

Learning Sound Categories: A Neural Model and Supporting Experiments

Frank H. Guenther^{1,2} and Jason W. Bohland¹

¹Boston University, Boston, MA, USA

²Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT

Our ability to discriminate sounds such as vowels is not uniform throughout acoustic space. That is, our auditory perceptual spaces are warped representations of acoustic space. One example of auditory space warping, the *perceptual magnet effect*, arises from exposure to the phonemes of an infant's native language. We have developed a neural model that accounts for this effect. The model is based on the idea that category learning during infancy changes the distribution of the firing preferences of neurons in auditory cortical maps and thus changes the discriminability of sounds from different parts of acoustic space. The model predicts that it should be possible to induce a perceptual magnet effect for non-speech stimuli. This prediction was verified by a psychophysical experiment in which subjects underwent categorization training involving non-speech auditory stimuli that were not "categorical" prior to training. The model further predicts that the magnet effect arises because prototypical vowels have smaller auditory cortical representations than non-prototypical vowels. This prediction was supported by a functional magnetic resonance imaging (fMRI) experiment involving prototypical and non-prototypical examples of the vowel /i/. The model thus provides an account of phoneme category learning that unifies observations from auditory psychophysics, cortical neurophysiology, and neural modeling.

1. Introduction

The vowels of the world's languages can be largely differentiated based on the first three formant frequencies of the acoustic signal. These formant frequencies can be viewed as a three-dimensional acoustic space, and different languages break up this space into different vowel systems. In English, for example, approximately twelve "pure" vowel categories occupy the vocalic region of formant space, whereas the Spanish and Japanese languages have only five vowels occupying this same space. When an infant enters the world amidst speakers of a particular language, she/he must somehow learn to parse speech signals into the set of phonemes of the native language. For example, she must learn to break formant space up into the appropriate number of discrete regions corresponding to the vowels of her local language, and she must learn where each vowel is located within this space.

Our ability to sort the speech signal into sound categories is aided by the fact that our ability to discriminate sounds is not uniform throughout acoustic space. We are generally better at discriminating between-category differences in sounds than within-category differences, a phenomenon often referred to as *categorical perception*, particularly when applied to consonant perception. Vowel perception is subject to a similar type of auditory space warping, often referred to as the *perceptual magnet effect* (Kuhl, 1991). In this effect, prototypical examples of a vowel or semi-vowel are more difficult to discriminate from each other than non-prototypical examples located near category boundaries. Unlike some instances of categorical perception, which appear to result from "hardwired" aspects of the auditory system, the perceptual magnet effect appears to arise due to linguistic experience, since 6-month-old American babies show the effect for an American vowel but not a Swedish vowel, while Swedish babies show the opposite effect (Kuhl et al., 1992).

We have developed, experimentally tested, and refined a neural model that explains the perceptual magnet effect as the result of changes to neural maps in auditory cortical areas (Guenther and Gjaja, 1996; Guenther, Husain, Cohen, and Shinn-Cunningham, 1999; Guenther, Nieto-Castanon, Tourville, and Ghosh 2000). These changes are hypothesized to occur during vowel category learning in infancy. In this article, we describe the model and

present the results of psychophysical and brain imaging experiments that support its account of the perceptual magnet effect and, more generally, the effects of categorization training on sensory cortical maps.

2. Effects of stimulus exposure on sensory cortical maps

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, Ochs, Allard, and Guíc-Robles, 1990). In the auditory realm, Recanzone, Schreiner, and Merzenich (1993) repeatedly exposed monkeys to tones in a particular frequency range as part of a tone discrimination task. This resulted in an increase in the area of auditory cortex preferentially activated by sounds in the trained frequency range, and there was a concomitant increase in the discriminability of these training tones.

The finding that sensory neural maps grow with heavy stimulus exposure has been explained by neural network models commonly referred to as self-organizing feature maps. Figure 1 schematizes a typical self-organizing feature map. Roughly speaking, a self-organizing feature map contains a subcortical layer (or layers) of cells and a cortical layer of cells. The subcortical layer represents incoming sensory stimuli. Cells in the cortical layer compete with each other through inhibitory connections, with only the cells receiving the largest total input from the subcortical layer becoming active when a stimulus is presented. The amount of input to each cortical cell depends on the synaptic weights between the subcortical layer and the cortical layer. When a stimulus is presented, the weights projecting to the cortical cells that “win” the competition are changed in such a way that those cells become even more likely to win the competition the next time the same stimulus pattern is presented to the network.

In the standard formulation of a self-organizing feature map, 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). Furthermore, it is widely believed that, all else equal, larger sensory cortical representations lead to better discriminability of the represented stimuli. For example, we have a larger somatosensory cortical representation (i.e., more cortical area per unit of skin surface) for our fingertips than our forearms, and we are better at two-point discrimination with a fingertip than with the forearm. This relationship makes sense when one considers that neural representations involving larger numbers of cells can better average out the noisy signals of individual neurons.

Although prototypical vowels are presumably much more commonly experienced by a listener than non-prototypical vowel-like sounds, the perceptual magnet effect indicates that listeners are *worse* at discriminating prototypical vowels. This clearly conflicts with the classical formulation of a self-organizing feature map, since in a classical self-organizing feature map there would be a larger representation for prototypical vowels, leading to better discriminability as compared to non-prototypical vowel-like sounds having smaller cortical representations. As described in the next section, we have developed a neural model based on the self-organizing feature map architecture that accounts for the perceptual magnet effect.

3. A neural model of the effects of category learning on cortical maps

The model consists of two layers: a formant representation layer and an auditory map layer. These correspond to the thalamic map and cortical map in the self-organizing feature map shown in Figure 1. It is assumed that the early stages of the auditory system process the incoming speech signal and extract important auditory information such as formant frequencies, which are represented in the first layer of the model. Each formant layer cell projects to each cell in the auditory map layer. Cells in the auditory map are tuned to a particular point in the formant frequency space, and thus fire maximally for that particular input.

In the simulations reported below, learning of the auditory map relies on a modification of Kohonen’s self-organizing feature map algorithm (1982). In Kohonen’s formulation, the *learning rate*, η , is a scalar that decays exponentially in time and modulates the amount of synaptic weight change that can occur in a given time step. In our model, η is also modulated by the Euclidean distance between the current input and the nearest category

center. As a result, the learning rate will be minimal when the input coincides with a category center and maximal for inputs that are near the boundaries of categories.

The “output” of the model is a *population response* in the auditory map layer. We assume that all cells with tuning centers within some distance r from the current input contribute to a population code that denotes the perceived speech sound. The population response is the weighted average of the active cells’ tuning centers, with each cell’s contribution weighted by the cell’s activation level.

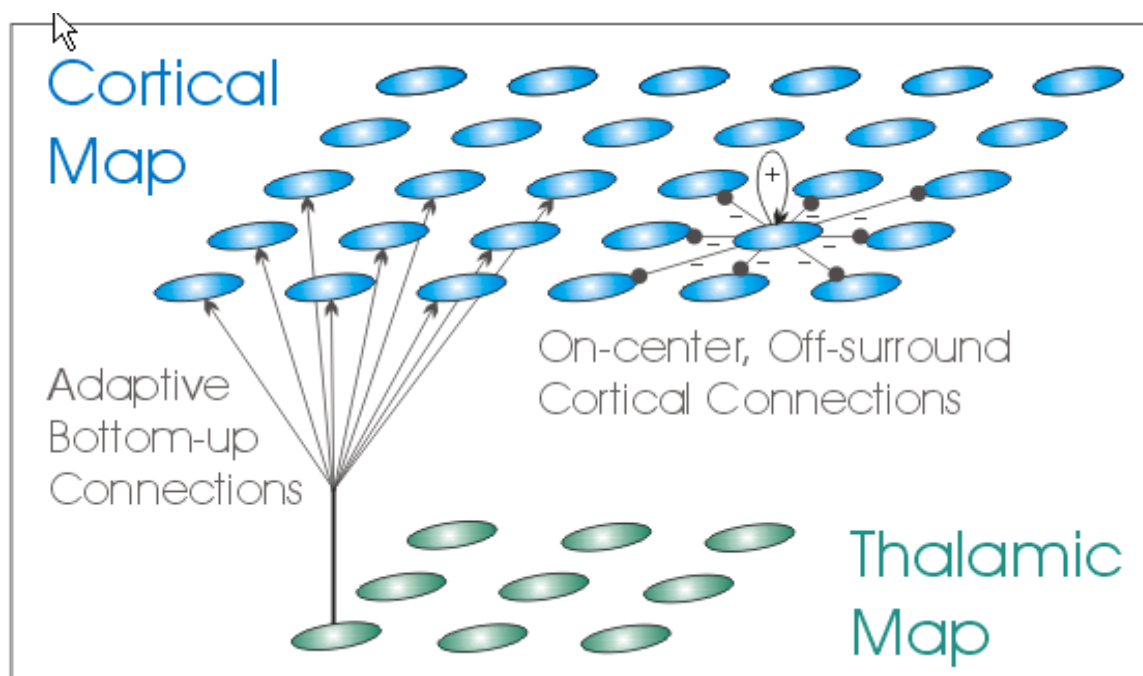


Figure 1. The basic architecture of a self-organizing feature map neural network.

According to this model, it is the type of training an infant undergoes with speech sounds, not just the distribution of training stimuli, that leads to a shrinking of the neural map for speech sound stimuli. The left side of Figure 2 schematizes the effects of phoneme category learning on auditory cortical maps. We hypothesize that infants are presented with examples of a phoneme that form a roughly Gaussian distribution centered on prototypical examples of the sound. This is schematized in the top left panel of Figure 2. We further hypothesize that the process of learning to treat these examples as members of the same sound category leads to a decrease in the size of the auditory cortical representation for sounds from near the center of the category, as schematized in the bottom left panel of Figure 2. In other words, category learning leads to a decrease in the number of cells representing category prototypes in auditory cortical maps. This contrasts with the effects of discrimination training, shown in the right half of Figure 2. As reported by Recanzone et al. (1993), discrimination training leads to an increase in the size of the cortical representation for the most frequently encountered stimuli.

The model’s explanation of the perceptual magnet effect is simple and straightforward: prototypical examples of a category are more difficult to discriminate from each other than non-prototypical examples because they have a smaller representation in auditory cortical maps (see also Bauer et al., 1996). The model further posits that this reduced cortical representation results from phoneme category learning during infancy. In particular, learning to treat sounds from a particular region of acoustic space as members of the same category leads to a decrease in the size of the auditory cortical representation of sounds near the center of that region.

4. Testing model predictions with auditory psychophysical experiments

The model attributes the perceptual magnet effect to neural map formation properties that are not unique to speech stimuli. This leads to the prediction that it should be possible to induce a perceptual magnet-like effect if

categorization training is carried out using non-speech stimuli. We have performed a psychophysical experiment to test this prediction (Guenther et al., 1999). In this experiment, subjects performed a category learning task in a 45-minute training session involving bandpass-filtered acoustic noise stimuli that varied in center frequency of the pass band. Each subject's ability to discriminate these sounds was estimated before and after training using a d' measure. The stimuli were not perceived as speech-like by experimental subjects, and they were not perceived "categorically" prior to training.

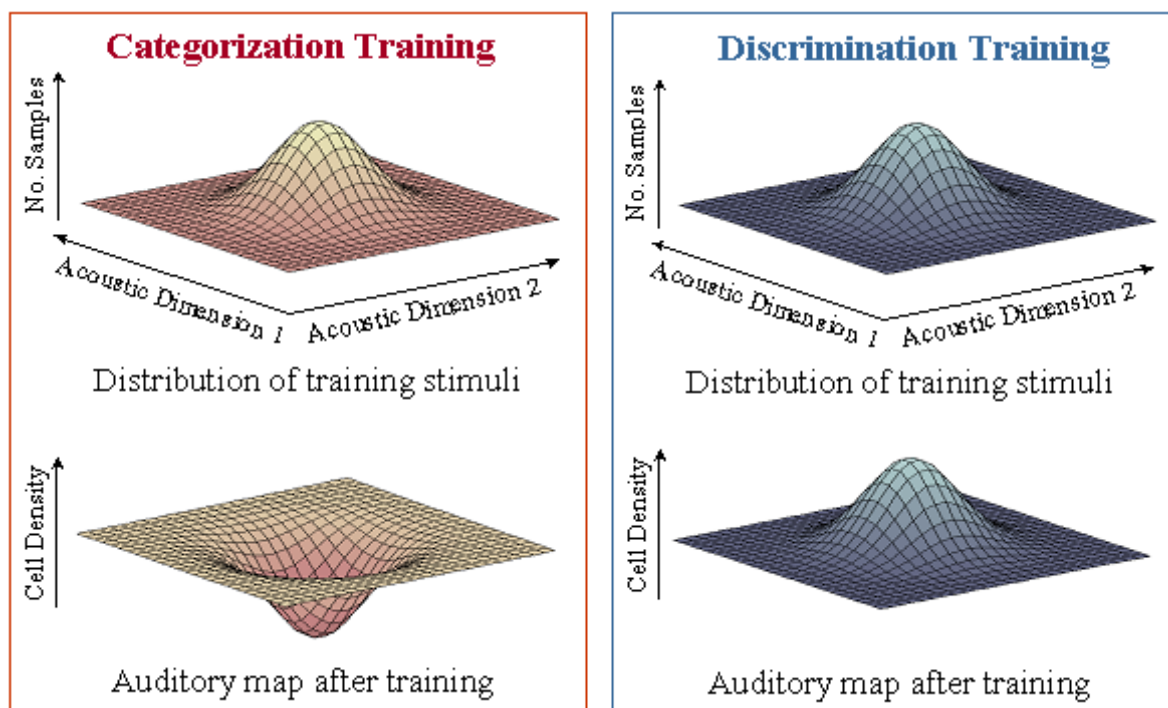


Figure 2. Hypothesized changes in the neural map in auditory cortex as a result of categorization training (left) and discrimination training (right). The x and y axes of all plots correspond to two acoustic or auditory dimensions, such as the first two formant frequencies of a vowel sound. The z axis corresponds to the number of training stimuli from that region of frequency space (top plots) or the number of cells in the map devoted to each region of frequency space (bottom plots). According to the model, 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.

Seven stimuli were generated for each of two regions of frequency space, referred to as the control region and the training region. The ability to discriminate the center stimulus from each range from the other stimuli in the range was measured before and after training. During each trial of the training task, the subject was presented with a sequence of two, three, or four stimuli. One of these stimuli was from the training region, and the rest were from other parts of frequency space. The subject had to choose which of the stimuli in the sequence belonged to the training region, or "category". Sounds from the control region were not heard during training. Subjects generally got significantly better at the task over the roughly 45-minute training session.

As shown in the top half of Figure 3, categorization training led to a decreased ability to discriminate stimuli from the region of frequency space corresponding to the center of the training category, as in the perceptual magnet effect for vowels described by Kuhl and colleagues. 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 from within the training region as compared to their ability to discriminate stimuli from within the control region. This verifies the model's prediction that a perceptual magnet effect can be induced using non-speech stimuli in a categorization training task.

In a second experiment, the same stimuli from the categorization training experiment were used in a discrimination training task rather than a categorization training task. On each trial, the subject was presented

with two stimuli from the training region and was asked to report if the stimuli were “same” or “different”. The results of this experiment are shown in the bottom half of Figure 3. Whereas categorization training led to a decrease in the discriminability of the training stimuli, discrimination training with the same set of stimuli led to increased discriminability. This indicates that it is the nature of the training task, and not just the distribution of the training stimuli, that leads to the perceptual magnet effect seen in the first experiment.

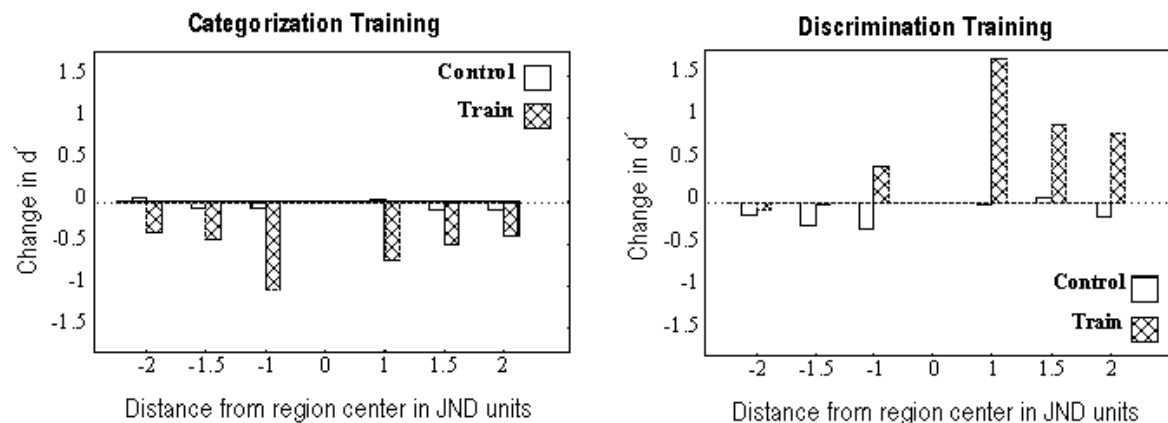


Figure 3. Psychophysical results indicating opposite effects on discriminability (as measured by d') for categorization training and discrimination training. Effects of categorization training on discriminability of stimuli within a newly formed category are shown in the left panel. Effects of discrimination training on discriminability are shown in the right panel. Categorization training leads to a decrease in discriminability for stimuli from the center of the training region, whereas discrimination training leads to an increase in discriminability for stimuli from the center of the training region. JND=Just Noticeable Difference (measured in Mels prior to training).

The model described above was simulated and compared to the experimental results of Guenther et al. (1999). In this simulation, the first layer conveys the center frequencies of the narrow band noise stimuli that were presented to subjects in a psychophysical experiment. Stimuli were presented at various distances (represented in units of the just noticeable difference (JND) as measured at the beginning of the experiment) from the prototype (i.e., the center of the training region) according to the experimental protocol while the model learned the auditory map.

According to Zohary (1992), the standard deviation of noise in a population code response is inversely proportional to the square root of the number of cells (N) used in the population response. In the simulations herein, we assume that all cells tuned to stimuli within a fixed distance r from the input stimulus contribute to the population code response for that stimulus. Because there are fewer cells representing stimuli from the region of frequency space near the category prototype, the neural representation of stimuli located near the prototype is more susceptible to noise than the neural representation of stimuli located farther from the category center.

In a testing phase following training, pairs of stimuli were presented to the model as in the psychophysical experiment. Gaussian noise was added as described by Zohary (1992). The noisy population responses for both stimuli were calculated, and if they were within some sensitivity measure, ϵ , the stimuli were judged to be the same. Otherwise, the stimuli were judged different. Resulting d' scores were calculated and compared with scores obtained using a map with a uniform distribution of cells in the frequency space (corresponding to the pre-training d' test in the psychophysical experiment). The change in d' was then fit to experimental results using a linear least squares fit (see Figure 4). The resulting r^2 value was 0.9061; the model, therefore, accounted for over 90% of the variance in the experimental results.

5. Testing model predictions with brain imaging experiments

According to our model, the perceptual magnet effect arises because there are fewer cells in the auditory cortical map representing prototypical examples of a phoneme category as compared to non-prototypical examples (i.e., those far away from category prototypes). The model further predicts that the reduced cortical representation for category prototypes arises due to exposure to sounds from an infant's native language. These aspects of the model are illustrated in Figure 5 for the region of formant frequency space corresponding to the American English phonemes /r/ and /l/ and the Japanese phoneme /r/. The left half of Figure 5 shows the results of a simulation in which the model was exposed to sounds drawn from a distribution approximating the distribution of /r/ and /l/ sounds in American English. The formant representation for this simulation consisted of the second and third formant frequencies (F2 and F3). The valleys in the density surface indicate that fewer cells are devoted to representing prototypical sounds (shown as diamonds in the F2-F3 plane), while sounds near category boundaries have larger representations. The right panel shows the results of a simulation in which the model was presented with stimuli chosen to approximate the distribution of Japanese /r/ sounds in formant frequency space. Here, there is a single dip in the map corresponding to the center of the Japanese /r/ category. This dip occurs at the approximate location of the boundary between American English /r/ and /l/. The model thus provides a partial account for why Japanese speakers who learn English as adults have difficulty discriminating English /r/ and /l/: they have relatively few auditory cortical cells tuned to the region of formant space corresponding to this boundary as a result of repeated exposure to the Japanese phoneme /r/.

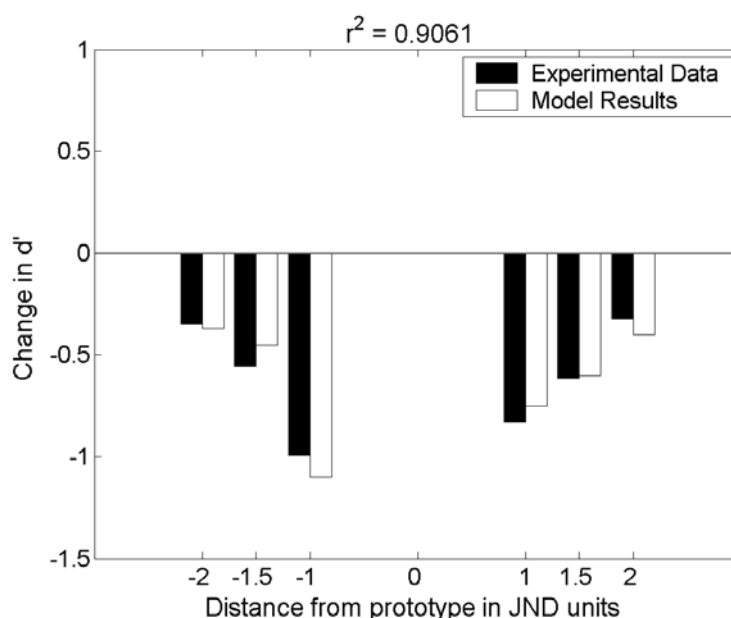


Figure 4: Comparison of simulation results with experimental results of Guenther et al (1999).

The model's prediction that the magnet effect arises because fewer cells in auditory cortex are responsive to prototypical examples of a phoneme was tested in a functional magnetic resonance imaging (fMRI) experiment while subjects listened to prototypical and non-prototypical examples of the American English vowel /i/ (Guenther, Nieto-Castanon, Tourville, and Ghosh, 2000). Subjects were nine right-handed native speakers of American English ages 18-55 with no history of language or other neurological disorders. The experimental protocol was approved by the Boston University committee on human subjects. Informed consent was obtained from all subjects. Subjects were stimulated binaurally with two synthetic vowel stimuli, a prototypical /i/ stimulus and a non-prototypical /i/ stimulus, presented in separate blocks. Stimuli were generated using the Sensyn speech synthesis software (Sensimetrics Corporation) with the following parameters: sampling frequency 8KHz, amplitude of voicing 60, and 4 formant frequencies (266Hz, 2294Hz, 3010Hz, 3300Hz for the prototypical stimulus, and 347Hz, 2095Hz, 3010Hz, 3300Hz for the non-prototypical stimulus, with bandwidths of 100Hz, 120Hz, 150Hz, and 300Hz respectively). These parameters were chosen to match synthetic vowels used to demonstrate the perceptual magnet effect psychophysically (Kuhl, 1991). Stimuli were presented in a block paradigm consisting of alternating 30-second blocks of prototypical vowels and non-prototypical vowels

separated by 30-second silent intervals for a total run length of 5-1/2 minutes. Subjects were told to attend to the stimuli by listening for differences from sound to sound.




Figure 5: Simulation results showing cell densities in an auditory map trained with American English /r/ and /l/ sounds (left) and Japanese /r/ sounds (right). The height of the surface represents the number of cells with tuning curve centers in a small region of formant space. Data averaged over 10 simulations and smoothed with a 4 x 4 spatial filter.

Data were obtained using a 1.5T General Electric Signa imager. Imaging sessions began with the acquisition of anatomical images that were later used to parcellate the regions of interest. T2-weighted functional images encompassing the entire peri-sylvian cortex were acquired using an asymmetric spin-echo echo-planar imaging sequence ($\tau=-25\text{ms}$, $TE=70\text{ms}$, $TR=2\text{s}$, matrix size 64×64 , 5mm thick contiguous slices with in-plane resolution= $3.1\times 3.1\text{mm}$). Individual functional runs were realigned (motion-corrected) using rigid body transformations to the first image in each scan, then coregistered with a structural T1 scan for each subject. Two runs were rejected for scanner data collection problems not detected during scanning. The remaining runs were visually inspected to meet noise and residual motion criteria, then tested for paradigm-correlated observed motion. Three runs showed excessive correlated motion and were thus removed from the analysis. Structural T1 images were parcellated individually for each subject to define 10 brain regions of interest (ROIs) on the basis of anatomical markers according to the procedure described by Caviness et al. (1996). The use of this parcellation procedure for each individual avoids the need for spatial averaging of the statistical parameter maps (and the subsequent loss of spatial resolution). The ROIs were ten peri-sylvian cortical areas, including areas known to become active during perceptual processing of auditory speech stimuli: Heschl's gyrus (HG), parietal operculum (PO), planum polare (PP), planum temporale (PT), anterior and posterior supramarginal gyrus (SGa, SGp), anterior and posterior superior temporal gyrus (T1a, T1p), and anterior and posterior middle temporal gyrus (T2a, T2p). HG, PT, and T1 are commonly considered to be auditory areas. PO, SG, and T2 are multimodal areas that become active during some speech or language tasks. Data reduction was applied to each ROI to obtain one temporal activation profile characterizing the response of all voxels within a given region. Significance of specific contrasts for each ROI activation profile were obtained using the general linear modeling (GLM) framework within the SPM statistical analysis package.

The main result of this experiment was that listening to prototypical examples of the vowel caused significantly less activation than listening to non-prototypical examples across the auditory cortical areas ($p<0.05$). Figure 6 shows the results for auditory cortical areas in the temporal lobes. Less activation is seen for the prototypical vowel than for the non-prototypical vowel throughout most of the auditory cortical areas. This is particularly evident in the auditory cortical areas in the right hemisphere, including HG (primary auditory cortex) and PT (a secondary auditory cortical region). This result supports the model's simple explanation for the perceptual magnet effect: prototypical vowels are more difficult to discriminate from each other than non-prototypical vowels because they have a smaller cortical representation, and smaller cortical representations are more susceptible to noise in the neural signals.

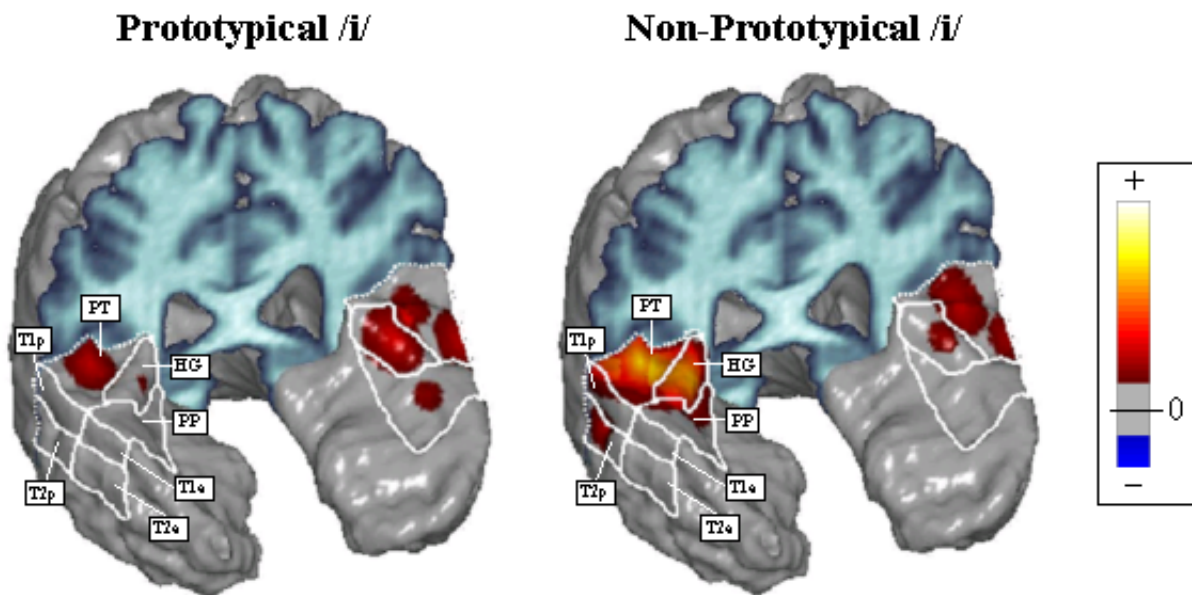


Figure 6. Temporal lobe activations as measured with fMRI while subjects listened to prototypical and non-prototypical examples of the vowel /i/. The frontal lobe has been removed to expose the temporal plane. As predicted by the neural model described in Section 3, less auditory cortical activation is seen while listening to prototypical /i/ sounds than while listening to non-prototypical /i/ sounds, particularly in the right hemisphere. HG = Heschl's gyrus; PT = planum temporale; PP = planum polare; T1a, T1p = anterior/posterior superior temporal gyrus; T2a, T2p = anterior/posterior middle temporal gyrus.

6. Conclusions

The experiments described in the current article were designed to investigate learned warpings of auditory perceptual space by testing a neural network model of the perceptual magnet effect. This model posits that phoneme category learning in infancy leads to the perceptual magnet effect because it causes a reduction in the size of the auditory cortical representation of prototypical examples of a vowel category. The model's assertion that general neural map formation properties are responsible for the effect implies that it should be possible to induce the effect in non-speech stimuli. This prediction was verified by a psychophysical experiment showing that subjects learning a new category for non-speech auditory stimuli get worse at discriminating central examples of the category from each other, and the model's fit to the data accounted for over 90% of the variance. An fMRI analysis revealed that listening to prototypical examples of the vowel /i/ leads to less activation in peri-sylvian cortical areas than listening to non-prototypical examples.

Taken together, these results strongly support the following assertions of the Guenther et al. (1999) neural model of auditory map formation:

- Categorization training leads to a relative decrease in the size of the cortical representation of prototypical examples of a category.
- Similarly, speech sound category training in infancy leads to a decrease in the size of the cortical representation of prototypical examples of some phonemes.
- This decreased representation is responsible for the perceptual magnet effect.

The model thus unifies observations from auditory psychophysics, cortical neurophysiology, and neural modeling to provide an account of how an infant learns the phoneme categories of his/her native language.

ACKNOWLEDGEMENTS

We gratefully thank Dr. Julie Goodman and the Massachusetts General Hospital NMR Center for the use of their imaging facilities. We also thank Dr. Fatima Husain, Alfonso Nieto-Castanon, Satrajit Ghosh, and Jason

Tourville for their contributions to the research described herein. Address correspondence to: Frank Guenther, Department of Cognitive and Neural Systems, Boston University, 677 Beacon Street, Boston MA 02215. Email: guenther@cns.bu.edu.

REFERENCES

- Bauer, H.-U., Der, R., and Herrmann, M. (1996). Controlling the magnification factor of self-organizing feature maps. *Neural Computation*, **8**, 757-771.
- Caplan, D., Gow, D., & Makris, N. (1995). Analysis of lesions by MRI in stroke patients with acoustic-phonetic processing deficits. *Neurology*, **45**, 293-298.
- Caviness, V.S. Jr., Meyer, J., Makris, N., & Kennedy, D.N. (1996). MRI-based topographic parcellation of human neocortex: An anatomically specified method with estimate of reliability. *J. Cog. Neurosci.*, **8**, 566-587.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, **23**, 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**, 1111-1121.
- Guenther, F.H., Husain, F.T., Cohen, M.A., and Shinn-Cunningham, B.G. (1999). Effects of categorization and discrimination training on auditory perceptual space. *Journal of the Acoustical Society of America*, **106**, 2900-2912.
- Guenther, F.H., Nieto-Castanon, A., Tourville, J.A., and Ghosh, S.S. (2000). The representation of prototypical and non-prototypical vowels in peri-sylvian cortical areas. *Society for Neuroscience Abstracts*, vol. 26, part 2, 1971.
- Jenkins, W.M., Merzenich, M.M., Ochs, M.T., Allard, T., and Guic-Robles, E. (1990). Functional reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally controlled tactile stimulation. *Journal of Neurophysiology*, **63**, 82-104.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, **43**, 59-69.
- 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**, 93-107.
- 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**, 606-608.
- 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**, 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**, 87-103.
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striata cortex. *Kybernetik*, **14**, 85-100.
- Zohary, E. (1992). Population coding of visual stimuli by cortical neurons tuned to more than one dimension. *Biological Cybernetics*, **66**, 265-272.



Frank Guenther is an Associate Professor of Cognitive and Neural Systems at Boston University in Boston, Massachusetts. He is also a Research Affiliate in the Speech Communication Group at the Massachusetts Institute of Technology, and a faculty member in the Harvard University-MIT Speech and Hearing Bioscience and Technology Program. He received a masters degree in engineering from Princeton University in 1987, and a PhD in computational neuroscience from Boston University in 1992. Dr. Guenther's research program involves neural modeling of speech perception and production in normal and disordered adults. These models are tested and refined using a range of experimental techniques, including functional imaging of brain function (fMRI and MEG), electromagnetic articulometry, and auditory psychophysics.



Jason Bohland is a second-year PhD student in the Department of Cognitive and Neural Systems at Boston University. His primary research focuses are neural modeling of speech perception and speech production, and the use of functional imaging methods in modeling. He received a masters degree in Electrical Engineering from the University of Cincinnati in 2000, specializing in intelligent