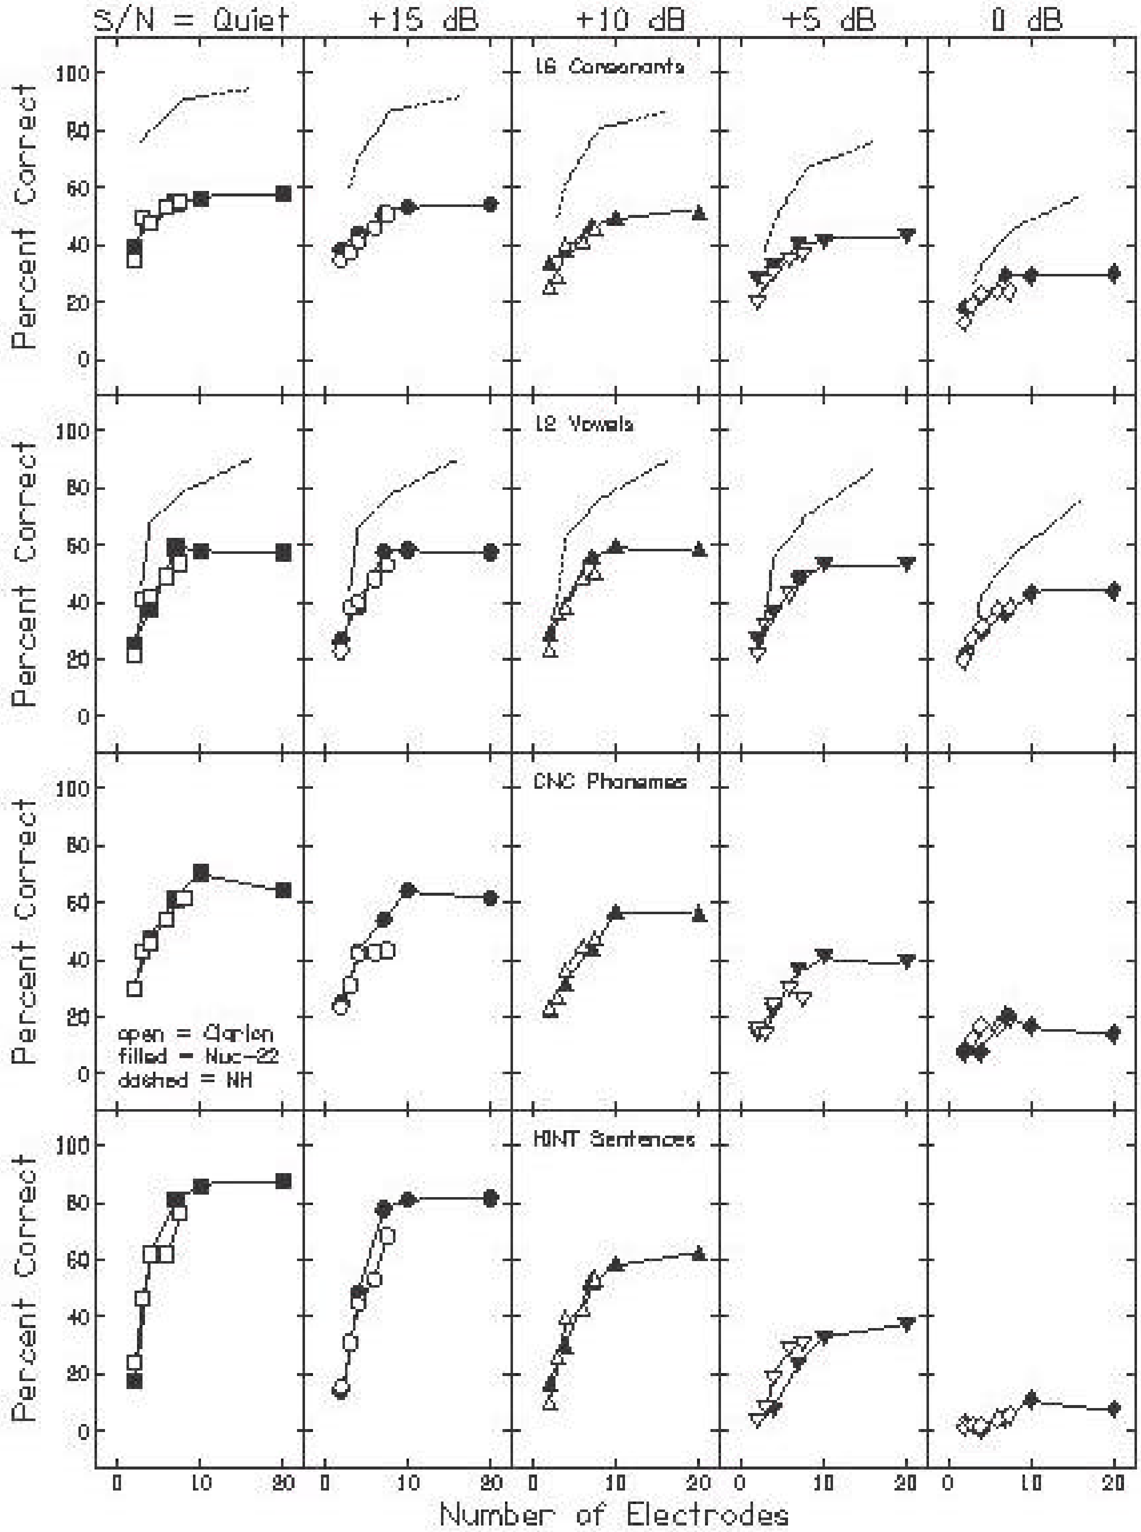

FIGURE 4: Recognition of consonants, vowels phonemes in CNC words and HINT sentences by 10 patients with the Nucleus-22 device with the SPEAK processing strategy and by 9 patients with the Clarion device (4 with CIS and 5 with SAS). Performance was measured as a function of the number of electrodes used in the processor and as a function of the signal-to-noise level.

One particularly interesting result in Figure 4 is that there was no significant difference in performance between patients with the Nucleus and Clarion devices either as a function of noise level or as a function of the number of electrodes. In spite of the differences in electrode design and placement, and differences in the speech processing strategies (SPEAK vs CIS vs SAS) average performance was similar for the same number of electrodes. Even the range of performance across subjects was similar for the two devices in this relatively small sample. This result suggests that the number of electrodes is the key factor in determining overall speech recognition performance.

Long-term Learning Effects

One of the most important issues in fitting speech processors of cochlear implants is the effect of learning. Most experiments to assess the effects of speech processor parameters are done in “acute” laboratory studies with little or no time for the subject to adapt to the new strategy. However, it is well documented that subjects’ performance improves over time with experience. Subjects generally show an initial rapid improvement with a new speech processor and in most cases reach asymptote after a few months (e.g., Tyler et al, 1997; Pelizzone et al., 1999; Lawson et al., 1999). However, the time constants of these learning effects vary widely. In the Tyler et al. study, 24 Nucleus-22 subjects and 25 Ineraid subjects showed most of the increase by one month of use, suggesting an initial time constant of a few weeks. Pelizzone et al., and Lawson et al. measured performance in Ineraid patients when they were switched over from their normal 4 channel compressed analog processor to a 6-channel CIS processor. Pelizzone et al. observed a time constant of approximately 50-60 days, while Lawson et al observed time constants of 100-200 days. It is not clear why these studies observed such a large difference in the time constants of learning under similar processor manipulations.

Another issue is the possibility that subjects can adapt to changes in the frequency-to-electrode assignments. Rosen et al. (1998) recently showed considerable learning in normal-hearing listeners when they were trained for only a few 20-minute sessions listening to a 4-band noise processor in which the frequency map had been shifted by 6 mm relative to the normal tonotopic map. Subject performance with these shifted maps initially dropped dramatically, but then recovered partially after only 2 or 3, 20-minute training sessions. This result was encouraging for fitting implant speech processors, in that it suggested that errors in the frequency-to-electrode mapping could be quickly corrected by learning. To test this hypothesis we shifted the frequency-to-electrode map in three Nucleus-22 subjects who had used the SPEAK strategy for more than three years. The frequency map was shifted by one octave from the map they used in their normal clinical system and they wore the shifted map for a period of three months. Performance on a variety of speech measures was assessed weekly.

I. METHOD

A. Subjects

Three post-lingually deafened adults using the Nucleus-22 implant device participated in this study. All had at least four years experience using the SPEAK speech processing strategy and all were native speakers of American English. All implant subjects had 20 active electrodes available for use. Two subjects (N4 and N7) used frequency allocation table 9 (150-10,823 Hz) in their clinically assigned processor and one subject (N3) used frequency allocation table 7 (120 Hz - 8,658 Hz). All participants had extensive experience in speech recognition experiments.

B. Experimental Speech Processor Mapping

In the Nucleus-22 implant device, the MAP consists of the psychophysical parameters for each individual patient stored in the memory of the speech processor; it is used to accommodate perceptual differences between patients and compensates for differences between electrodes within the same patient. The MAP determines how acoustic information is transformed into patterns of electric stimulation. A primary function of the MAP is to translate the spectral information found in the incoming acoustic signal into the appropriate electrodes to be stimulated. During the clinical fitting, frequency ranges are assigned to the electrodes; usually the default frequency allocation Table 9 (frequency range: 150-10,824 Hz) is assigned to patients. The experimental speech processor was implemented in each subject’s Spectra 22 speech processor (i.e., hardware device) by assigning frequency allocation Table 1 (frequency range: 75 –5411 Hz), thereby shifting the filterbank analysis bands one octave lower. 20-electrode processors were used; threshold and maximum comfort levels (T- and C-levels) for each electrode were re-measured for each subject, as well as the optimal microphone sensitivity. Subjects volunteered to wear the experimental processor as their “everyday” map, allowing them to experience the new processor in a variety of settings outside the laboratory environment.

C. Test Materials

Speech intelligibility was assessed using five sets of test materials, including two closed-set identification tasks and three open-set recognition tasks. The two closed-set identification tasks included multi-talker vowel and consonant identification. Vowel recognition was assessed in a 12-alternative identification paradigm, including 10 monophthongs (/i I E Q u U A Ã • ÎÕ /) and 2 diphthongs (/o e/), presented in a /h/vowel/d/ context (heed, hawed, head, who’d, hid, hood, hud, had, heard, hoed, hod, hayed). The tokens for these closed-set tests were digitized natural productions from 5 men, 5 women, and 5 children, drawn from speech samples collected by Hillenbrand et al. (1995). Consonant recognition was assessed in a 16-alternative identification paradigm for the consonants /b d g p t k l m n f s ʃ v z θ dZ/, presented in a /a/-consonant-/a/ context. Three speakers (1 male, 2 female) produced two exemplars of each consonant.

Three open-set recognition tasks included one monosyllable word test and two sentence tests. The CNC word recognition test was used to evaluate open-set phoneme and word recognition. The material contains 10 lists of 50 monosyllabic words for a total of 150 phonemes. Listener responses were scored for phonemes correctly identified. Recognition of words in sentences was measured using the Hearing in Noise Test (HINT) sentences and the DARPA TIMIT acoustic-phonetic continuous speech corpus. The sentences were of easy-to-moderate difficulty for HINT and of moderate-to-hard difficulty for TIMIT. Two different sets of 10 HINT sentences and three different sets of 20 TIMIT sentences were used for each testing session. Note that no TIMIT sentence recognition was conducted for subject N4 because the test materials were not available at that time.

D. Procedures

Tests of speech intelligibility with the experimental speech processor were administered at the time of fitting, then at 1, 2, and 4 days post-fitting, and then weekly thereafter, until the end of the 3-month test period. Speech intelligibility tests using the clinically assigned MAP were administered immediately before fitting with the experimental MAP; post-experiment control tests were conducted immediately following the re-installation of the clinically assigned MAP.

For vowel and consonant identification, each test block included 180 tokens (12 vowels*15 talkers) for vowel identification and 192 tokens (16 consonants * 3 talkers * 2 exemplars * 2 repeats) for consonant identification. For each trial, a stimulus token was chosen randomly, without replacement, and presented to the subject. Following presentation of each token, the subject responded by pressing one of 12 buttons in the vowel test or one of 16 buttons in the consonant test, each marked with one of the possible responses. The response buttons were labeled in a /h/-vowel-/d/ context for the vowel recognition task and /a/-consonant-/a/ context for the consonant recognition task. For CNC monosyllabic word recognition, a stimulus token was chosen randomly from a 50-word list for each trial. Following presentation of each stimulus token, the subject responded by repeating the word as accurately as possible; the experimenter tabulated correctly identified phonemes and words. For HINT sentence recognition, a list was chosen pseudo-randomly from among 26 lists, and sentences were chosen randomly, without replacement, from the 10 sentences within that list. For TIMIT sentence recognition, a list was chosen pseudo-randomly from among 50 lists, and sentences were again chosen randomly, without replacement, from the 20 sentences within that list. The subject responded by repeating the sentence as accurately as possible; the experimenter tabulated correctly identified words and sentences. Individual data points represent two runs for both vowel and consonant identification, one run for CNC word recognition, two runs for HINT sentence recognition, and three runs for TIMIT sentence recognition. No feedback was provided for all measures, and subjects were instructed to guess if they were not sure, although they were cautioned not to provide the same response for each guess. To avoid any headset microphone effects, all speech tokens were presented to the subject via

Cochlear’s Audio Input Selector (AIS) device; the gain of the AIS was set to maximum. All speech stimuli were presented at a comfortable loudness level and all stimuli were equated in terms of rms level.

II. RESULTS

Figure 5 shows the mean and individual scores for vowel, consonant, and CNC phoneme recognition as a function of time. Note that scores are plotted as the difference between post-fitting and pre-fitting scores; pre-fitting scores were measured just before implementation of the experimental processor. For multitalker vowel identification, scores dropped dramatically for all subjects immediately post-fitting with the experimental processor. The performance drop ranged from 20% to 35%, with a mean of 32%. The mean deficit gradually reduced to 20% during the first week, after which no significant improvement was observed. For multitalker consonant identification, a mean drop of only 15% was observed immediately post-fitting. The mean deficit reduced to about 10% after 20 days and 5% after 50 days. For consonant identification, similar performance to baseline levels was achieved after 3 months (90 days). For CNC phonemes, a mean drop of 25% was observed immediately post-fitting. The mean deficit gradually reduced to about 5% after 2 months. In contrast to the learning pattern in vowels and consonant recognition, CNC phoneme recognition showed large variation across subjects. Two of the three subjects showed asymptotic performance similar to the baseline performance after 60 days of exposure to the new processor. However, one subject (N4) showed asymptotic performance after only 7 days and remained 30% lower than the baseline performance.

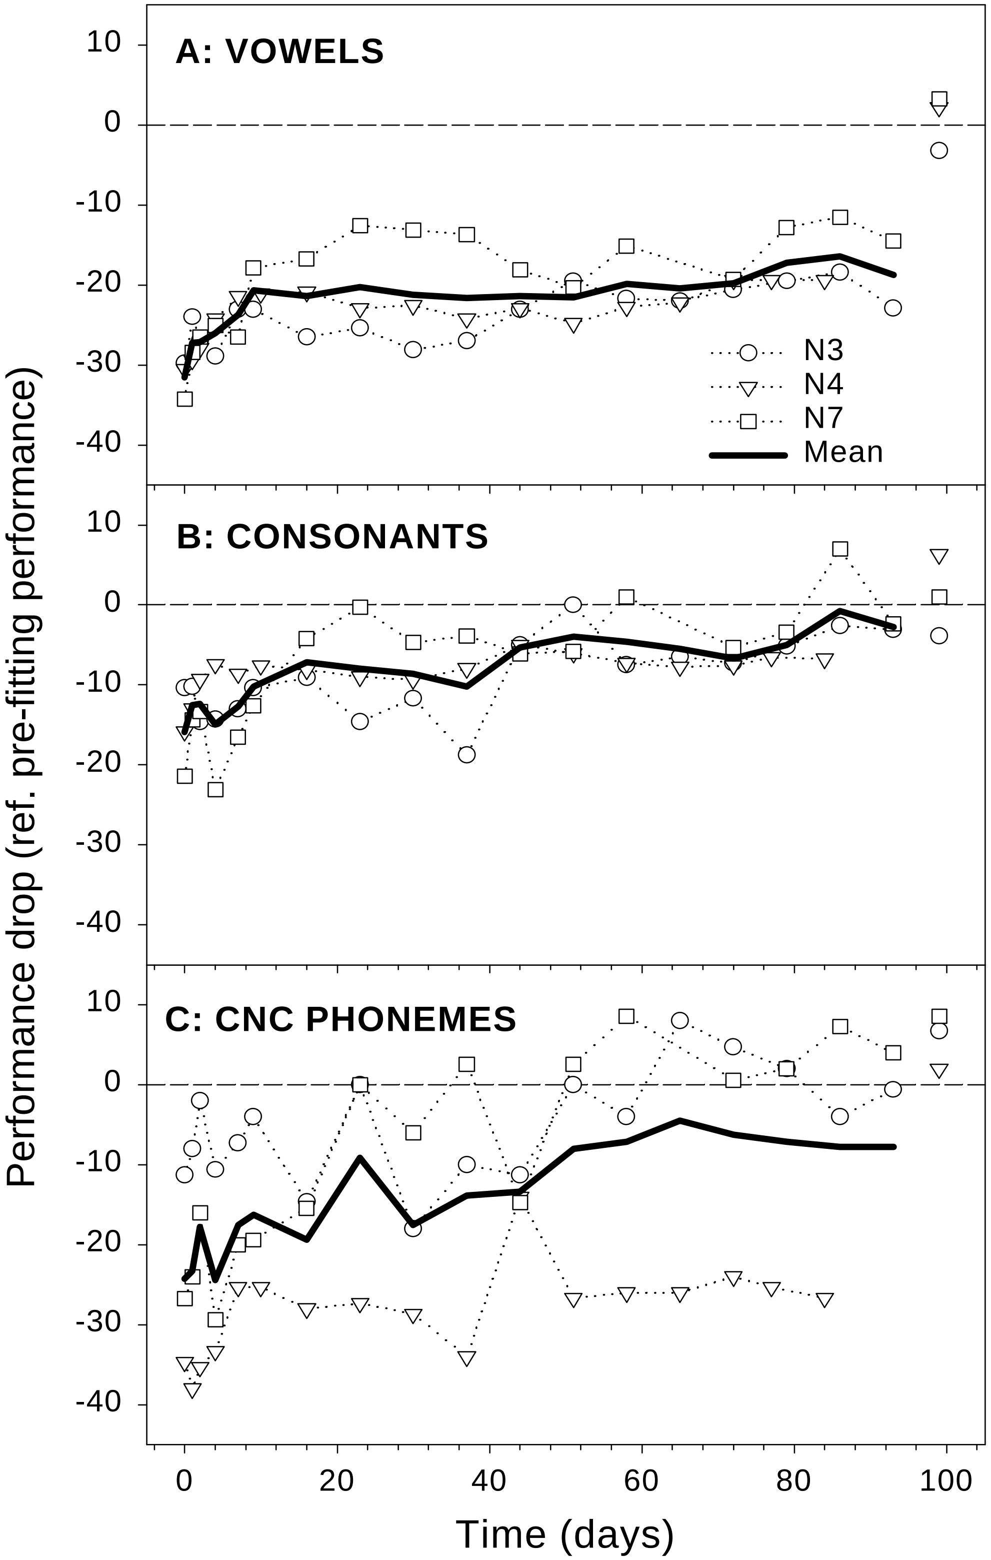

FIGURE 5: Vowel, consonant, and CNC phoneme recognition scores over time. Data are presented as the difference in scores between baseline performance with the clinical processor and the experimental processor.

Figure 6 shows mean and individual scores of HINT and TIMIT sentence recognition as a function of time. For HINT sentence recognition, a mean drop of 35% was observed immediately following implementation of the new speech processor. The deficit reduced to 20% after 1 week of exposure to the new processor and to 5% two months later, after which no improvement was observed. For TIMIT sentence recognition, a mean drop of 40% was observed immediately post-fitting. Note that only two subjects' data are available (no TIMIT recognition scores were measured in subject N4 because the test materials were not available at that time). The mean deficit for TIMIT sentences reduced to 20% after 2 months of exposure to the new speech processor, after which no improvement was observed.

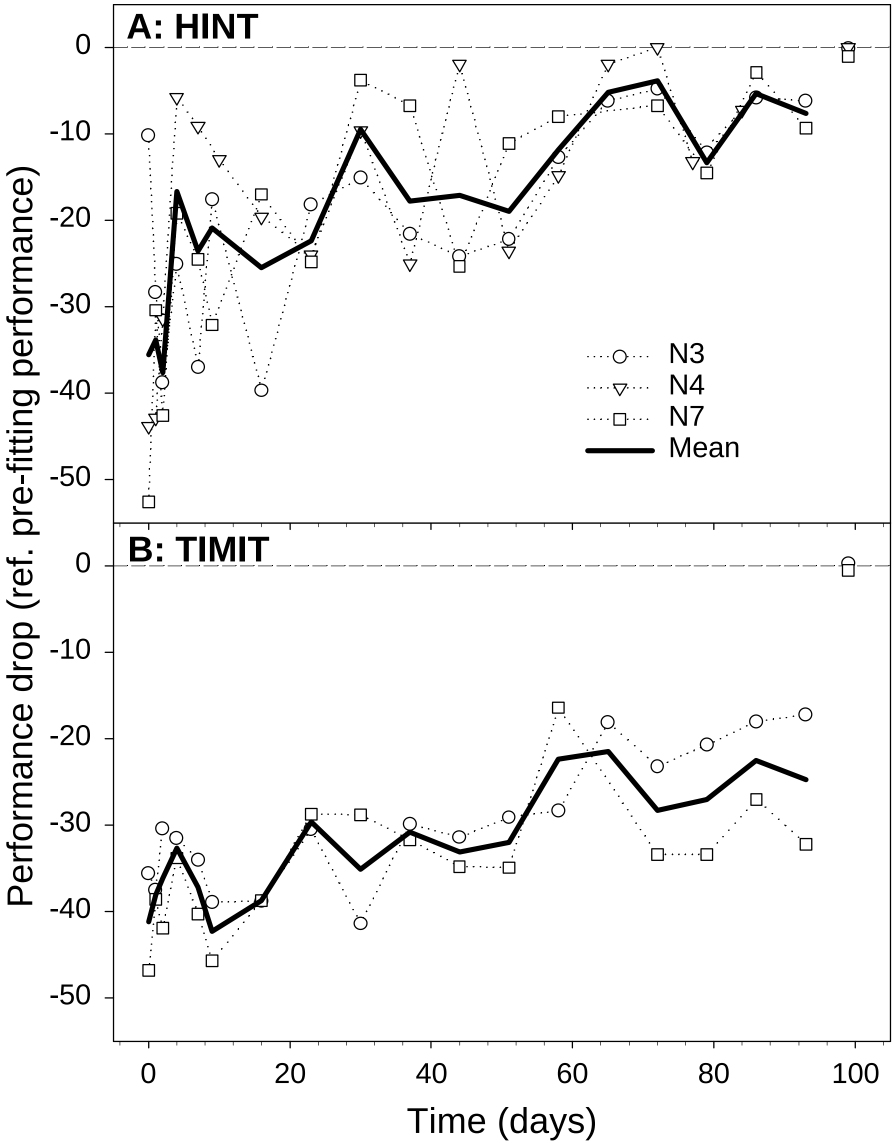

FIGURE 6: Sentence recognition over time with HINT and TIMIT sentences.

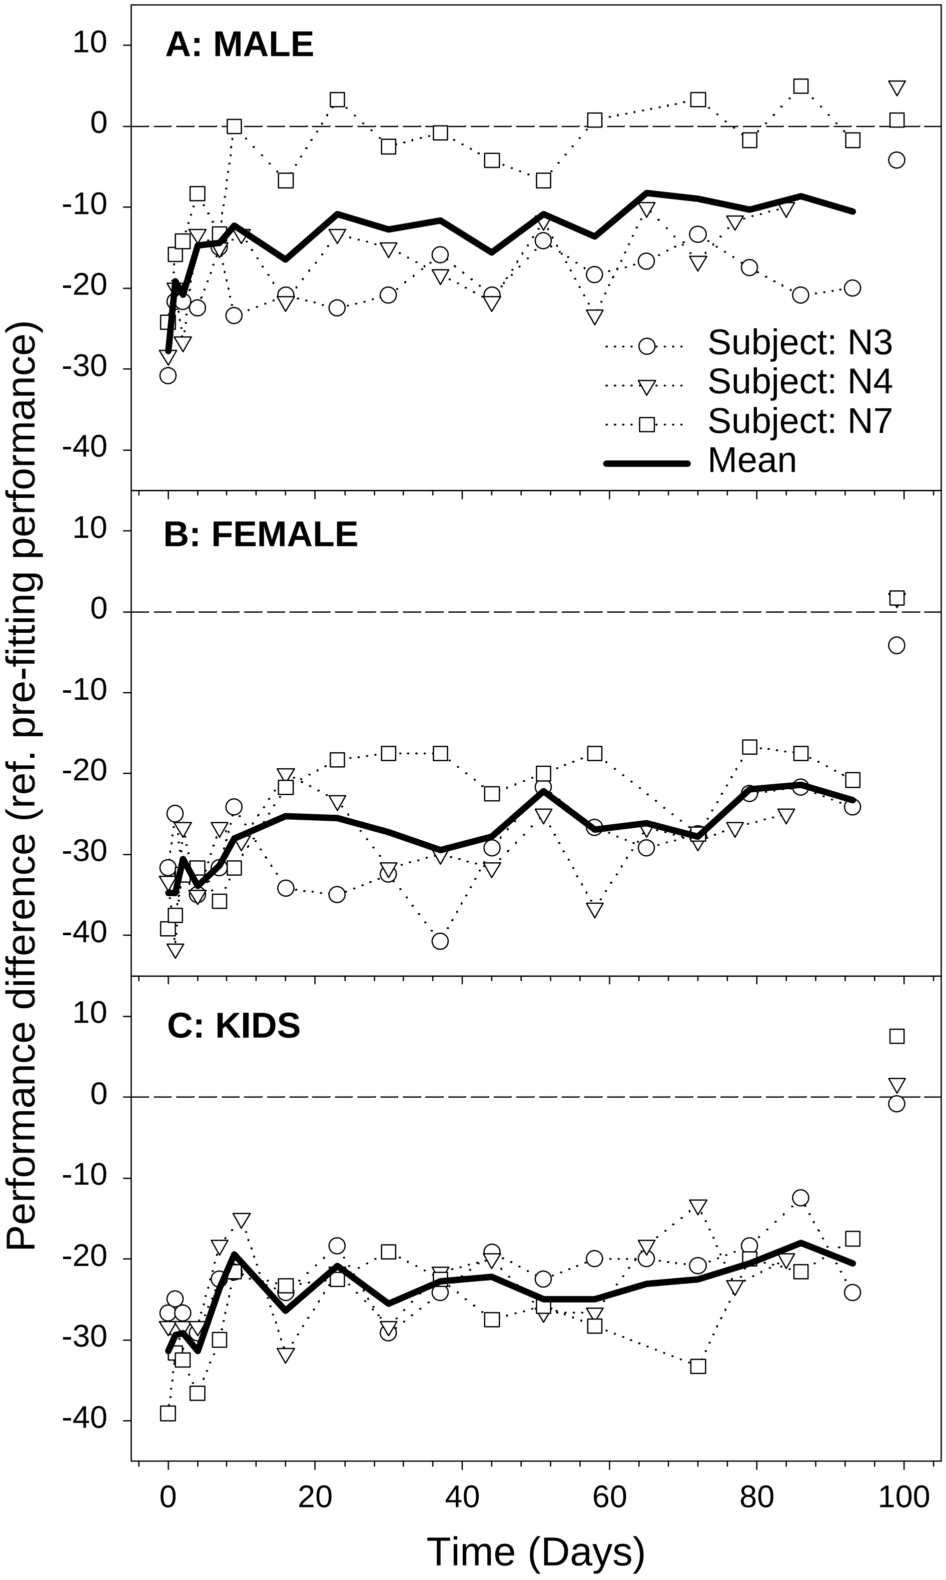

FIGURE 7: Vowel recognition scores over time for male, female and children talkers.

Figure 7 shows recognition scores of vowels produced by male, female, and child talkers as a function of time. For male talkers, mean performance dropped 28% immediately post-fitting. The mean deficit quickly reduced to 15% after only 3 days of exposure to the experimental processor, then gradually reduced to 12% after 3 months of exposure. Note that subject N7 reached an asymptotic performance level similar to his pre-fitting score. For female talkers, mean performance dropped 35% immediately post-fitting, and vowel scores remained 22% lower than the pre-fitting performance, even after 3 months of exposure. Similarly, the mean drop in vowel scores for child talkers reduced from 30% immediately post-fitting to 20% after 7 days of exposure, after which no further improvement was observed.

Figure 8 presents the consonant information received on the production-based categories of voicing, manner, and place as a function of time. For voicing information, the mean performance drop was 10% immediately post-fitting; however, after 2 months of exposure, the performance of the experimental processor was comparable to that of the normal clinical processor. Here also, large inter-subject variation was observed. Subject N7 actually demonstrated similar performance in the voicing category when using frequency Table 1 and Table 9, while subject N3 remained slightly lower with Table 1. A significantly different recognition pattern was observed in the reception of manner cues. The mean performance was 5% higher than the baseline performance immediately post-fitting; after 3 months, mean scores improved to 15% higher than the baseline performance. However, the post-fitting control performance regarding manner cues was also 15% higher than the pre-fitting performance, suggesting that some task-based learning might have occurred. For reception of consonant place cues, mean scores were 25% lower immediately post-fitting. The performance deficit gradually reduced to 8% after 3 months of exposure to the new processor.

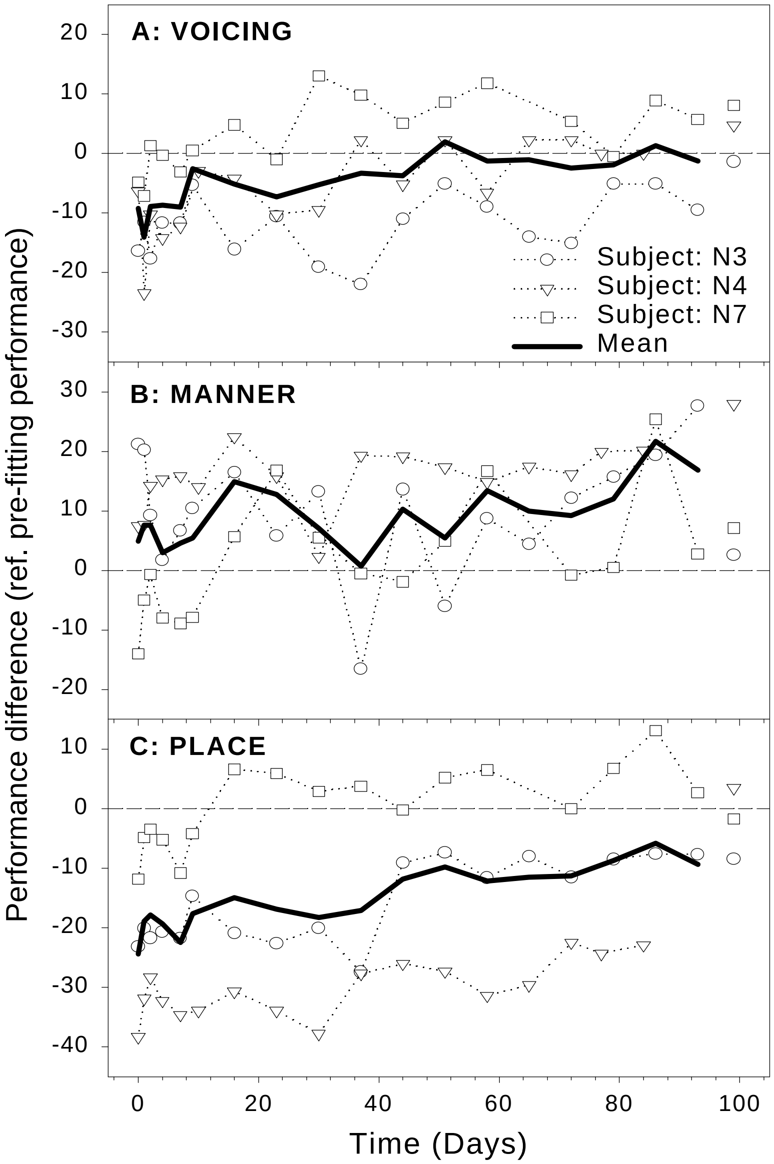

FIGURE 8: Consonant information received on the production-based categories of voicing, manner and place.

Discussion

The results show that implant listeners can partially adapt to a shifted frequency-to-electrode map, but that the adaptation is not complete after three months of daily listening. There is an initial rapid adjustment that occurs with a time constant of less than 10 days and, for some test materials, a longer adaptation with a time constant of 30-50 days. The initial rapid learning may be similar to that observed by Rosen et al. (1998), and may correspond to a “novelty factor”. Following an initial drop in performance the listeners may simply adjust to the quality of the shifted stimuli and may be able to “get past” the novelty of the new sound and use some cues which were available, but not used because their attention was diverted by the altered sound quality. Once past this initial adjustment period, performance either asymptotes, or shows slow improvement over the remaining 3 month trial, depending on the test materials. This latter process may be similar to the longer adjustment time observed by Lawson et al. (1999).

This result is not encouraging in terms of the implant patient's ability to adapt to a misadjusted frequency map. Researcher's have been optimistic about the possibility of implant listener's adaptability to frequency-shifted maps based on short-term experiments with normal-hearing listeners (Rosen et al., 1998) and based on historical visual experiments in the 1960s with orientation-shifting goggles, which indicated that only a short time period (possibly a few weeks) was required to adapt to an alteration in the sensory image. The results from the present experiment, however, show that for these three subjects even three months of daily use was not sufficient to adapt to a frequency map that was shifted by only one octave (4.5 mm). Such shifts could easily occur in cochlear implants due to the individual differences in insertion depth of the intracochlear electrodes. We emphasize that all other aspects of the implant processing were similar between the original map and the shifted map – the only difference was the shift in the frequency-to-electrode mapping. This result, together with earlier results on the effects of frequency shifts and frequency warping, suggest that these parameters are critical factors in fitting speech processors for cochlear implants. An incorrect assignment of frequency information to electrodes could cause a long-term or even permanent decrement in speech recognition performance.

Publications and Presentations in this Quarter

Fu, Q.-J. and Shannon, R.V. (1999). Phoneme recognition as a function of signal-to-noise ratio under nonlinear amplitude mapping by cochlear implant users, Journal of the Acoustical Society of America (Acoustic Research Letters Online), 106(2), L18-L23.

Shannon, R.V. (1999). Cochlear Implants and Auditory Brainstem Implants, Hearing Science for Practicing Otologists, Beijing, China, May 17-22. (Invited Speaker)

Zeng, F.-G. and Shannon, R.V. (1999). Psychophysical laws revealed by electric hearing, NeuroReport, 10(9), 1-5.

Plans for the Next Quarter

In the next quarter we will continue hardware and software development on the interface for the Nucleus-24 device. The Clarion research interface will be integrated into laboratory experimental software and experiments will be initiated.

Fan-Gang Zeng will visit our laboratory during the week of August 2-6 to complete experiments on amplitude manipulations in Nucleus-22 and Clarion cochlear implant subjects.

In August 1999 we will prepare four invited talks and five poster presentations at the 1999 Conference on Implantable Auditory Prostheses, August 30-Sept 3 at Asilomar Conference Center, Pacific Grove, California.

The data presented in this report on long-term learning and on the number of channels in noise will be further analyzed. One point of particular interest is that the implant patients with the highest level of performance demonstrate phoneme recognition at levels equivalent to that of normal-hearing listeners with the same number of spectral channels. This implies that these patients are optimally using all spectral channels. However, other implant patients do not reach this level of performance. The data from excellent, medium, and poor implant users will be analyzed to see which speech features are not being received properly in the poorer performing implant users and which phonetic confusions are causing the difference in performance.

We are preparing several manuscripts that we anticipate submitting for publication in the next quarter:

Friesen, L.M., Shannon, R.V., and Slattery, W.H. Effects of electrode location on speech recognition in Nucleus-22 cochlear implant listeners, J. Amer. Acad. Audiol.

Fu, Q.-J. and Shannon, R.V. Effect of acoustic dynamic range on phoneme recognition in cochlear implant listeners, Journal of the Acoustical Society of America (Acoustic Research Letters Online).

Fu, Q.-J, and Zeng, F.-G. Identification of temporal envelope cues in Chinese tone recognition, International Journal of Language and Communication Disorders.

Fu, Q.-J. and Shannon, R.V. Effects of dynamic range and amplitude mapping on phoneme recognition in Nucleus-22 cochlear implant users, Ear & Hearing.

Fu, Q.-J. and Shannon, R.V. Performance over time of Nucleus-22 cochlear implant listeners wearing speech processor with a shifted frequency-to-electrode assignment, J. Acoust. Soc. Amer.

REFERENCES

Fishman K., Shannon R.V. and Slattery W.H. (1997). Speech recognition as a function of the number of electrodes used in the SPEAK cochlear implant speech processor. **Journal of Speech and Hearing Research**, 40, 1201-1215.

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(3), 1889-1900.

Fu, Q.-J., Shannon, R.V., and Wang, X. (1998). Effects of noise and spectral resolution on vowel and consonant recognition: Acoustic and electric hearing, **Journal of the Acoustical Society of America**, 104(6), 3586-3596.

Hillenbrand, J., Getty, L., Clark, M., and Wheeler, K. (1995). Acoustic characteristics of American English vowels, **J. Acoust. Soc. Am.,** 97, 3099-3111.

Lawson, D., Wilson, B., and Zerbi, M. (1999). Speech processors for auditory prostheses, Second Quarterly Progress Report, NIDCD Contract N01-DC-8-2105, **Neural Prosthesis Program**, Bethesda, MD.

McDermott, H.J., McKay, C.M., and Vandali, A.E. (1992). A new portable sound processor for the University of Melbourne/Nucleus Limited Multichannel cochlear implant, **J. Acoust. Soc. Amer.**, 91, 3367-3371.

Pellizone, M., Cosendai, G., and Tinembart, J. (1999). Within-patient longitudinal speech reception measures with continuous interleaved sampling processors for Ineraid implanted subjects, **Ear & Hearing**, 20(3), 228-237.

Rosen, S., Faulkner, A., and Wilkinson, L. (1998). Perceptual adaptation by normal listeners to upward shifts of spectral information in speech and its relevance for users of cochlear implants, in **Speech, Hearing, and Language: Work in Progress**, Phonetics and Linguistics, University College London, Vol. 10.

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

Shannon, R.V., Zeng, F.-G., and Wygonski, J. (1998). Speech recognition with altered spectral distribution of envelope cues, **J. Acoust. Soc. Amer.**, 104(4), 2467-2476.

Tyler, R.S., Parkinson, A.J., Woodworth, G.G., Lowder, M.W., and Gantz, B.J. (1997). Performance over time of adult patients using the Ineraid or Nucleus cochlear implant, **J. Acoust. Soc. Amer.**, 102, 508-522.