

Effects of Categorization and Discrimination Training on Auditory Perceptual Space

Frank H. Guenther^{a,b}, Fatima T. Husain^a, Michael A. Cohen^a, and
Barbara G. Shinn-Cunningham^a

Journal of the Acoustical Society of America (1999) Vol. 106, pp. 2900-2912

^aDepartment of Cognitive and Neural Systems
Boston University
677 Beacon Street
Boston, MA, 02215
Fax Number: (617) 353-7755
email: guenther@cns.bu.edu

^bResearch Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

Psychophysical phenomena such as categorical perception and the perceptual magnet effect indicate that our auditory perceptual spaces are warped for some stimuli. This paper investigates the effects of two different kinds of training on auditory perceptual space. It is first shown that categorization training using non-speech stimuli, in which subjects learn to identify stimuli within a particular frequency range as members of the same category, can lead to a decrease in sensitivity to stimuli in that category. This phenomenon is an example of acquired similarity and apparently has not been previously demonstrated for a category-relevant dimension. Discrimination training with the same set of stimuli was shown to have the opposite effect: subjects became more sensitive to differences in the stimuli presented during training. Further experiments investigated some of the conditions that are necessary to generate the acquired similarity found in the first experiment. The results of these experiments are used to evaluate two neural network models of the perceptual magnet effect. These models, in combination with our experimental results, are used to generate an experimentally testable prediction concerning changes in the brain's auditory maps under different training conditions.

PACS numbers: 43.71.An (models and theories of speech perception), 43.66.Ba (models and theories of auditory processes).

Frank Guenther is supported in part by the Alfred P. Sloan Foundation and the National Institute on Deafness and other Communication Disorders (NIDCD R29 02852). Fatima Husain is supported by NIDCD R29 02852. Barbara Shinn-Cunningham is supported in part by AFOSR grant F49620-98-1-0108. The authors would like to thank Harlan Lane, David Pisoni, Jacob Beck, James Beale, Peter Chiu, and Kenneth Stevens for helpful discussions regarding this work.

1. Introduction

It is well-known that our perceptual spaces for some auditory stimuli, such as phonemes, are warped. That is, the perceptual distance between two stimuli, as evidenced by a subject's ability to discriminate them, is not always a straightforward function of their distance as measured along physical dimensions such as frequency or time¹. English stop consonants, for example, have long been known to exhibit categorical perception (see Jusczyk, 1986; Liberman, 1996; Lieberman and Blumstein, 1988; and Repp, 1984 for reviews). For example, if subjects are presented with synthetic speech stimuli created by varying the second formant transition in small steps through a range corresponding to the phonemes /b/, /d/, and /g/, they show very poor discriminability when two stimuli both fall within one of the categories and very good discriminability for stimuli that straddle category boundaries, even though the stimulus pairs in these two cases are equidistant in frequency space (Liberman, Harris, Hoffman, and Griffith, 1957; Eimas, 1963). Other experiments have shown similar categorical effects for voice onset time (VOT) distinctions between /d/ and /t/ (Liberman, Harris, Kinney, and Lane, 1961) and between /b/ and /p/ (Liberman, Harris, Eimas, Lister, and Bastian, 1961). Similar effects have also been reported for a variety of non-speech stimuli, including melodic musical intervals (Burns and Ward, 1978), simple visual shapes (Lane, 1965; Goldstone, 1994), and morphed faces along an intriguing "John F. Kennedy to Bill Clinton" continuum (Beale and Keil, 1995).

Researchers have also shown, relatively recently, that the perceptual space for some synthetic vowels and semivowels appears to be warped (e.g., Aaltonen, Eerola, Hellström, Uusipaikka, and Lang, 1997; Iverson, Diesch, Siebert, and Kuhl, 1994; Iverson and Kuhl, 1994, 1995; Kuhl, 1991, 1995; Kuhl, Williams, Lacerda, Stevens, and Lindblom, 1992; Sussman and Lauckner-Morano, 1995). Kuhl (1991) referred to this warping as a "perceptual magnet effect," thus distinguishing it from categorical perception. Roughly speaking, the effect is characterized by a warping of perceptual space such that acoustic patterns near phonemic category prototypes are perceived as closer together than equally spaced acoustic patterns that are further away from phonemic category prototypes. According to the Kuhl et al. account, the magnet effect differs from categorical perception in that it is characterized by differences in discriminability for prototypical vs. nonprototypical stimuli that fall within the *same* phonemic category. Specifically, better discrimination is found near non-prototypical members of a category than near prototypical members. However, other researchers have claimed that categorical perception and the perceptual magnet effect are essentially the same. For example, Lotto, Kluender, and Holt (1998) conclude from their study that the magnet effect "may be nothing more than further demonstration that general discriminability is greater for cross-category stimulus pairs than for within-category pairs" (p. 3648). By either account, though, the perceptual space for vowels and semivowels appears to be warped, although apparently not as dramatically as for consonants.

It is very likely that some of the warping of auditory space is "built in" to the auditory nervous system. Evidence for this comes from studies of auditory perception in animals and newborn infants. For example, the discriminability by chinchillas of changes in VOT for stimuli varying between [d α] and [t α] is nonuniform and peaks at a VOT of about 30ms, which is near the voiced/voiceless boundary in English (Kuhl and Miller, 1975, 1978; Kuhl, 1981). A similar result was also reported for macaque monkeys (Kuhl and Padden, 1982). Increased discriminability was also found at the /b/-/d/ and /d/-/g/ phonetic boundaries of a continuum of F2 transition onset frequencies in the macaque monkey (Kuhl and Padden, 1983). Eimas, Siqueland, Jusczyk, and Vigorito (1971) showed that human infants 1-4 months old produced evidence of categorical perception for the voiced/voiceless distinction, further suggesting that this effect is a consequence of auditory mechanisms that are present at birth.

1.1. Experience-based warping of auditory space

Other aspects of the warping of auditory space appear to arise from learning, rather than from built-in properties of the auditory system. Evidence for this view comes from cross-language studies, since differences in the locations of warping in auditory space across languages are presumably the result of learning driven by linguistic experience. One example of such a difference is the small but systematic difference in the VOT boundary for the voiced/voiceless distinction across languages (e.g., Lisker and Abramson, 1970). Another example is the language specificity of the warping of auditory space for vowels as measured in studies of the perceptual magnet effect. In a study of 6-month-old English and Swedish infants presented with English and Swedish vowel stimuli, Kuhl et al. (1992) found that infants had more difficulty discriminating between stimuli falling near a prototypical vowel from their native language than stimuli falling near a prototypical vowel in the non-native language.

The experiments described in the current article were designed to investigate learned warpings of auditory perceptual space. Because the experiments were designed in part to test neural network models of the perceptual magnet effect (as described in the next section), and because the magnet effect is one of the most heavily studied examples of a learned warping of auditory space, we will frequently refer to it when discussing our experimental results. We do not mean to imply by this that the perceptual magnet effect should be considered as a separate phenomenon from learned instances of categorical perception.

Lieberman (1957) identified two possible learning processes that might underly categorical perception. The first, *acquired distinctiveness*, is defined as an increase in perceptual sensitivity for items that are repeatedly categorized differently in a learning situation. Lieberman (1957) reported evidence for acquired distinctiveness in detecting duration differences for speech sounds vs. non-speech sounds, and later studies provided further examples of acquired distinctiveness for non-speech stimulus sets (e.g., Lane, 1965; Goldstone, 1994). The second possible learning process identified by Lieberman (1957) was *acquired similarity*, also referred to by some authors as *acquired equivalence*. In acquired similarity, sounds that were originally distinguishable from each other become less distinguishable after repeatedly being categorized together. It has been noted that very young infants are capable of making some acoustic distinctions that become more difficult to make later in life if those distinctions are not used to differentiate phonemes in the infant's native language (e.g., Eimas, 1975; Goto, 1971; Miyawaki et al., 1975; Werker and Tees, 1984). These results appear to be examples of acquired similarity for a category-irrelevant stimulus dimension; i.e., a physical dimension which does not provide any information about category membership.

Goldstone (1994) reported another example of acquired similarity for a category-irrelevant dimension in adults performing a categorization learning task utilizing visual stimuli. Participants were trained to categorize visual stimuli that differed along two dimensions (brightness and size). For some subject groups, only one stimulus dimension was relevant for the categorization task. Goldstone (1994) found one case of acquired similarity for a category-irrelevant dimension, but no instances of acquired similarity were found for category-relevant dimensions.

However, if acquired similarity is playing a role in learned instances of categorical perception and the perceptual magnet effect, it must involve *category-relevant* dimensions. The very notion of “nearer to the category boundary” that is commonly used to describe these phenomena implies that we are talking about category-relevant dimensions, such as formant frequencies for vowels. Although attempts have been made (e.g., Goldstone, 1994), acquired similarity for a category-relevant dimension has apparently not been shown experimentally (Lieberman, 1996, pp. 18-19).

1.2. Considerations from experimental and theoretical neuroscience

It seems reasonable to assume that infants are more commonly exposed to prototypical examples of a speech sound than non-prototypical examples during the learning process that leads to the perceptual magnet effect². Perhaps relatedly, many neurophysiological studies of sensory maps have shown that disproportionately large exposure to a particular type of stimulus typically leads to a larger cortical representation for that stimulus. For example, kittens reared in a visual environment consisting only of vertical stripes have more visual cortex cells tuned to vertical contours than kittens reared in a normal environment (e.g., Rauschecker and Singer, 1981). Analogous results have been found in other sensory modalities. Preferential stimulation of a digit in monkeys leads to a larger cortical representation for that digit in somatosensory cortex (Jenkins, Merzenich, and Ochs, 1984; Jenkins, Merzenich, Ochs, Allard, and Guic-Robles, 1990). In the auditory realm, Recanzone, Schreiner, and Merzenich (1993) found that repeatedly exposing monkeys to tones in a particular frequency range during learning of a tone discrimination task resulted in an increase in the area of auditory cortex preferentially activated by sounds in the trained frequency range and a concomitant increase in the discriminability of the training tones.

It is also commonly believed that, all else being equal, stimuli that have a larger cortical representation are more easily discriminated from one another than stimuli that have a smaller cortical representation. For example, the cortical representation of the fingers in human somatosensory cortex is disproportionately large when compared to the representation of the back, and, correspondingly, humans are typically much better at discriminating tactile stimuli with their fingers than with their backs (e.g., Kandel, 1985). Similarly, the primary visual cortex representation of the high-resolution foveal area of our retinas is much larger than the representation of the low-resolution visual periphery.

If one assumes that frequent exposure to a stimulus leads to a larger cortical representation, and that larger cortical representations lead to better discriminability, then one sees a paradoxical aspect of the perceptual magnet effect: in the magnet effect, discriminability of more frequently encountered stimuli (prototypical vowels) is *worse* than discriminability of less frequently encountered stimuli (non-prototypical vowels). Two recent neural network models posit explanations for the perceptual magnet effect in terms of experience-based formation of neural maps in the auditory system (Bauer, Der, and Herrmann, 1996; Guenther and Gjaja, 1996). These models are of interest because they make clear predictions about the organization of the brain that can be tested using recently available imaging techniques such as functional magnetic resonance imaging (fMRI) or positron emission tomography (PET).

A schematic of the Guenther and Gjaja (1996) model is provided in Figure 1. The model uses two layers of neurons, referred to as the formant representation and the auditory map, connected by a set of modifiable synapses. When presented with a vowel input, the formant frequencies of the vowel are represented by the formant representation cells. Signals projecting from these cells to the auditory map through modifiable synapses lead to the activation of a subset of the cells in the auditory map. The strengths of the synapses determine which cells become active in the auditory map. The strengths of the synapses are then modified in a manner that depends on the pre- and post-synaptic cell activities, thus changing the “firing preferences” of the cells (i.e., the vowel stimuli that maximally activate the cells) in the auditory map. This process is carried out repeatedly with new vowel stimuli during a training period. The Bauer et al. (1996) model has a similar structure and function, differing primarily in the equations governing changes in the synaptic weights projecting to cells in the auditory map.

Though developed independently, the basic idea behind the Guenther and Gjaja (1996) and Bauer et al. (1996) models is the same. Exposure to vowel sounds in the early months of life causes changes in the distribution of firing preferences of neurons in an infant’s auditory map. For example, in the Guenther and

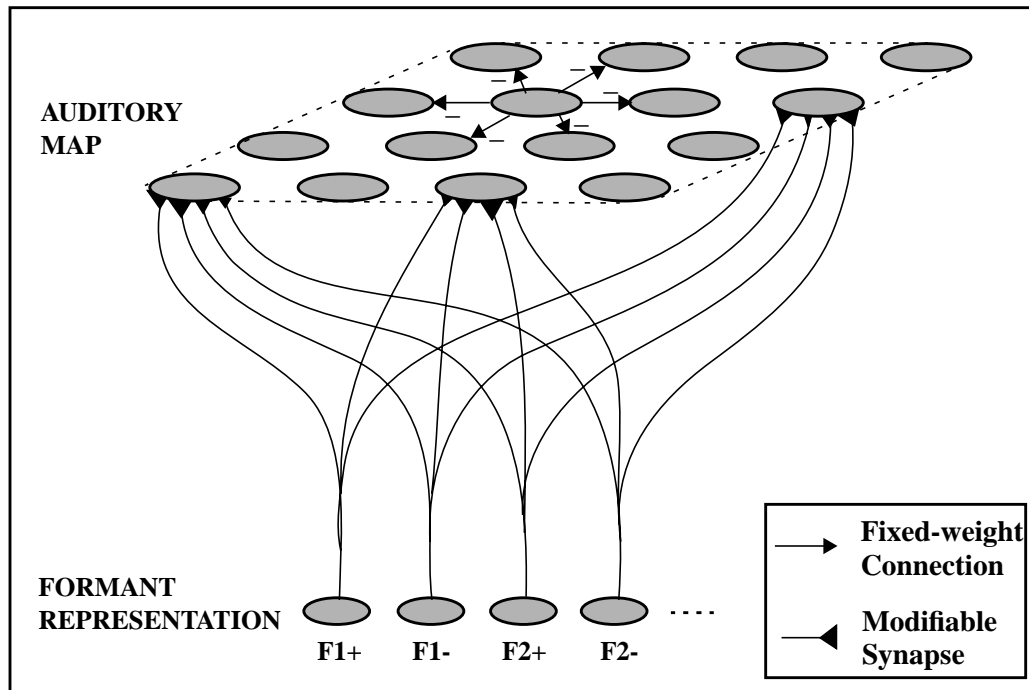


FIGURE 1. Schematic of the Guenther and Gjaja (1996) neural network model of the perceptual magnet effect. See text for details.

Gjaja (1996) model, more cells in the auditory map become tuned to the vowel sounds that the infant hears most often. It is these changes in the auditory map that underlie the perceptual magnet effect in both models.

Although the Guenther and Gjaja (1996) and Bauer et al. (1996) models are similar in many respects, they differ in how they account for the apparent paradox described above. In the Guenther and Gjaja (1996) model, the paradox is accounted for by differences in the training distributions for categorical stimuli as compared to non-categorical stimuli, where categorical stimuli are those that are typically perceived as members of a discrete set of categories (e.g., speech sounds) and non-categorical stimuli are those that are not typically perceived in this way (e.g., pure tones). Specifically, it is suggested that the training distribution of categorical stimuli has relatively sharp peaks near the category prototypes (i.e., infants hear many more examples of vowel-like sounds that fall near prototypical vowels than near non-prototypical vowels), as compared to the typically flatter distributions of non-categorical stimuli. This sharply peaked training distribution leads to a similarly peaked distribution of cell firing preferences in the neural map, and this in turn leads to a warping of perception toward the more prototypical exemplars due to population coding in the nervous system. The details of this process are presented in Guenther and Gjaja (1996); for the current purposes, it suffices to note that this model predicts that it is the distribution of training stimuli, not the type of training, that leads to the perceptual magnet effect.

In the Bauer et al. (1996) model, it is assumed that, for some stimuli, the neural map formation process leads to *smaller* cortical representations for the most frequently encountered stimuli, rather than the larger cortical representations reported in the neurophysiological studies described above. Although not treated by Bauer et al., we infer here that differences in the learning situation for categorical stimuli as compared to non-categorical stimuli lead to this difference in how the cortical representation changes size for these stimuli. In other words, whereas discrimination training leads to a larger cortical representation for the

most frequently encountered stimuli, categorization training leads to a smaller cortical representation for the most frequently encountered stimuli. Although the differential effects of different types of training in the Bauer et al. model may seem more intuitive from a learning perspective, the Guenther and Gjaja model is more in line with the traditional view of neural map formation in the computational neuroscience literature.

1.3. Goals of the current experiments

The first purpose of the current studies was to observe whether it was possible to induce acquired similarity for a category-relevant dimension of non-speech stimuli (auditory noise stimuli) using a categorization training task. This type of induced “perceptual magnet effect” is predicted by the Guenther and Gjaja (1996) model since this model attributes the reduced discriminability near a category prototype to neural map formation principles that are not specific to speech. Although this sort of acquired similarity had been identified as a possible learning mechanism underlying categorical perception several decades ago (e.g., Liberman, 1957; Lane, 1965), it apparently has not been demonstrated experimentally (Goldstone, 1994; Liberman, 1996). A second purpose of the current study was to investigate some of the learning conditions that are necessary to reduce sensitivity for frequently encountered stimuli, if it is indeed possible to induce such an effect. A final purpose of this study was to test between the Guenther and Gjaja (1996) and Bauer et al. (1996) neural models of the perceptual magnet effect in order to form a clear and testable hypothesis concerning the properties of the nervous system that lead to this effect. Most of the experimental results reported herein have been presented in preliminary form in conference publications (e.g., Husain and Guenther, 1998a,b).

2. Experiments

Four experiments were performed. All experiments consisted of four phases: a calibration phase in which a subject’s detection threshold for auditory stimuli like those used in later phases of the experiment was determined, a pre-test phase to determine baseline sensitivity, a training phase, and a post-test phase to measure any change in sensitivity that may have resulted from training. Experiments I, II, and III used the same testing procedure; these experiments differed only in the type of training the subjects underwent during the training phase. The pre- and post-tests for Experiment IV were modified slightly from the others. The common aspects of the experimental design are treated in the following paragraphs.

Participants

Subjects were male and female adults between the ages of 18 and 50 with no history of speech, language or hearing disorders. Subjects were compensated at the rate of \$8 an hour. Each subject participated in a single experimental session, consisting of a calibration phase lasting approximately 15 minutes, a pre-test lasting approximately 15 minutes, a training session lasting approximately 45 minutes, and a post-test lasting approximately 15 minutes, for a total session length of approximately 1.5 hours. No subjects were used in more than one experiment. A subject’s results were excluded from analysis if the subject did not perform within a previously determined criterion on the training task, as described below. Subjects had no prior knowledge regarding the purpose of the experiment.

Apparatus and Stimuli

The stimuli for all experiments were narrow-band filtered samples of white noise with different center frequencies. The center frequencies of the pass band ranged between 1000 Hz and 3500 Hz. The band-

widths of the stimuli were chosen to be equal in mel space, with the stimulus at 2500 Hz having a bandwidth of 100 Hz and the bandwidths of all stimuli falling within the range of 90-130 Hz. The stimuli were generated at a sampling rate of 16 kHz using Entropic's ESPS/Waves software on a Sun SPARCstation 10 by filtering white noise through a bandpass filter (a linear phase finite impulse response filter created using a weighted mean square error criterion) with a fall-off of approximately 20 dB per 100 Hz.

The sound files were sent through an Ariel D/A converter to both speakers of a set of headphones worn by the subject while sitting in a quiet room. The sounds were played at a level the subjects deemed most comfortable³, typically around 75 dB. Subjects' responses were entered using the mouse and keyboard of the computer that controlled the presentation of stimuli.

Calibration Phase

Each individual subject's threshold for discriminating the stimuli was established at the beginning of the experiment. This was done to account for rather large inter-subject differences in the ability to discriminate between the stimuli⁴. An adaptive up-down staircase method (AX same-different paradigm) was used to determine the discrimination threshold. Stimuli for this procedure consisted of narrow-band white noise centered at different frequencies around 2500 Hz with a bandwidth of 100 Hz. The step size that shifted the center frequency of the noise stimuli was fixed at 5 Hz. Thresholds were determined both for frequencies lower than and greater than 2500 Hz. The final threshold was the average of these two thresholds. This threshold, specified in mel units, was used as an estimate of the Just Noticeable Difference (JND) for that particular subject throughout the range of frequencies used in the experiment.

The stimuli for the remainder of the experiment were generated based on this JND measure, as shown in Figure 2. The hashmarks on the x axis of this figure are spaced one JND apart. First, a reference stimulus, labeled "Milestone B" in Figure 2, was located at 3200 Hz. This stimulus and six additional stimuli spaced ± 1 , 1.5, and 2 JND from it constitute the "training region" of frequency space. Next, a second reference stimulus, milestone A, was located at a frequency 11 JNDs less than milestone B. Milestone A and stimuli spaced ± 1 , 1.5, and 2 JND from it constitute the "control region" of frequency space. (The spacing used for the stimuli in Experiment IV were slightly different; this will be addressed in the description of that experiment.) In the pre- and post-tests, the subject's sensitivity to stimuli in both the control region and the training region were measure by estimating d' between the milestones A and B and their neighboring stimuli. Stimuli in the training region were involved in the training phase in a manner specific to the particular experiment; stimuli in the control region were not encountered during training. Comparison of the difference between pre- and post-test results for the training and control regions provides information about the effects of training on the listener's sensitivity to the training region stimuli. This design partially controls for shifts in response bias which may occur over the roughly 45 minute training period separating the pre-test and post-test. The training regimes used in the experiments required the use of stimuli from outside of the training region but not in the control region. These additional stimuli were chosen from a uniform distribution over two regions of frequency space labeled "band edges" in Figure 2: a 4-JND-wide region between the training region and the control region, and a 4-JND-wide region located above the training region in frequency space. There was a separation of 1.5 JNDs between the band edges and both the training and control regions.

Pre- and Post-Test Phases

Tests were conducted to measure subjects' sensitivities to differences in the auditory stimuli for both the control and training regions before and after training. The pre- and post-test sessions for the same

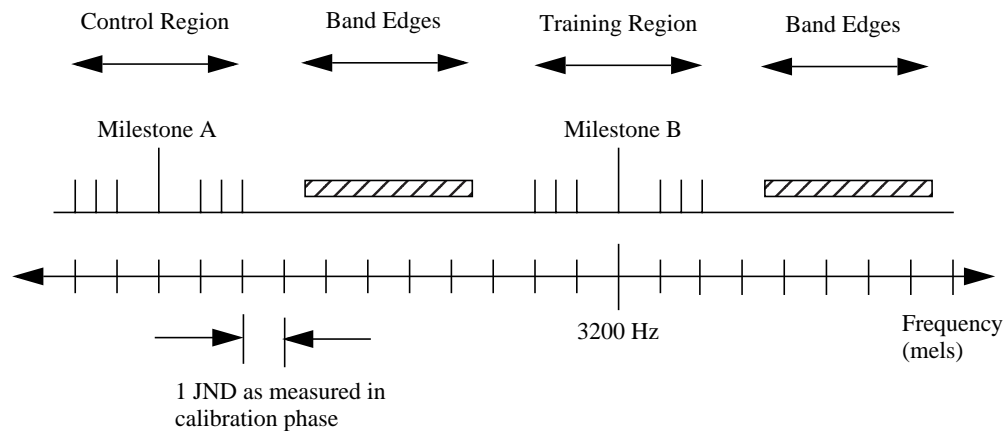


FIGURE 2. The range of frequency space within which different types of stimuli were generated. Milestone A and its neighbors form the control region and milestone B and its neighbors form the training region. Regions spanning 4 JNDs on either side of the training region are called “band edges”. See text for details.

experiment were identical. The tests measured discriminability around milestone A and milestone B (see Figure 2). Tests were conducted in two blocks of 64 trials each, using an AX same-different paradigm. One block of trials measured sensitivity in the control region, while the other block measured sensitivity in the training region. The order of presentation of the blocks was varied for different subjects, with roughly half the subjects performing tests with the control block first and the other half performing tests with the training block first. Subjects were not provided feedback concerning the correctness of their responses in the pre- and post-tests.

Each trial within a block was composed of a pair of stimuli. The first stimulus of the pair was always the milestone. The second stimulus of the trial was either the milestone again or any of its neighbors in the corresponding region, as shown in Table 1. A total of 7 distinct pairs of stimuli were generated for each block, with 6 being ‘different’ and 1 being ‘same’. There were 8 repetitions of each ‘different’ pair and 16 repetitions of the ‘same’ pair, for a total of 64 trials per block. The length of each of the stimuli in a trial was 500 msec, as shown in Figure 3. In the first three experiments, the inter-stimulus interval (ISI) was 750 msec long with a brief burst of white noise, 250 msec long, in the middle. In the fourth experiment, the ISI was 250 msec and there was no white noise between the two stimuli. The noise burst was added in the first three experiments to favor a categorical mode of sensory processing; this topic is discussed further in the description of the fourth experiment (Section 2.4).

Trials were presented in random order. Subjects indicated whether they thought the tones they heard were the same or different by pressing the ‘s’ or ‘d’ key on the keyboard. Subjects generally completed a test in 15 minutes.

The change in sensitivity between pre- and post-tests was analyzed for the group of subjects. Analysis of each subject’s sensitivity (d') was performed using both hit and false-alarm rates. Group d' scores were then calculated from these individual measures to produce a collapsed d' measure (Macmillan and Creelman, 1991; Macmillan and Kaplan, 1985). This measure has been used by researchers such as Sussman and Lauckner-Morano (1995) to investigate the perceptual magnet effect. Pairwise t -tests (Howell, 1992)

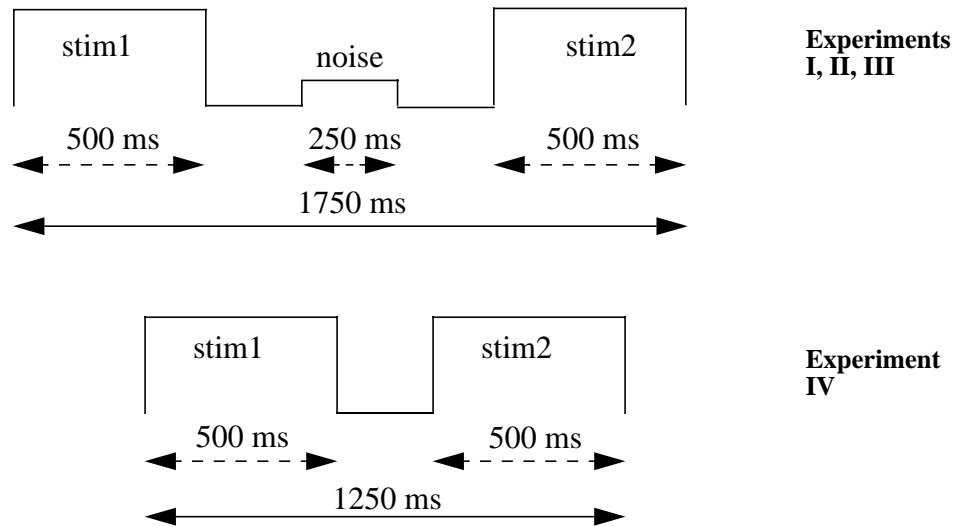


FIGURE 3. Time course of a test trial. In experiments 1, 2, and 3, the ISI was 750 msec long with a distractor noise of 250 msec. In Experiment 4, the ISI was 250 msec long with no distractor noise.

Table 1. Generation of pairs of stimuli for the AX Same-Different discrimination tests. M stands for milestone, and subscripts denote the distance from the milestone in JND units.

A	X	Type of Trial	Number
M	M	Same	16
M	M ₋₁	Different	8
M	M _{-1.5}	Different	8
M	M ₋₂	Different	8
M	M ₊₁	Different	8
M	M _{+1.5}	Different	8
M	M ₊₂	Different	8
Total			64

compared pre- and post-training d' scores to test for significant change in both the training region and the control region (whose sounds did not occur during the training session). Changes in sensitivity for the training region were then compared to changes in sensitivity for the control region.

Discriminability was compared across groups of subjects using the G statistic (Gourevitch and Galanter, 1967, p. 27) which allows for comparison of group d' measures. The G statistic tests the significance of the difference of the pre- and post-training d' scores by considering the number of observations per data point (10 subjects \times 8 trials = 80 observations per data point).

Training Phase

The type of training varied for each experiment, and the different training paradigms are explained along with the relevant experiments below. All experiments shared the following criterion for inclusion of a subject's results in the analysis: the subject must have responded correctly on half the trials of each of the ten training subsessions which comprised the training phase. If the subject did not meet this criterion, it was assumed that he/she did not succeed in learning the training task, and his/her results were thus excluded from the statistical analyses.

2.1. Experiment I

The main goal of the first experiment was to investigate whether it is possible to induce a decrease in discriminability along a category-relevant dimension of a set of non-speech stimuli that were repeatedly encountered during a training session. This would constitute a demonstration of acquired similarity along a category-relevant dimension, and it would also be in keeping with models of the perceptual magnet effect that attribute the effect to neural map formation properties that are not specific to speech (Guenther and Gjaja, 1996).

2.1.1. Training

In the training phase of Experiment I, subjects were trained to choose sounds that belonged to the training region (i.e., milestone B and its neighbors) from a list of sounds. Specifically, subjects were told that they were to learn to identify sounds from a category, referred to as the "prototype category" and corresponding to the training region of frequency space in Figure 2, and that during training they would have to choose the prototype category sound from a list of sounds that included only one member of the prototype category. Since the subjects were taught to treat the training region sounds as members of the same category, we will refer to this type of training as *categorization training*. The subjects underwent two types of training trials: (1) listening trials in which they heard example sounds from the training region and did not have to make any response, and (2) identification trials in which they identified one sound from a list of sounds as belonging to the training region. During a listening trial, subjects heard 4 sounds randomly chosen from a set of 9 sounds which were evenly spaced in 0.5 JND increments within the training region. These included the milestone B and its six neighbors used in the testing procedure, plus the two stimuli falling ± 0.5 JND from the milestone. During an identification trial, subjects heard a short list of sounds, only one of which came from the training region. The other sounds that comprised the identification trial were generated from the "band edges" regions flanking the training region (see Figure 2). These sounds were randomly chosen from a set of 18 sounds: 9 sounds spaced 0.5 JND apart from within the band edge region lower in frequency than the training region, and 9 sounds spaced 0.5 JND apart from the band edge region higher in frequency than the training region. As noted earlier, the band edge regions did not overlap with either the training or control regions, and no sounds from the control region were presented during training.

Subjects could choose whether to perform a listening or an identification task for any given trial, with the stipulation that no more than 15 listening trials could be performed in any subsession. The subsession ended when 30 identification trials were completed. Each subject performed 10 such subsessions, and sub-

jects received feedback about the correctness of their responses. Task difficulty was increased over the ten subsessions by increasing the length of the list of sounds from which the subject had to identify the training region sound: a 2-sound list was used in the first 3 subsessions, a 3-sound list was used in the next 3 subsessions, and a 4-sound list was used in the last 4 subsessions. Subjects generally completed the entire training phase in 45 minutes. Two of the twelve subjects who participated in the experiment performed below the established criterion during the training session and their data were thus excluded from subsequent analysis.

2.1.2. Results


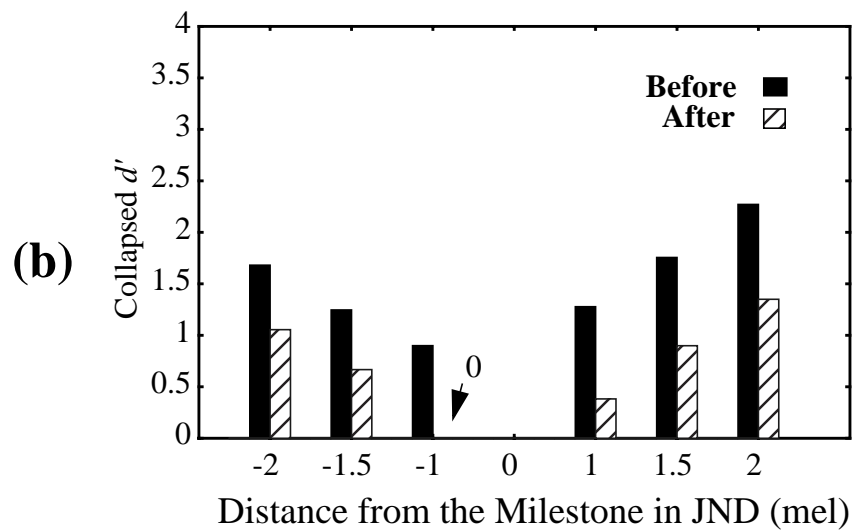
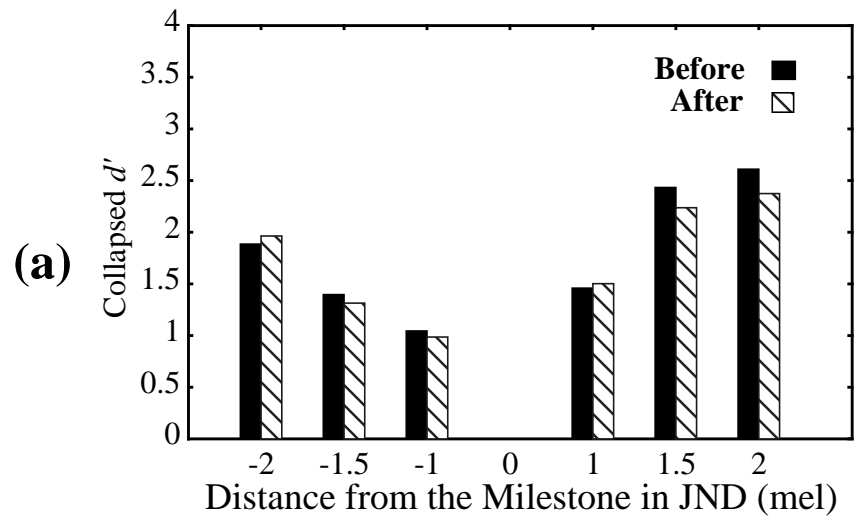
Figure 4(a) shows the collapsed d' values for sounds in the control region before and after training, and Figure 4(b) shows the same results for the training region. Subjects were significantly worse ($p < 0.05$) at discriminating stimuli in the training region after training compared to before training ($t(5) = -12.4$; $p < 0.05$) but not in the control region ($t(5) = -1.48$; $p > 0.05$). Figure 4(c) compares the change in d' for the control and training regions. The change in d' was calculated as the percentage increase or decrease in the d' from pre-test to post-test. This figure indicates that the change in sensitivity for the training set of stimuli was significantly more negative ($t(5) = -5.14$; $p < 0.05$) than the change in sensitivity for the control region. All ten subjects showed a decrease in sensitivity for the training region, and eight of the ten showed a larger sensitivity decrease in the training region than in the control region.

Discriminability before and after training was also compared across groups using Gourevitch and Galanter's (1967) G statistic. Overall, as seen in Table 2, there was a general pattern for sensitivity to worsen (indicated by the negative values) for the training region. On the other hand, sensitivity for the control region, across all the comparison steps, did not change significantly.

Table 2. G statistic comparison for Experiment I. Asterisks denote statistically significant ($p < 0.05$) changes in sensitivity.

Stimulus	Control	Training
	G score	G score
-2 JND	0.27	-2.37*
1.5JND	-0.29	-2.25*
-1 JND	-0.20	-3.54*
1 JND	0.15	-3.46*
1.5 JND	-0.63	-3.28*
2 JND	-0.73	-3.20*

FIGURE 4. [ON FOLLOWING PAGE] (a) The collapsed d' score for the control region of Experiment I, before and after training, as a function of distance from the milestone. (b) The collapsed d' score for the training region of Experiment I, before and after training, as a function of distance from the milestone. (c) Change in sensitivity after training for the control and training regions in Experiment I. Subjects showed a significant decrease in sensitivity for the stimuli in the training region but not in the control region.



2.1.3. Discussion

The results of the first experiment indicate that it is possible to induce acquired similarity along a category-relevant dimension if an appropriate training regime is utilized. Although the training region stimuli were encountered more frequently than the control region stimuli during the experiment, subjects showed a reduction in their ability to discriminate stimuli in the training region as compared to the control region. As mentioned in the Introduction, the perceptual magnet effect also appears to be a case of acquired similarity along category-relevant dimensions (formant frequencies) for more heavily experienced stimuli. The main result of Experiment I might thus be interpreted as a case of inducing a “perceptual magnet-like” effect in a non-speech modality, as predicted by the Guenther and Gjaja (1996) neural model of the perceptual magnet effect. Although Bauer et al. (1996) do not address the issue of whether the conditions leading to the magnet effect in their model are speech-specific, the results of Experiment I are not inconsistent with the Bauer et al. model if one assumes that a reduced cortical representation with heavy exposure can occur for non-speech stimuli as well as speech stimuli (see the General Discussion).

2.2. Experiment II

The second experiment tested whether a training regimen different from that used in Experiment I could result in a different effect on the subjects’ sensitivity to the training stimuli even though a similar distribution of sounds is presented during training. In this experiment, a discrimination training paradigm was used in which subjects were repeatedly asked to report whether they thought two sounds were the same or different. Subjects were given feedback concerning the correctness of their responses. One might expect that this sort of training would lead to an increase in the ability to discriminate the sounds encountered during training, as opposed to the decrease in discriminability seen in Experiment I for approximately the same distribution of training sounds. As discussed further below, such a result would be inconsistent with the Guenther and Gjaja (1996) model of the perceptual magnet effect, since that model suggests that it is the shape of the distribution of vowel-like stimuli encountered by an infant that leads to the magnet effect, not the type of training.

Twelve adults with normal hearing participated in the second experiment. Two subjects performed below the established criterion during the training session, and their data were thus excluded from subsequent analysis.

2.2.1. Training

The set of training stimuli for Experiment II was generated in an identical fashion to the training set for Experiment I, consisting of sounds from the training region and band edge regions but not the control region. Care was taken to insure that the number of times each subject heard each training sound was approximately the same as in Experiment I⁵. During the training session, subjects listened to pairs of stimuli and indicated whether they thought the two stimuli in the pair were the same or different by pressing the ‘s’ or ‘d’ key of the computer keyboard. Each training trial was of the same form as the pre- and post-test trials as described at the beginning of Section 2 (see top half of Figure 3) except that subjects were provided with feedback about the correctness of their response. Each subsession consisted of 45 trials, 15 of which involved pairs of stimuli that were the same and 30 of which involved pairs of stimuli that differed. There were 10 such subsessions within the training session, and the task difficulty of the subsessions increased as follows: the initial 3 subsessions required subjects to discriminate stimuli that were 2 JNDs apart, the next 3 subsessions involved stimuli spaced 1.5 JNDs apart, and the final 4 subsessions involved stimuli that were 1 JND apart. Subjects generally completed the training session in about 45 minutes.

2.2.2. Results

Figure 5 shows the main results for Experiment II. Figure 5(a) shows the results of the pre- and post-tests for the control region. Subjects became significantly worse at discriminating stimuli within the control region ($t(5)=-2.54, p<0.05$). This differs from the effects of training on the control region in Experiment I, where no significant change in d' was measured, though there was a very small negative change in sensitivity for the control region in that experiment. Although it is unclear why there are relatively small negative changes in d' for the control region in all four experiments (though not statistically significant in Experiment I), we suspect that this may be due to subject fatigue toward the end of the approximately 1.5 hour experimental session. Of course, fatigue would be expected to affect the training region as well, but the larger changes induced by training would make this small fatigue effect difficult to detect. Because we are primarily interested in the relative effects of training on one region of frequency space (the training region) as compared to another (the control region), the source of the small negative changes in d' for the control region was not investigated further in this paper.

Figure 5(b) shows the results of the pre- and post-tests for the training region. Subjects showed a significant increase in d' ($t(5) = 2.29, p<0.05$) after training. The increase in d' was significantly greater for the training region as compared to the control region ($t(5)=3.23, p<0.05$; see Figure 5(c)).

The general pattern for sensitivity to improve for the training region, but not the control region, is also indicated by the G scores listed in Table 3. Note that for the training region, the most positive change in sensitivity occurred to the right of the prototype of the training region. In fact, the sensitivity for the -2 JND and -1.5 JND stimuli did not change significantly. Perhaps relatedly, subjects as a group showed far fewer errors for the -2 JND and -1.5 JND stimuli during the pre-test than they showed for the other four stimuli, with only 9 total errors for the -2 JND stimulus and 20 total errors for -1.5 JND stimulus as compared to 38, 53, 32, and 27 errors respectively for the -1, 1, 1.5, and 2 JND stimuli. We thus suspect that the lack of an increase in d' for the -2 and -1.5 JND stimuli was a ceiling effect due to the very high level of sensitivity for these stimuli even before training, which was in turn apparently due to inaccuracies in calibrating the JNDs for a subject across the entire range of frequencies used in the study.

2.2.3. Discussion

The results of this experiment indicate that the same distribution of training stimuli that led to a *decrease* in sensitivity for the training region in Experiment I can lead to an *increase* in sensitivity if the training regime is changed to a discrimination training task. This is a case of acquired distinctiveness along a category-relevant dimension (see also Goldstone, 1994). Possible implications of this result for neural models of the perceptual magnet effect are treated in the General Discussion.

2.3. Experiment III

The third experiment was designed to elaborate on the training conditions required to induce the acquired similarity along a category-relevant dimension that was demonstrated in Experiment I. The specific question this experiment sought to answer was whether training with only a single exemplar from a

FIGURE 5. [ON FOLLOWING PAGE] (a) The collapsed d' scores for the control region of Experiment II, before and after training. (b) The collapsed d' score for the training region of Experiment II, before and after training. (c) Change in sensitivity after training for the control and training regions in Experiment II.

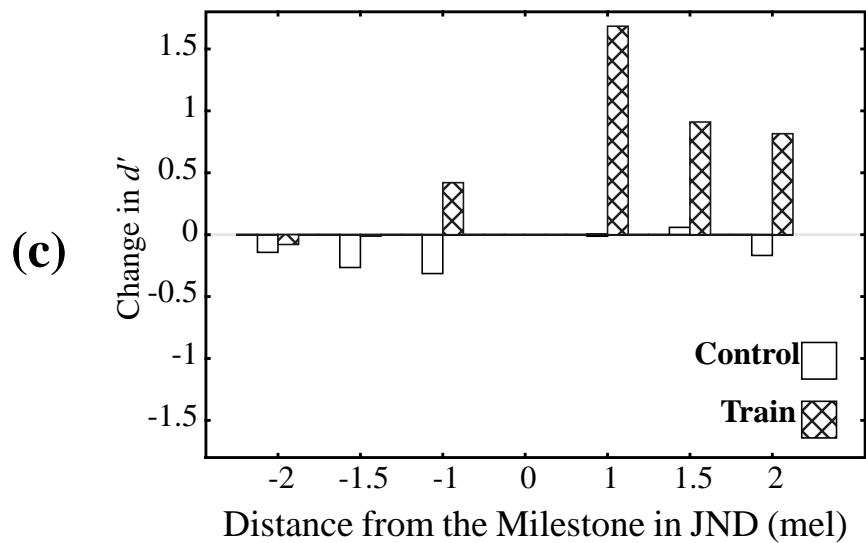
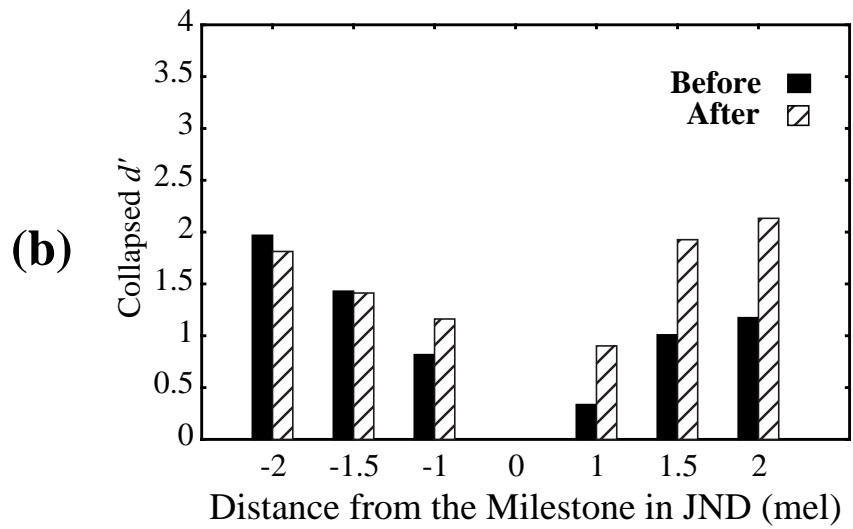
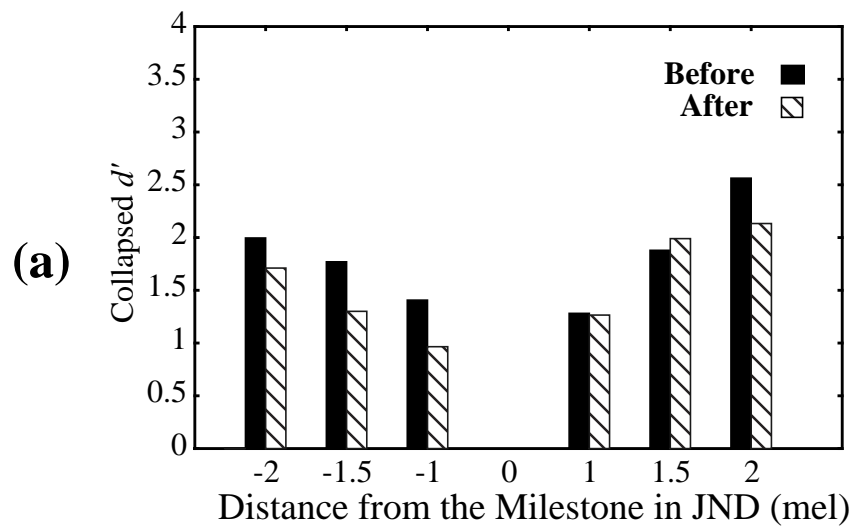


Table 3. *G* statistic comparison for Experiment II. Asterisks denote statistically significant ($p < 0.05$) changes in sensitivity.

Stimulus	Control	Training
	<i>G</i> score	<i>G</i> score
-2 JND	-0.48	-0.01
-1.5 JND	-1.21	0.45
-1 JND	-1.15	1.87*
1 JND	0.46	2.74*
1.5 JND	0.89	3.83*
2 JND	-0.80	3.81*

category is sufficient to induce decreased sensitivity in its immediate region of acoustic space. It is possible that a listener must experience many exemplars from the same category in order to induce acquired similarity. This scenario makes sense if one takes the view that acquired similarity is a case of learning to ignore differences between exemplars of the same category; if subjects hear only one exemplar of a category, there are no differences between category exemplars to learn to ignore.

Eleven adults participated in the third experiment. One subject's performance did not meet the established criterion, and this subject's results were thus not included in the analysis.




2.3.1. Training

This experiment involved a categorization training regime that differed from that of Experiment I in only one respect: instead of hearing different exemplars from the training region when performing either a listening or identification trial, subjects always heard the same exemplar, milestone B (see Figure 2).

2.3.2. Results

Figure 6 shows the main results for Experiment III. Figure 6(a) shows pre- and post-test results for the control region. As in Experiment II, subjects became significantly worse at discriminating stimuli within the control region ($t(5) = -2.98, p < 0.05$). Again, general fatigue may have been a factor in this decrease in sensitivity. Subjects also became significantly worse at discriminating stimuli in the training region (Figure 6(b); $t(5) = -2.04, p < 0.05$). More importantly, the change in sensitivity for the training region was not significantly different from the change in sensitivity for the control region ($t(5) = 0.30, p > 0.05$; see Figure 6(c)). In other words, using only a single exemplar from the training region during training did not lead to a significant decrease in discrimination performance for the training region as compared to the control region.

FIGURE 6. [ON FOLLOWING PAGE] (a) The collapsed d' scores for the control region of Experiment III, before and after training. (b) The collapsed d' scores for the training region of Experiment III, before and after training. (c) Change in sensitivity after training for the control and training regions in Experiment III.



The G scores are shown in Table 4. The change across most of the testing distances after training was not significant for either the training or the control region, except in the case of the 1.5 JND step to the right of the control milestone. The decrease for the training region was not significantly larger than the decrease for the control region. This indicates that training with only one exemplar of a category does not lead to significant acquired similarity in the immediate region of that exemplar.

Table 4. G statistic comparison for Experiment III. Asterisks denote statistically significant ($p < 0.05$) changes in sensitivity.

Stimulus	Control	Training
	G score	G score
-2 JND	1.51	-0.16
-1.5 JND	-0.02	1.05
-1 JND	-0.65	1.50
1 JND	0.8	-0.38
1.5 JND	2.44*	-0.92
2 JND	0.75	0.29

2.3.3. Discussion

In both Experiments I and III, the overall change for the control and the training regions was in the negative direction. However, the decrease in sensitivity for the training region in Experiment I was highly significant when compared to the change in the control region, while the analogous comparison in Experiment III was not significant. This suggests that a single category exemplar is not sufficient to induce acquired similarity in the neighborhood of the category exemplar, or at minimum that a single exemplar does not induce as much acquired similarity as multiple exemplars. Perhaps relatedly, Goldstone (1994) was not successful in using two exemplars to induce acquired similarity of a category-relevant dimension for visual stimuli differing along two dimensions. A possible explanation for the success in inducing acquired similarity in Experiment I and the failure to do so in Experiment III and Goldstone (1994) is that many exemplars of a category, not just one or two, are needed to noticeably decrease sensitivity along a category-relevant dimension.

2.4. Experiment IV

Several investigators have suggested that the brain's representation of sounds can be broken into two different memory modes: a continuous auditory memory mode that consists of a reasonably accurate representation of a sound that decays relatively rapidly after the stimulus goes away or is interrupted by a new auditory stimulus, and a more "discretized" or "categorical" mode that can be maintained in memory for a longer period of time, e.g. for comparison to a second stimulus in a discrimination task with a relatively large interstimulus interval (ISI). When investigating speech sounds, Pisoni (1973) referred to the different memory forms as *auditory mode* and *phonetic mode*. In a model of sound intensity discrimination, Durlach and Braida (1969) delineated two memory modes that they termed *sensory-trace mode* and *context-coding mode*; these modes are roughly analogous to Pisoni's auditory mode and phonetic mode, respectively. Macmillan, Goldberg, and Braida (1988) extended the Durlach and Braida (1969) model to explain exper-

imental results involving speech stimuli. Since we are not dealing with speech stimuli directly in this experiment, we will use the terms *sensory-trace mode* and *context-coding mode* here.

The purpose of Experiment IV was to determine whether the acquired similarity induced in Experiment I could be better characterized as a result of changes in the sensory-trace mode or the context-coding mode of auditory memory. It is usually assumed that increasing the ISI and/or adding a brief noise burst between two stimuli interferes with the sensory-trace mode of memory more than context-coding mode (e.g., Repp, 1984; Werker and Pegg, 1982). Given the relatively long ISI of Experiment I and the use of a noise burst between the two stimuli in a discrimination trial, one might reasonably conclude that the effect measured in that experiment primarily involved the context-coding mode of auditory memory. In Experiment IV, the ISI during discrimination training was reduced and the inter-stimulus noise was removed in order to better gauge whether the acquired similarity demonstrated in Experiment I is also manifested in the sensory-trace mode of auditory memory.

2.4.1. Training and testing

The training and testing stimuli used in Experiment IV are shown in Figure 7. The training regime for Experiment IV was identical to that of Experiment I, and the training stimuli were generated in the exact same fashion as in that experiment. The testing procedure for Experiment IV involved an ISI of 250 ms and there was no distractor noise between the two stimuli (see Figure 3). In a pilot experiment, it was determined that these manipulations allowed subjects to discriminate the test stimuli almost perfectly. This invalidated the d' measures since they are only accurate if a significant number of errors are made during testing. In order to obtain an accurate d' measure with the shorter ISI, the stimuli used in the testing sessions of Experiment IV had to be more closely spaced than they were in the earlier experiments. Test stimuli for Experiment IV were located at 0.75, 1.125, and 1.5 JND units⁶ above and below the milestones in the control and training regions, as compared to a spacing of 1, 1.5 and 2 JND units in Experiment I. The placement of the milestones and the positioning of the band edges regions were not affected by this change.

2.4.2. Results

Figure 8 shows the collapsed d' scores for the control region (Figure 8(a)) and training region (7(b)) before and after training. A significant decrease in sensitivity occurred for both the control region ($t(5)=-5$, $p<0.05$) and the training region ($t(5)=-3.8$, $p<0.05$). The change in the training region was not significantly different from the change in the control region ($t(5)=-0.63$, $p>0.05$; see Figure 8(c)). The G scores for Experiment IV are presented in Table 5, with the group change in d' reaching significance for only one test stimulus (-1 JND in the training region).

2.4.3. Discussion

The results of this experiment indicate that the use of a shorter ISI and no noise burst between the two stimuli in the sensitivity testing trials essentially eradicates the acquired similarity found in Experiment I, despite the use of the same training regime as in that experiment. Since decreasing the ISI and removing the noise burst presumably favors a sensory-trace memory mode over a context-coding memory mode, this result suggests that the acquired similarity seen in Experiment I was primarily associated with the context-coding mode of auditory short term memory. This result is consistent with the hypotheses of Macmillan et al. (1988), Pisoni (1973), Repp (1984), and Werker and Pegg (1992) that a shorter ISI can diminish the categorical nature of the responses made by an observer.

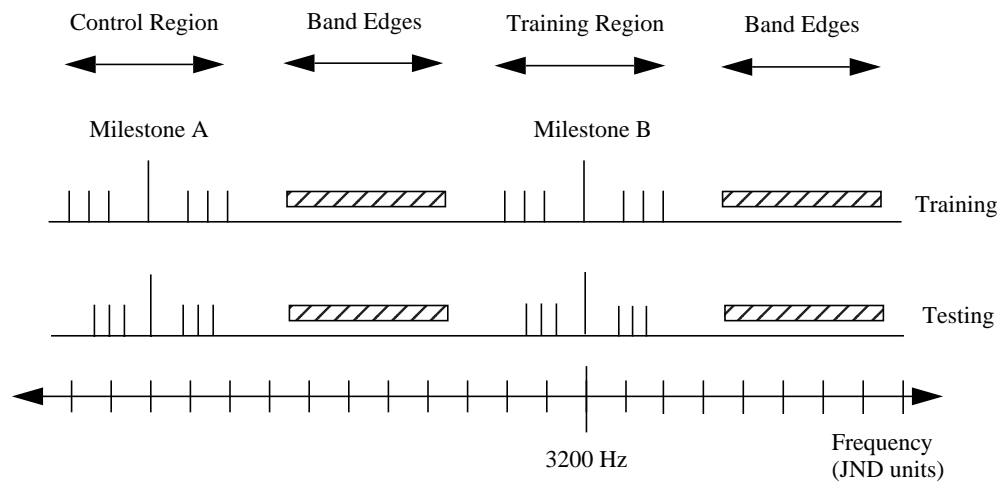
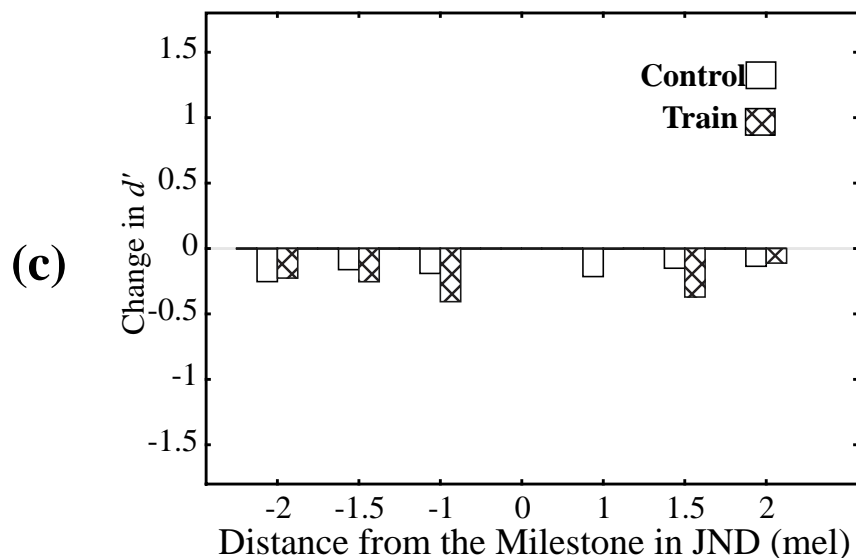




FIGURE 7. The training and testing stimuli for Experiment IV. Training stimuli were generated in exactly the same manner used in Experiment I. Testing stimuli were more closely spaced than in Experiments I-III to compensate for increased discriminability of the test sounds due to the shorter ISI and removal of the inter-stimulus noise burst. See text for details.

Table 5. *G* statistic comparison for Experiment IV. Asterisks denote statistically significant ($p < 0.05$) changes in sensitivity.

Stimulus	Control	Training
	<i>G</i> score	<i>G</i> score
-2 JND	1.56	1.39
-1.5 JND	0.54	1.42
-1 JND	0.23	1.90*
1 JND	0.32	-0.53
1.5 JND	0.31	1.35
2 JND	0.45	0.10

FIGURE 8. [ON FOLLOWING PAGE] (a) The collapsed d' scores for the control region of Experiment IV, before and after training. (b) The collapsed d' scores for the training region of Experiment III, before and after training. (c) Change in sensitivity after training for the control and training regions in Experiment IV.



3. General Discussion

Figure 9 is a composite plot of the total d' measures collapsed across subjects before and after training in all four experiments. The left side of this figure illustrates that the change in sensitivity in the control region due to training in all four experiments was negative, though this change was relatively small and did not reach statistical significance in Experiment I. Because the control region stimuli were not presented during training, we suspect that these small negative changes in d' were the result of generally poorer performance in the post-test as compared to the pre-test, perhaps due to subject fatigue near the end of the roughly 1.5 hour-long experimental session (see Section 2.2.2).




FIGURE 9. Total d' in the control and training regions before and after training for Experiments I through IV.

The right half of Figure 9 illustrates the d' measures for the training region before and after training. The results of the first two experiments indicate that, depending on the training regime, it is possible to induce either an increase or a decrease in the discriminability of a set of auditory stimuli. The first experiment indicated that categorization training, in which subjects were asked to identify sounds belonging to a small region of frequency space as members of the same category, led to a decrease in the discriminability of stimuli within this small range. That is, subjects exhibited acquired similarity along the category-relevant dimension of center frequency of the narrow-band noise stimuli. The third and fourth experiments helped elucidate some of the necessary conditions for attaining this acquired similarity. In Experiment III, the small range of frequencies corresponding to the learned category in Experiment I was shrunk down to a single exemplar during training. This eliminated the acquired similarity seen in Experiment I, suggesting that a listener needs to be exposed to different examples of a category during training, not just a single exemplar, in order to decrease the listener's ability to discriminate between stimuli falling near the center

of the category. In Experiment IV, it was shown that a testing regime that favors a hypothesized sensory-trace mode of auditory memory over a context-coding mode (e.g., Durlach and Braida, 1969; Pisoni, 1973; Macmillan, Goldberg, and Braida, 1988) weakens the acquired similarity effect of training, suggesting that categorization training primarily affects the context mode of memory processing.

If we are to fully understand the neural processes that lead to experience-based warpings of auditory space such as the perceptual magnet effect, it is important to formulate and test hypotheses regarding the neural mechanisms underlying these phenomena. The Guenther and Gjaja (1996) and Bauer et al. (1996) models of the perceptual magnet effect attribute it to neural map formation properties in auditory brain areas such as the primary auditory cortex. According to both of these models, the learning process during which infants develop phonemic categories involves a change in the distribution of firing preferences of cells in auditory cortex. This change in the auditory neural map for vowel-like sounds is hypothesized to underlie the perceptual magnet effect.

Because the Guenther and Gjaja (1996) model posits that the magnet effect results from neural map formation properties that are not specific to speech stimuli, it predicts that exposing a listener to new, non-speech auditory stimuli within a training regime that appropriately mimics the learning of phonemic categories by an infant should lead to a similar change in the distribution of firing preferences of cells coding these stimuli in auditory cortex. This change in the auditory neural map should in turn result in a measurable “perceptual magnet-like” effect for these auditory stimuli. That is, we should see a decreased ability for subjects to discriminate the training stimuli. The results of Experiment I supported the prediction of decreased discrimination in the heavily experienced training region due to categorization training. Though not predicted by Bauer et al. (1996), the results of Experiment I are not inconsistent with their model if one assumes that the reduced cortical representation for heavily experienced sounds that underlies the magnet effect in the model results from a particular kind of training, rather than from speech-specific neural mechanisms, as we hypothesize below.

The results of Experiment II indicated that the decrease in sensitivity was related to the categorical nature of the training task used in Experiment I, since a discrimination training task led to an *increase* in the ability to discriminate training stimuli in Experiment II. This result conflicts with the Guenther and Gjaja (1996) model, since this model posits that it is the distribution of training stimuli, not the type of training, that leads to the magnet effect. Bauer et al. (1996) do not speculate on what training conditions might be required to induce a perceptual magnet-like effect, but their model allows for different training conditions to have different effects on the size of the representation of training stimuli in the neural map. Based on the results of the current experiments, we propose that discrimination training and categorization training have opposite effects on the size of the neural representation of the training stimuli. This hypothesis, in combination with the neural map model of Bauer et al. (1996), is schematized in Figure 10. The left side of the figure corresponds to a categorization training situation, as in Experiment I. The top and bottom panels schematize the auditory map as a function of acoustic space before and after training, and the middle panel schematizes the distribution of training stimuli in acoustic space. In categorization training, heavy exposure to a set of training sounds leads to fewer cells coding these sounds in the auditory map, and the resulting smaller cortical representation diminishes a listener’s ability to differentiate sounds in this region of acoustic space. This is how the Bauer et al. (1996) model, with an appropriate parameter choice that leads to a negative magnification factor for the cortical representation, accounts for the perceptual magnet effect. The right side of Figure 10 corresponds to a discrimination training situation, as in Experiment II. Here, more cells in the map become tuned to the most frequently encountered training stimuli, and the resulting larger cortical representation increases the listener’s ability to differentiate sounds in this region of acoustic space. This learning situation corresponds to the “classical” formulation of a self-organizing feature map in the computational neuroscience literature, in which increased exposure to a set of stimuli leads to a larger

cortical representation for those stimuli (e.g., von der Malsburg, 1973; Grossberg, 1976; Kohonen, 1982), and can also be accounted for by using a positive magnification factor in the Bauer et al (1996) model. We are currently testing predictions of the hypothesis illustrated in Figure 10 using functional magnetic resonance imaging techniques.

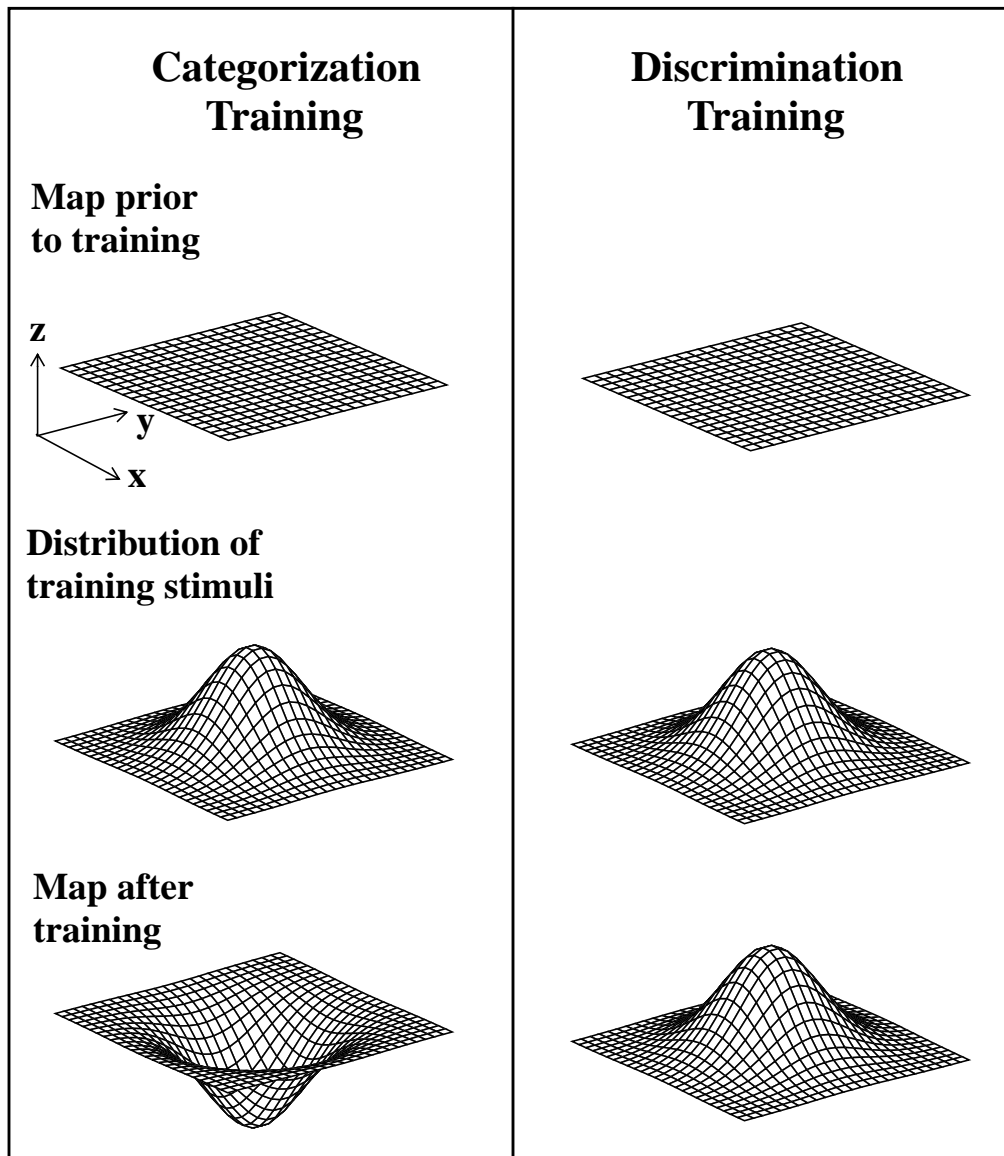


FIGURE 10. Hypothesized changes in the neural map in auditory cortex as a result of categorization training (left; Experiment 1) and discrimination training (right; Experiment 2). The x and y axes of all plots correspond to two acoustic dimensions, such as the first two formant frequencies. The z axis corresponds to the number of cells in the map devoted to each region of frequency space (top and bottom plots) or the number of training stimuli from that region of frequency space (middle plots). Categorization training leads to a decrease in the number of cells coding the most frequently encountered stimuli, whereas discrimination training leads to an increase in the number of cells coding the most frequently encountered stimuli.

References

- Aaltonen, O., Eerola, O., Hellström, Å, Uusipaikka, E., and Lang, A.H. (1997). Perceptual magnet effect in the light of behavioral and psychophysiological data. *Journal of the Acoustical Society of America*, **101**, 1090-1105.
- Bauer, H.-U., Der, R., and Herrmann, M. (1996). Controlling the magnification factor of self-organizing feature maps. *Neural Computation*, **8**, pp. 757-771.
- Beale, J. M., and Keil, F. C. (1995). Categorical effects in the perception of faces. *Cognition*, **57**, pp. 217-239.
- Burns, E.M., and Ward, W.D. (1978). Categorical perception -- phonemenon or epiphenomenon: Evidence from experiments in the perception of melodic musical intervals. *Journal of the Acoustical Society of America*, **63**, pp. 456-468.
- Durlach, N.I., and Braida, L.D. (1969). Intensity perception. I. Preliminary theory of intensity resolution. *Journal of the Acoustical Society of America*, **46**(2) pt. 2, pp. 372-383.
- Eimas, P.D. (1963). The relation between identification and discrimination along speech and non-speech continua. *Language and Speech*, **6**, 206-217.
- Eimas, P.D. (1975). Auditory and phonetic coding of the cues for speech: Discrimination of the (r-l) distinction by young infants. *Perception & Psychophysics*, **18**, pp. 341-347.
- Eimas, P.D., Siqueland, E.R., Jusczyk, P., and Vigorito, J. (1971). Speech perception in infants. *Science*, **171**, 303-306.
- Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology*, **123**, pp. 178-200.
- Gourevitch, V. and Galanter, E. (1967). A significance test for one-parameter isosensitivity functions. *Psychometrika*, **32**, pp. 25-33.
- Goto, H. (1971). Auditory perception by normal Japanese adults of the sounds "l" and "r". *Neuropsychologia*, **9**, pp. 317-323.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, **23**, pp. 121-134.
- Guenther, F.H., and Gjaja, M.N. (1996). The perceptual magnet effect as an emergent property of neural map formation. *Journal of the Acoustical Society of America*, **100**, pp. 1111-1121.
- Howell, D. (1992). *Statistical methods for psychology*. Belmont, California: Duxbury Press.
- Husain, F.T., and Guenther, F.H. (1998a). Inducing a "perceptual magnet"-like effect in a non-speech modality. Program of the 16th International Congress on Acoustics and the 135th Meeting of the Acoustical Society of America, *Journal of the Acoustical Society of America*, **103**(5), Pt. 2, p. 2982.
- Husain, F.T., and Guenther, F.H. (1998b). Experimental tests of the neural models of the perceptual magnet effect. *Proceedings of the Second International Conference on Cognitive and Neural Systems*, Boston University, May 28-30, 1998.
- Iverson, P., Diesch, E., Siebert, C., and Kuhl, P.K. (1994). Cross-language tests of the perceptual magnet effect for /ɹ/ and /l/. *Journal of the Acoustical Society of America*, **96**(5), Pt. 2, p. 3228.
- Iverson, P., and Kuhl, P.K. (1994). Tests of the perceptual magnet effect for American English /r/ and /l/. *Journal of the Acoustical Society of America*, **95**(5), Pt. 2, p. 2976.
- Iverson, P., and Kuhl, P.K. (1995). Mapping the perceptual magnet effect for speech using signal detection theory and multidimensional scaling. *Journal of the Acoustical Society of America*, **97**, pp. 553-562.

- Jenkins, W.M., Merzenich, M.M., and Ochs, M.T. (1984). Behaviorally controlled differential use of restricted hand surfaces induce changes in the cortical representation of the hand in area 3b of adult owl monkeys. *Society for Neuroscience Abstracts*, **10**, p. 665.
- Jenkins, W.M., Merzenich, M.M., Ochs, M.T., Allard, T., and Guíc-Robles, E. (1990). Functional reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally controlled tactile stimulation. *Journal of Neurophysiology*, **63**, pp. 82-104.
- Jusczyk, P. (1986). Speech perception. In K.R. Boff, L. Kaufman, and J.P. Thomas (Eds.), *Handbook of perception and human performance* (Chapter 27). New York: Wiley.
- Kandel, E.R. (1985). The central representation of touch. In E.R. Kandel and J.H. Schwartz (Eds.), *Principles of neural science* (pp. 316-330). New York: Elsevier.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, **43**, pp. 59-69.
- Kuhl, P.K. (1981). Discrimination of speech by nonhuman animals: Basic auditory sensitivities conducive to the perception of speech-sound categories. *Journal of the Acoustical Society of America*, **70**, 340-349.
- Kuhl, P.K. (1991). Human adults and human infants show a 'perceptual magnet effect' for the prototypes of speech categories, monkeys do not. *Perception & Psychophysics*, **50**, pp. 93-107.
- Kuhl, P.K. (1995). Mechanisms of developmental change in speech and language. In Elenius, K., and Branderud, P. (eds.): *Proceedings of the XIIIth International Congress of Phonetic Sciences*, vol. 2, pp. 132-139. Stockholm: KTH and Stockholm University.
- Kuhl, P.K., and Miller, J.D. (1975). Speech perception by the chinchilla: Voiced-voiceless distinction in alveolar plosive consonants. *Science*, **190**, 69-72.
- Kuhl, P.K., and Miller, J.D. (1978). Speech perception by the chinchilla: Identification functions for synthetic VOT stimuli. *Journal of the Acoustical Society of America*, **63**, 905-917.
- Kuhl, P.K., and Padden, D.M. (1982). Enhanced discriminability at the phonetic boundaries for the voicing feature in macaques. *Perception and Psychophysics*, **32**, 542-550.
- Kuhl, P.K., and Padden, D.M. (1983). Enhanced discriminability at the phonetic boundaries for the place feature in macaques. *Journal of the Acoustical Society of America*, **73**, 1003-1010.
- Kuhl, P.K., Williams, K.A., Lacerda, F., Stevens, K.N., and Lindblom, B. (1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science*, **255**, pp. 606-608.
- Lane, H. (1965). The motor theory of speech perception: A critical review. *Psychological Review*, **72**, pp. 275-309.
- Lieberman, A.M. (1957). Some results of research on speech perception. *Journal of the Acoustical Society of America*, **29**, pp. 117-123.
- Lieberman, A.M. (1996). *Speech: A special code*. Cambridge, MA: MIT Press.
- Lieberman, A.M., Harris, K.S., Eimas, P., Lisker, L., and Bastian, J. (1961). An effect of learning on speech perception: The discrimination of durations of silence with and without phonemic significance. *Language and Speech*, **4**, 175-195.
- Lieberman, A.M., Harris, K.S., Hoffman, H.S., and Griffith, B.C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, **54**, 358-368.
- Lieberman, A.M., Harris, K.S., Kinney, J.A., and Lane, H. (1961). The discrimination of relative onset time of the components of certain speech and nonspeech patterns. *Journal of Experimental Psychology*, **61**, 379-388.

- Lieberman, P, and Blumstein, S.E. (1988). *Speech physiology, speech perception, and acoustic phonetics*. Cambridge, U.K.: Cambridge University Press.
- Lisker, L., and Abramson, A.S. (1970). The voicing dimension: Some experiments in comparative phonetics. *Proceedings of the Sixth International Congress of Phonetic Sciences, Prague, 1967* (pp. 563-567). Prague: Academia.
- Lotto, A.J., Kluender, K.R., and Holt, L.L. (1998). Depolarizing the perceptual magnet effect. *Journal of the Acoustical Society of America*, **103**, 3648-3655.
- Macmillan, N. and Creelman, C. (1991). *Detection theory: A user's guide*. Cambridge: Cambridge University Press.
- MacMillan, N.A., Goldberg, R.F., and Braida, L.D. (1988). Resolution for speech sounds: Basic sensitivity and context memory on vowel and consonant continua. *Journal of the Acoustical Society of America*, **84**, pp. 1262-1280.
- Macmillan, N. and Kaplan, H. (1985). Detection theory analysis of group data: Estimating sensitivity from average hit and false alarm rates. *Psychological Bulletin*, **98**, pp. 185-199.
- Miyawaki, K., Strange, W., Verbrugge, R., Liberman, A. M., Jenkins, J. J., & Fujimura, O. (1975). An effect of linguistic experience: the discrimination of /r/ and /l/ by native speakers of Japanese and English. *Perception and Psychophysics*, **18**, pp. 331-340.
- Pisoni, D.B. (1973). Auditory and phonetic memory codes in the discrimination of consonants and vowels. *Perception & Psychophysics*, **13**, pp. 253-260.
- Rauschecker, J.P., and Singer, W. (1981). The effects of early visual experience on the cat's visual cortex and their possible explanation by Hebb synapses. *Journal of Physiology*, **310**, pp. 215-239.
- Recanzone, G.H., Schreiner, C.E., and Merzenich, M.M. (1993). Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *Journal of Neuroscience*, **13**, pp. 87-103.
- Repp, B. (1984). Categorical perception: Issues, methods, and findings. In N.J. Lass (Ed.), *Speech and language: Advances in basic research and practice* (Vol. 10). New York: Academic Press.
- Sussman, J.E., and Lauckner-Morano, V.J. (1995). Further tests of the "perceptual magnet effect" in the perception of [i]: Identification and change/no-change discrimination. *Journal of the Acoustical Society of America*, **97**, pp. 539-552.
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striata cortex. *Kybernetik*, **14**, pp. 85-100.
- Werker, J.F. and Pegg, J.E. (1992). Infant speech perception and phonological acquisition. In Ferguson, C.A., Menn, L., and Stoel-Gammon, C. (Eds.), *Phonological Development: Models, research, implications*. Parkton, MD: York Press.
- Werker, J.F., and Tees, R.C. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development*, **7**, pp. 49-63.

Footnotes

1. The term “warping” in this article refers to nonuniformities above and beyond the roughly logarithmic relationship between perceptual space and frequency in Hertz, as estimated by the Bark and Mel scales.
2. It is difficult to ascertain the typical distribution of speech sounds heard by an infant during the first years of life, particularly given that most infants are exposed to “motherese” in which phonemes are often spoken in an exaggerated fashion as compared to casual speech. If one simply counts the number of occurrences in the Peterson and Barney vowel formant frequency database of /i/ sounds falling within a 120-mel radius of the prototypical and non-prototypical /i/ sounds used by Kuhl (1991), one finds that there are indeed more /i/ examples near the prototypical /i/ than near the non-prototypical /i/. This evidence should be viewed as weak, however, as no female or child utterances of /i/ in the database fall within the 120-mel radius of either the prototypical or non-prototypical /i/ of Kuhl (1991).
3. Steps were taken to ensure that stimuli differing in center frequency were played at the same absolute intensity level (dB SPL).
4. Typical JND measures determined in the calibration phase ranged between 10 Hz and 50 Hz for the different subjects.
5. Because subjects in Experiment I could choose to perform fewer than 15 listening trials during each training sub-session, the total number of times that each sample was heard during training varied from subject to subject. However, subjects usually used all 15 listening trials per sub-session. The training stimulus distribution for Experiment II was thus chosen to match the training distribution for Experiment I under the assumption that all listening trials were used.
6. Because the JND was estimated in the calibration phase using a longer ISI and an inter-stimulus noise burst, it is expected to be larger than the JND for the stimuli as presented during the pre- and post-tests, which had a shorter ISI and no inter-stimulus noise burst.