Speech, Hearing and Language: work in progress

Volume 11

The perceptual magnet effect is not specific to speech prototypes: new evidence from music categories.

Sarah BARRETT

Department of Phonetics and Linguistics
UNIVERSITY COLLEGE LONDON

The perceptual magnet effect is not specific to speech prototypes: new evidence from music categories.

Sarah BARRETT

Abstract
Previous work on prototypicality in music has led to the claim that music prototypes act in the opposite way to speech prototypes – as anchors rather than magnets. In one such study, professional musicians were given discrimination tasks in which they had to distinguish acoustically-similar sounds in the context of both a prototypical and non-prototypical C-major chord (Acker et al., 1995). In contrast to what has been found for American English listeners for various speech-sound categories (e.g. Kuhl, 1991), professional musicians show enhanced discrimination in the region of the prototype. The present study questions whether the performance of such musically-trained subjects is representative of the average listener's perception of music categories. Here, 10 non-musicians as well as 10 musicians were given a discrimination task in which they were required to distinguish prototypical and non-prototypical C-major chords from a series of acoustically-similar variants. Unlike the musicians who showed enhanced discrimination in the context of the prototype, the non-musicians showed reduced discrimination. The results have implications for the applicability of the perceptual magnet effect to domains other than speech and are interpreted in terms of a new theory called A&R theory, which suggests that prototypes have a dual role in the perceptual system depending upon the amount of attention paid to them by the listener.

Background

It has been claimed that the best exemplars or 'prototypes' of phonetic categories exert an attractor effect on surrounding sounds in the same speech category, making it difficult to tell the difference between the prototype and acoustically-similar sounds (e.g. Iverson & Kuhl, 1995; Kuhl 1991). This has been called the perceptual magnet effect (Kuhl, 1992) and it has been used to account for one of the most fundamental problems in speech perception: mapping the relationship between the acoustic variability in the speech signal and the stable linguistic categories that listeners report that they effortlessly hear. The assumption is made that incoming speech-sounds are identified by reference to the category prototype acting as a perceptual magnet because they are completely or relatively indistinguishable from it. Poorer examples of the same speech-sound category, non-prototypes, do not function as perceptual magnets. The perceptual magnet effect has been confirmed in discrimination experiments using both adults and infants: when both a prototype and a non-prototype of the same speech category are compared with a number of similar-sounding variants of that category, listeners are significantly worse at distinguishing the prototype from its variants than the non-prototype (e.g. Grieser & Kuhl, 1989, for babies; Iverson & Kuhl, 1995, for adults). This 'contraction' of the perceptual space around the category prototype has been assumed to be specific both to humans (e.g. Kuhl, 1991) and to the speech modality (e.g. Acker, Pastore, & Hall, 1995). Evidence for this latter assumption comes from work on music.

Work on music

Acker et al. (1995) replicated Kuhl's (1991) perceptual magnet experiment using sets of C-major chords. In the first portion of this experiment, five musically-experienced subjects rated the goodness of a number of different examples of C-major chords. They rated the examples for how 'in-tune' they sounded. Each subject's highest rated exemplar and a poorly rated exemplar served as their prototype (P) and non-prototype (NP) respectively in a subsequent discrimination task. In setting up this discrimination task, thirty 'comparison' chords were generated around each of the P and NP. These 'comparison' chords were created by holding the C frequency constant in the P/NP and then varying the E and G frequencies in both sharp (i.e. increasing) and flat (i.e. decreasing) directions. Similar to Kuhl's experiment where vowel stimuli varied along either one (F1 or F2) or two dimensions (F1 and F2), musical stimuli in Acker et al.'s experiment varied along either the E or G dimension alone, or along both dimensions simultaneously. For the actual discrimination task, the P and NP were paired both with themselves and with each of their thirty comparison chords. Listeners were given a two-interval forced-choice task in which they heard two pairs of chords, one of which contained two identical chords (the P or NP paired with itself) and one of which contained two different chords (the P or NP paired with one of their 30 comparison chords). Subjects had to decide which of the two pairs was the 'different' pair.

Unlike in the discrimination tests with speech, results from music categories showed that discrimination was significantly better in the context of the musician's prototypical C-major chord than in the context of their non-prototypical C-major chord. In other words, the musicians could tell more sounds apart from their prototype than from their non-prototype. On the basis of these findings, the suggestion was made that music prototypes function in the opposite way to speech prototypes - as anchors, rather than as magnets.

Interpreting the Findings

Although the results of Acker et al.'s experiment suggest that enhanced discrimination occurs in the region of a representative C-major chord, the generality of these results is questionable. Their listeners were all musically-experienced. Each had received a minimum of 10 years of training in Western music (Acker et al., 1995, p. 866). Although Acker et al.'s rationale for selecting such a subject pool was a reasonable one[1], it is not inconceivable that their subjects were more successful at discriminating in the context of the prototypical C major chord because their professional training had included extensive ear work involving the recognition of in-tune and out-of-tune musical chords. In other words, it is unclear whether the enhanced discrimination found near a music prototype in the Acker et al. experiment is due more to the nature of the subjects they tested than to the nature of the prototype under investigation. Given that this experiment is the only explicit test of the perceptual magnet effect in music to date, no satisfactory resolution of this issue can be found by recourse to the literature.

[1] One of the things that Acker et al. (1995) were looking at was the difference between an exemplar and a prototype model, the former of which necessitates previous experience with the stimuli in question. For these reasons, they used musically-experienced listeners.

Further investigation is therefore warranted. The claim that a perceptual magnet effect is not found for musical categories could be strengthened by the finding that non-musicians (who have received no formal ear-training to recognise in-tune and out-of-tune musical chords) also show enhanced discrimination near a music prototype. On the other hand, if non-musicians were to show evidence for reduced discrimination near a music prototype (characteristic of the perceptual magnet effect identified for speech) then this would have important implications for the applicability of the perceptual magnet effect to a general auditory mode of processing.

The present investigation

Ten musically-experienced listeners (here referred to as 'musicians') and ten listeners with very limited musical experience (here referred to as 'non-musicians') participated in the present investigation. There are two experiments reported in this paper. The first experiment involved an adaptation of Acker et al.'s goodness rating test and was used to isolate a less representative (i.e. non-prototypical) exemplar of a C-major chord for each of the musicians. Only the musicians participated in this experiment; the non-prototypes for the non-musicians were selected by taking the average of the musicians' responses[2] in this goodness rating task. Furthermore, only the non-prototypes were personalised. An approximation of the equal-tempered C major chord (261.6 Hz [C]; 329.6 Hz [E]; 392 Hz [G]) was used as the prototype for both musicians and non-musicians (at 261.6 Hz; 329 Hz; 392 Hz). This was close to the average frequencies selected for the prototypical stimulus in Acker et al's experiment. Because all subjects would be required to discriminate around the same prototype, this meant that there was some sort of commonality between the musicians and non-musicians.

Both musicians and non-musicians participated in the second experiment. This involved an adaptation of Acker et al.'s discrimination task using an AX same-different design in which the prototypical and non-prototypical musical chords selected in Experiment 1 were paired both with themselves and with a number of similar-sounding comparison chords that were either sharper or flatter than the P and NP.

Experiment 1

As mentioned in the previous section, the goal of this experiment was to isolate each individual musician's less representative example of a C major chord to be used as the non-prototype in a subsequent discrimination task. In the same way that speakers of a language may have different pronunciation preferences for various speech sounds[3],

---

[2] It was not necessary for non-musicians to take part in Experiment 1 since a pilot experiment by the present author suggested that they were inconsistent when asked to locate a clear boundary on the C major-minor continuum. 10 non-musicians were given a categorical perception task using a C-major-C-minor continuum. They did not show evidence of two clear categories, and were fairly inconsistent in their responses. Although Howard et al. (1992) suggest that non-musicians have some awareness of the concepts 'major' and 'minor' and are able to relate them to other terms such as 'happy' or 'sad', it is obvious that they do not have two clearly-defined categories

[3] In other work by the present author, Southern British English speakers selected different tokens of /u:/, /lu:/ and /ju:/ as their prototypes and non-prototypes in a goodness rating task (Barrett , 1997).

musicians may have different preferences for where they place the boundary between a major and minor chord, and consequently, for what they consider to be in-tune and out-of-tune. This was a concern raised by Acker et al. who also gave their subjects a goodness rating task.

In selecting the NP for each listener, a modification of Acker et al.'s goodness rating task was used whereby an NP is selected based on a distance metric from the musician's major-minor category boundary.

Subjects

A group of 10 subjects (5 male, 5 female) participated in this experiment. Individuals in this group had all received some form of musical training, although the amount of training varied across subjects. Subjects ranged from 20 to 26 years and they were all either undergraduate or graduate students at the University of Cambridge. They were each paid £2 for their participation in this 20 minute experiment.

Stimuli

The stimuli were a series of synthesised 3-note-chords that went from a perfectly in-tune C-major to a perfectly in-tune C-minor. Each chord was composed of 3 simultaneously presented square waves of different fundamental frequencies. The frequencies for the C-major were adapted from the textbook description of an in-tune C-major chord with the highest note at 392 Hz (G), the middle note at 329.5 Hz (E)[4] and the lowest note at 261.6 Hz (C) [Burns & Ward (1982)]. Similarly, the frequencies for the C-minor were taken from the textbook description of an in-tune C-minor chord. This has the same frequencies for the highest and lowest notes as the C-major chord, but the middle note is at 311.6 Hz (E♭). To create the continuum between C-major and C-minor, the middle note was varied in 1.5 Hz steps. This created a series of 12 steps which are listed in Table 1.

392	392	392	392	392	392	392	392	392	392	392	392
329	327.5	326	324.5	323	321.5	320	318.5	317	315.5	313	311.5
261.6	261.6	261.6	261.6	261.6	261.6	261.6	261.6	261.6	261.6	261.6	261.6

Table 1. *Chords in the major/minor continuum at distances of 1.5 Hz*

The square waves (and subsequent chords) were generated using interactive signal processing commands in ESPS (developed by the Entropic Research Laboratory) running on a Silicon Graphics computer. First, each wave was generated using a program called 'testsd' with the default sampling rate of 8000 Hz. Each wave was one second in duration. Once the individual square waves had been created, they were combined into chords using a program called 'addsd'. The duration of each resultant 3-note chord was also one second.

[4] For this task, 329 Hz was selected instead of 329.6 Hz for the E frequency to allow for a wider category in the discrimination task. Pilot tests with a group of musicians suggested that a C-major chord with a middle frequency of 329 Hz was still considered a prototypical example of the C-major category.

Procedure

The C-major category boundary was measured by locating each individual musician's 50% cross-over point. This is the frequency point at which musicians classify a chord as major or minor with equal probability.

To determine the location of their 50% cross-over point, subjects had to identify each of the twelve chords as either 'major' or 'minor'. They first heard the sequence of twelve chords in descending order of frequency and were required to make their identifications out loud. The experimenter noted each response on a sheet of paper. Subjects then heard the sequence of 12 chords in the reverse order and were again required to make their responses out loud. Once more, the experimenter noted each response on a sheet of paper.

Results and Discussion

Individual musicians differed as to the location of their 50% crossover point. The range was between 320-322.5 Hz.

The non-prototype C major for each individual musician (to be used in Experiment 2) was set at 5 Hz greater than each subject's 50% crossover point; this non-prototypical value was felt to be well within each subject's C major category, given the range possible. Although this may not have been each subject’s ‘least’ preferred C major chord, Kuhl has been criticised for using a subject’s least preferred exemplar as the non-prototype because it is often too close to the category boundary to ensure that all comparison exemplars in the subsequent discrimination task are from the same category (e.g. Sussman & Lauckner-Moreno, 1995).

Experiment 2

The goal of Experiment 2 was to confirm Acker et al.’s claim that musicians perform better in a discrimination task in the context of a prototypical chord than in the context of a non-prototypical chord. It also investigated whether non-musicians, who have received no formal training in music, would show evidence for this enhanced discrimination around the C-major prototype. The discrimination task was similar in procedure to that used by Acker et al. except that here an AX design was used and the chords were only varied along one dimension (the E frequency) rather than along two dimensions (the E & G frequencies)[5].

Subjects

20 subjects participated in this experiment. Ten of these subjects were the musicians used in Experiment 1. The remaining ten subjects (4 males, 6 females) made up the non-musicians group. None of these individuals had received any formal instruction in music above what they might have received during their school years. Members of this group ranged from 21 to 25 years and they were all graduate students at the University of Cambridge. All subjects (both musicians and non-musicians) were paid £3 for their participation in this 45-minute experiment.

---

[5] Kuhl & Iverson (1995) were successful in demonstrating a perceptual magnet effect for prototypical tokens of /i/ using stimuli that varied only along one dimension.

Stimuli
The non-prototype for the non-musicians was set at 326.5 Hz. This was determined by taking the average of the musicians' cross-over points which was 321.5 Hz, and increasing this by 5 Hz.

For each individual musician, and for the non-musicians as a group, 6 comparison chords were synthesised in equal steps of 1.5 Hz[6] from both the prototype and non-prototype in both sharp and flat directions. Consequently, for three of the comparison chords, the middle note was either 1.5 Hz, 3 Hz, or 4.5 Hz greater in frequency than the middle note of the P or NP. For the other three comparison chords, the middle note was either 1.5 Hz, 3 Hz or 4.5 Hz smaller in frequency than the middle note of the P or NP. The small step size of 1.5 Hz was chosen due to the narrowness of the C major category for most subjects.

Procedure
Subjects were comfortably seated in a sound-treated room at the University of Cambridge. The chords were played out through a Silicon Graphics computer and subjects listened binaurally over Sennheiser HD 520II headphones with an impedance of 300 ohms.

Subjects completed a discrimination task for both the P and NP, with the order of presentation counterbalanced across subjects. They had a short break between P and NP tasks.

For each discrimination task, subjects heard a randomised series of pairs of chords. There were 60 ‘same’ trials where the P (or NP) was paired with itself and 60 ‘different’ trials where the P (or NP) was paired with one of the 6 comparison chords. On half of the ‘different’ trials the P (or NP) occurred first in the pair, and on the other half, it occurred second. There was a 1 second ISI between members of a pair and a 2 second ITI between successive pairs. The last 500 ms of this ITI was a female voice saying the word 'next'. The number of each trial was signalled on the computer screen immediately prior to presentation. Subjects judged whether the two chords were the 'same' or 'different' and made their response by clicking the mouse on one of two panels on a button box on the computer screen. The panels bore the labels 'same' and 'different'. No feedback was given.

There were 5 practice trials before the main experiment began and subjects were given immediate visual feedback, 'SAME'/'DIFFERENT' (as appropriate) on these practice trials, but responses were not included in the analysis.

Results
The responses were scored within the framework of signal detection theory [Macmillan & Creelman (1991)]. As shown in Table 2, d-primes for individual subjects were calculated for both the P and NP experimental blocks. For both musicians and non-musicians, a separate one-tailed t-test for paired observations was

---

[6] At such a low frequency (approximately 300 Hz) Hertz and Mels are roughly equivalent. For this reason, the step sizes were left in Hertz and not converted into Mels, as had been done in previous perceptual magnet work on speech [e.g. Kuhl , 1991; Acker et al., 1995, did not convert from Hertz into Mels in the discrimination task.

conducted on the data. The musicians in Table 2 have been divided into two groups (see later text for explanation).

For all 10 of the non-musicians, discrimination was significantly worse around the prototype than around the non-prototype [t (9) = 4.67, p < .01]. This is the opposite to what was found for the musicians in Acker et al.'s experiment.

When the data for the 10 musicians were analysed, there was no significant difference between discrimination around the P and discrimination around the NP. Two trends do however emerge in the musicians' data which might account for this non-significant finding. For 6 of the 10 musicians (labelled as Group 1 in Table 2), discrimination appears to be worse around the non-prototype than around the prototype, and this difference was found to be significant [t (5) = 3.98, p < .01]. For the remaining 4 musicians (labelled as Group 2 in Table 3), the opposite occurred. Discrimination was significantly worse around the prototype than around the non-prototype, and this difference was also significant [t (3) = 4.67, p < .01]. Furthermore, although the non-musicians and Group 2 musicians both showed evidence for a perceptual magnet effect (i.e. reduced discrimination around the P), average discrimination around the prototype was significantly worse for the non-musicians [t (13) = 9.65, p < .005].

Musicians

		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Group 1:	**P**	3.45	2.68	3.18	1.83	2.54	2.56				
	NP	2.92	2.65	2.87	1.50	1.87	2.28				
Group 2:	**P**							1.59	2.64	2.01	2.36
	NP							1.94	3.21	2.39	3.21

Non-musicians

		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
	P	1.63	1.76	1.79	0.81	2.01	1.54	2.72	1.55	1.52	1.81
	NP	2.06	1.77	2.39	1.29	2.37	1.67	2.39	2.51	2.48	2.02

Table 2: *Individual d-primes for musicians (listed in terms of Group 1 and Group 2) and non-musicians for both the P and NP. Mean discriminations are as follows: For the non-musicians, P = 1.71, NP = 2.19; for the Group 1 musicians, P = 2.71, NP = 2.35; for the Group 2 musicians, P = 2.15.*

Discussion

These results show that for more than half of the musicians (Group 1), a C major prototype acted like a perceptual anchor. In other words, discrimination was enhanced in the region of the prototype so that it was easily discriminated from similar-sounding versions of a C-major chord. This is identical to the findings of Acker et al.

The discrimination performance of the Group 2 musicians and all of the non-musicians is more surprising. These subjects showed evidence of reduced discrimination around the C major prototype. They were significantly worse and distinguishing the prototype from its similar-sounding variants than the non-prototype. In other words, prototypical C major was acting like a perceptual magnet. The consequences that this finding has for the uniqueness of the perceptual magnet effect are addressed in the General Discussion.

The discrepancy in discrimination performance amongst the musicians is difficult to interpret. A post-hoc examination of the subject histories, however, revealed that the 6 musicians who showed enhanced discrimination around the music prototype (Group 1) had much more musical experience than the remaining 4 who showed reduced discrimination (Group 2). These six Group 1 subjects had all had at least as much musical training as the experienced musicians in Acker et al.'s study [i.e. a minimum of 10 years of experience with Western Music - see Acker et al., 1995, p. 866)] and were all actively pursuing a musical career at the time of testing. The remaining 4 musicians had had some previous experience with music but were not as actively engaged in musical-related activities at the time of testing. A similar distinction in musicality was made in a cross-cultural study by Lynch, Eilers, Olller, & Urbano (1990). In Lynch et al.’s comparison of adults' and infants' responses on within-category discrimination of Western and Javanese scales, subjects were divided into three categories based on their amount of musical training: non-musician, amateur musician, and professional musician. The professional and amateur musicians had approximately as much experience as the Group 1 and 2 musicians in Experiment 2 respectively. Although the main purpose of Lynch et al.'s study was not to test for a perceptual magnet effect, their data from subjects' discrimination of a 2 Hz change in frequency in the Western scales do lend support to the interpretation of the findings from Experiment 2: in this study, discrimination of this 2 Hz change was below chance for both the non-musicians and the amateur musicians, but above chance for the professional musicians. This suggests that within the category of 'musician', the amount of individual training in music can significantly affect the performance of the individual subject on certain discrimination tasks. This suggestion is also confirmed by Howard, Rosen, & Broad’s (1992) categorical perception study where the performance of individual subjects on an identification/discrimination paradigm was directly related to their degree of musicality.

Given the points raised in this discussion, for the remainder of this paper, the Group 1 musicians who showed enhanced discrimination in the region of the C-major prototype will be referred to as 'professional musicians' and the Group 2 musicians who showed reduced discrimination in the region of the C major prototype will be referred to as 'amateur musicians'.

General Discussion

It has been claimed in the literature that music prototypes act as perceptual anchors: there is enhanced discrimination near these prototypes making them easily discriminable from acoustically-similar variants of the same musical category. In contrast, speech prototypes act as perceptual magnets: there is reduced discrimination near these prototypes so that they are not easily discriminable from acoustically-similar variants.

The results from Experiment 2 question this link between the type of discrimination (i.e. reduced or enhanced) and the type of prototype (i.e. speech or music). The data suggest that whether or not there is reduced or enhanced discrimination in the region of a prototype does not depend on the nature of the prototype per se but rather on the nature of the listener. Professional musicians, for example, treat a prototypical C major chord as a perceptual anchor. In contrast, non-musicians and amateur musicians treat exactly the same prototype as a perceptual magnet. Before tackling this issue of prototype duality, it is necessary to pause and analyse the terminology that has been used in the prototype literature.

Examining the term 'perceptual magnet'

The distinction has been made throughout the prototype literature between the perceptual magnet which 'attracts' similar sounding stimuli (through reduced discrimination in the perceptual space) and the perceptual anchor which 'repels' similar-sounding stimuli (through enhanced discrimination in the perceptual space). It could be argued, however, that it is misleading to use the term 'magnet' to describe a prototype that only 'attracts' similar-sounding stimuli, given that in real life [especially noticeable in the physical sciences], magnets can both attract and repel.

In order to recognise this link with the physical sciences and to reinforce the dual nature of the prototype (interpreted from the results of Experiment 2), it would be more appropriate to abandon the term 'perceptual anchor' at this point and divide the term 'perceptual magnet' into two separate concepts: 'perceptual attractor' and 'perceptual repellor'.

The term 'perceptual attractor' is appropriate for the type of prototype (originally identified by Kuhl for speech-sounds) which is hard to distinguish from its similar-sounding variants, and exhibits a reduced discrimination in the surrounding perceptual space. The non-musicians and amateur musicians in Experiment 2 treated prototypical C major as a perceptual attractor. Similarly, Kuhl’s American English Speakers treat prototypes of /i:/ as perceptual attractors (e.g. Kuhl, 1991).

In contrast, the term 'perceptual repellor' is more appropriate for the type of prototype (originally identified by Acker et al. for music) which is easy to distinguish from its similar-sounding variants and which shows enhanced discrimination in the surrounding perceptual space. The professional musicians in Experiment 2 treated prototypical C major as a perceptual repellor.

This distinction between the perceptual attractor and the perceptual repellor preserves Kuhl's magnet analogy yet emphasises the prototype's adaptive significance: within the same magnet there are potential forces of both attraction and repulsion that are activated according to the needs of the listener.

Explaining how a prototype can be both an attractor and a repellor

One way to explain the discrepancy in discrimination performance exhibited by the subjects in Experiment 2 is to consider what purpose each type of discrimination serves for the listener. In other words, it is possible that the same prototype is meant to have two (albeit opposing) roles in the perceptual system - as either a perceptual attractor or a perceptual repellor - and that the nature of the prototype's role at a given point in time depends entirely upon the needs of the listener and, consequently, on the amount of attention that must be paid to the prototype by this listener.

Although it is unique here in its application to prototype theory, the concept of selective attention is not new in the speech perception literature. Many researchers have discussed how selective attention to a certain dimension causes a stretching in the perceptual space underlying that dimension (e.g. Jusczyk, Pisoni, & Mullenix, 1992; Nosofsky 1984, 1986; Pisoni, Aslin, Perey & Hennessy, 1982). Jusczyk et al., for example, demonstrate empirically how attention affects the discrimination of young infants. They found that when new-born babies were habituated on syllables that were perceptually very similar such as /pa:/, /ta:/, /ka:/ (called 'fine-grained distinctions' by Jusczyk et al., 1992, p. 254), they were easily able to detect the addition of a new syllable to the set no matter how similar or dissimilar it was to the pre-exposed stimuli. However, when infants were habituated on stimuli that were perceptually very dissimilar, such as /bi:/, /ba:/, /bu:/ (called 'coarse-grained distinctions' by Jusczyk et al., 1990, p. 254), they were unable to detect the addition of a new syllable that was perceptually very similar to their habituation set. Juszcyk et al. suggest that attention to a coarse-grained distinction (through habituation) resulted in perceptual stretching along items aligned on this dimension. Because the focus was taken away from a fine-grained distinction there was perceptual shrinking along items aligned on this dimension making them less distinguishable.

It is possible that this same concept is true of the prototype. If a lot of attention is paid to the prototype, there will be a stretching (i.e. enhanced discrimination) in the perceptual space and the prototype will be easily discriminated from similar-sounding variants of the same category. In contrast, if prototypical sounds are heard often, but not much attention is paid to the detailed properties of the prototype itself, there will be a shrinking (i.e. reduced discrimination) in the perceptual space and the prototype will not be easily discriminated from similar-sounding variants.

This point can be made clearer with an example: it is obvious that a professional musician needs to be able to make fine discriminations around a music prototype: if one instrument is out of tune, it can throw off an entire orchestra (cf. Acker et al., 1995, p. 871). It is not particularly surprising therefore that professional musicians pay a lot of attention to music prototypes, treating them as perceptual repellors and finding them very easy to distinguish from similar-sounding variants.

For the amateur musician and the non-musician, in contrast, music prototypes have a completely different role. For these individuals, the primary purpose of listening to music is to be entertained. They do not need to pay a lot of attention to the detailed properties of the prototype itself. In other words, they do not need to recognise whether or not a musical chord is exactly ‘in tune’. For amateur musicians and non-musicians therefore, music prototypes act primarily as magnets and show reduced discrimination in the perceptual space. The finding that the average discrimination around the C major prototype is significantly better for amateur musicians than for non-musicians reflects the fact that these amateur musicians have had some formal instruction in music (e.g. in ensemble playing where they are required to tune up and keep in tune) although not as much as the professional musicians.

This duality in the role of the prototype is also paralleled in speech. Under normal circumstances, the fine tuning (i.e. focused attention) of the professional musicians is not necessary in listening to speech. The initial goal of speech perception is to recognize; interpretation occurs at a later stage. For recognition purposes, it does not matter how qualitatively correct the incoming speech sound is as long as it is sufficiently good to activate a prototype in memory. Speech sounds can be successfully identified even if they aren't precisely 'in tune', and thus in general, phonetic prototypes act as perceptual magnets. A poor (but recognisable) example of an /i:/ is still considered to be a member of the category /i:/, and most people need be no more precise than this.

Although existing research suggests that speech prototypes function as perceptual magnets (e.g. Kuhl, 1991), it is plausible to argue that there are times when the speech prototype could adopt the alternative role of perceptual repellor and thus become easy to discriminate from a number of similar-sounding variants. The best experimental evidence of this to date is babies' abilities to discriminate nearly every contrast on which they have been tested regardless of whether it is native to their language system or not (e.g. Werker & Tees, 1984) and in the success with which adults are trained to discriminate non-native contrasts (e.g. Pisoni et al., 1982). Although it is less clear how the babies' discrimination is directly related to attentional factors, it is obvious in the latter case that adults' attention is focused due to the training tasks, making it easier for them to discriminate around a prototype.

In conclusion, prototypes have an adaptive significance. If the system needs a prototype to be easily distinguished (as in the case of a professional musician), it will adapt accordingly. This concept of adaptive significance parallels Lindblom's (1990) H&H Theory of speech production. The main principle of this theory is that during the course of speech production, the talker makes a running estimate of the listener's need for signal information and then adapts his/her articulations to suit this need. These adaptations occur along a continuum with the more forcefully articulated (and less coarticulated) 'hyper' forms at one end and the less forceful (and more coarticulated and reduced) 'hypo' forms at the other end. The main thrust of Lindblom's theory seems to be that the speaker does not expend more energy in speech production than is absolutely necessary.

This same principle can be borrowed to explain the perception of prototypes: the listener does not need to expend more energy in focusing their attention in speech perception than is absolutely necessary. Given that it is not necessary for a non-musician to recognise how precisely in-tune a C major chord actually is, for example, the music prototype adapts to suit the needs of this individual and shows reduced discrimination characteristic of a perceptual magnet. In other words, no energy is wasted by the listener in trying to make fine discriminations around the prototype since these fine discriminations are not essential.

The next section outlines a theory in which this principle of selective attention has been incorporated into a theory of perception based on the concept of the dual role of the prototype as either a perceptual attractor or a perceptual repellor.

Introducing A&R Theory

Let us call this perceptual theory 'A&R Theory' ('A' for 'perceptual attractor'; 'R' for 'perceptual repellor'). Having suggested in the previous section that the role of the prototype is directly related to the amount of attention paid to it by the listener, let us assume that the perception of prototypes, or, more accurately, the perception of the perceptual space around the prototype, is organised along a continuum where the position of the prototype at any given time along this continuum is directly determined by the amount of attention paid to it by the listener. A&R Theory is illustrated in Figure 1:

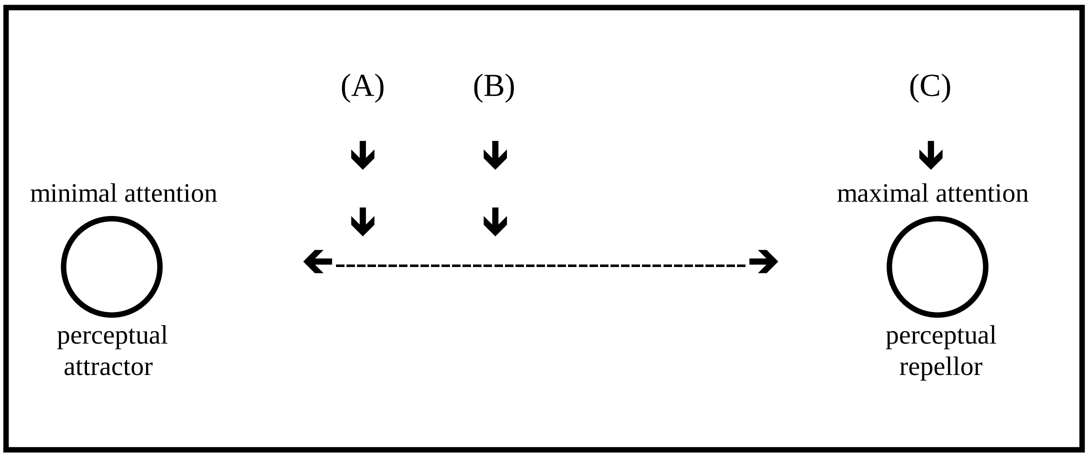

Figure 1. *Schematic representation of A&R Theory. (A) represents the hypothesised position of the non-musicians. They are shifted rightwards from the point of minimal attention due to their participation in a laboratory-controlled discrimination task. (B) represents the hypothesised position of the amateur musicians who pay more attention to the C major prototype but whose prototype still acts as a perceptual attractor. (C) represents the hypothesised position of the professional musicians who pay maximal attention to the music prototypes. One would assume that the average individual listening to speech (outside the laboratory) would fall at the minimal attention end of the continuum.*

At one end of the continuum is the point of maximal attention. Here the greatest amount of attention is being paid to the prototype and so the perceptual space around the prototype is stretched. Because discrimination is enhanced in the region of the prototype due to this perceptual stretching, the prototype functions like a perceptual repellor. In contrast, at the other end of the continuum is the point of minimal attention. Here the least amount of attention is paid to the prototype with the consequence that the perceptual space is shrunk around it. Discrimination is reduced in the region of the prototype and so it functions as a perceptual attractor. As mentioned above, the position of the prototype along this perceptual continuum at any given point in time is determined by the amount of attention paid to it by the listener. At the one extreme, for example, is the professional musician [here denoted by the letter (C)] who is highly attentive to incoming music prototypes. At the other extreme is the average person listening to speech or the average person listening to music whose phonetic or music prototypes function as attractors since not much attention is required by the system other than for recognition or entertainment respectively. Notice, however, that (A) in Figure 1 reflects the position of the non-musicians from Experiment 2 and not the population of non-musicians (who would most likely fall at the point of minimal attention). Because the non-musicians in Experiment 2 were being tested under controlled laboratory conditions, it is likely that their attention would have been focussed more onto the C major prototype than if they had been listening to music under normal uncontrolled circumstances.

A compromise in position is reached in the case of the amateur musician [denoted by the letter (B)] who pays neither maximal nor minimal attention to the music prototypes but whose C major prototype acts as a perceptual attractor. The finding from Experiment 2 that the discrimination of the amateur musicians around the P is significantly better than the discrimination of the non-musicians is reflected in the shifted position of the amateur musician's prototype along the perceptual continuum in Figure 2, away from the non-musicians. Because these amateur musicians still show a perceptual magnet effect, their position is closer to the point of minimal attention than the point of maximal attention.

The fact that prototypes are situated along a continuum means that they can function flexibly - depending upon the way that the listeners are using their system. This flexibility is reflected in the data from Experiment 2 where the C major prototype itself does not change, but its role in the system does. The similarities here between Lindblom's H&H Theory and A&R Theory are deliberate and reflect the often-debated commonality between the perception and production of speech.

Summary

In summary, the claim has been advanced through the experimental findings in this paper that perception is mediated by prototypes which can function as either perceptual attractors or perceptual repellors. According to A&R theory, the role of the prototype is determined by the amount of attention that is paid to it. The flexibility of the prototype is such that it can alter its role depending upon the needs of the listener. A&R Theory accounts for both the perception of speech and the perception of non-speech and thus reflects the commonalities that exist between the two systems.

Investigations are underway by the present author to use A&R Theory as an explanation for the infants' transition from a generalised to specialised perception in the first year of life (e.g. Werker & Tees, 1984). One proposed interpretation is that infants' native-language prototypes are initially repellor-like since the pre-linguistic infant can devote much of its attention to what speech actually sounds like. This focussing of attention causes a perceptual stretching in the region of the prototype, making it easier to distinguish all sounds of all languages regardless of whether they are phonemic in the baby's particular language environment (e.g. Trehub, 1976). It is only once the infant becomes distracted by higher-order linguistic processing such as developing its first words [at approximately 10-months of age (de Villiers & de Villiers, 1978)] or learning the rhythmic structure of its native language [at around 9-months of age (Jusczyk, Hohne, & Mandel, 1995)] that its attention is distracted from what speech actually sounds like, and its prototypes begin to act as attractors. This repellor-to-attractor conversion causes a shrinking in the perceptual space around the prototype, thus specialising the perceptual abilities of the infant listener.

Acknowledgements

The research contained in this paper was carried out as part of a doctoral dissertation at the University of Cambridge. The author wishes to thank Dr. Sarah Hawkins for her invaluable guidance throughout this project.

References

Acker, B.E., Pastore, R.E., & Hall, M.D. (1995). Within-category discrimination of musical chords: perceptual magnet or anchor? Perception & Psychophysics, 57(6), 863-74.

De Villiers, J. & de Villiers, P. (1978). Language Acquisition. Cambridge, Mass: Harvard University Press.

Grieser, D. & Kuhl, P.K. (1989). Categorization of speech by infants: Support for speech-sound prototypes. Developmental Psychology, 25, 577-88.

Howard, D., Rosen, S., & Broad, V. (1992). Major/minor triad identification and discrimination by musically trained and untrained listeners. Music Perception, 10(2), 205-220.

Iverson, P. & 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, 553-62.

Iverson, P. & Kuhl, P.K. (1996). Influences of phonetic identification and category goodness on American listeners' perception of /r/ and /l/. Journal of the Acoustical Society of America, 99(2), 1130-1139.

Jusczyk, P.W., Pisoni, D.B., & Mullenix, J. (1992). Some consequences of stimulus variability on speech processing by 2-month-old infants. Cognition, 43, 253-291.

Jusczyk, P.W., Bertoncini, J., Bijeljac-Babic, R., Kennedy, L.J., & Mehler, J. (1990). The role of attention in the speech of young infants. Cognitive Development, 5, 265-86.

Jusczyk, P.W., Hohne, E.A., & Mandel, D.R. (1995). Picking up regularities in the sound structure of the native language. In W. Strange (ed.): Speech Perception and Linguistic Experience: Theoretical and Methodological Issues in Cross-language Speech Research. Baltimore: York Press, 91-120.

Kuhl, P.K. (1983). Perception of auditory equivalence classes for speech in early infancy. Infant Behavior and Development, 6, 263-85.

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. (1992). Infants' perception of speech: Development of a new theory. Proceedings of the International Conference on Spoken Language Processing, Alberta, Canada, 449-55.

Kuhl, P.K. (1993). Innate predispositions and the effects of experience in speech perception: The Native Language Magnet Theory. In D. de Boysson-Bardies et al, S. de Schonen, P. Jusczyk, P. MacNeilage, and J. Morton (eds.): Developmental Neurocognition: Speech and Face Processing in the First Year of Life. The Hague: Kluwer Academic Publishers.

Kuhl, P.K. & Iverson, P. (1995). Linguistic experience and the `perceptual magnet effect'. In W. Strange (ed.): Speech Perception and Linguistic Experience: Issues in Cross-Language Research. Baltimore: York.

Lindblom, B.E.F. (1990). Explaining phonetic variation: A sketch of the H&H theory. In W. Hardcastle and A. Marchal (eds.): Speech Production and Speech Modeling. Dordrecht: Kluwer.

Lynch, M.P., Eilers, R.E., Oller, E.K., & Urbano, R.C. (1990). Innateness, experience, and music perception. Psychological Science, 1(4), 272-76.

Macmillan, N.A. & Creelman, C.D. (1991). Detection theory: A user's guide. New York: Cambridge University Press.

Nosofsky, R.M. (1984). Choice, similarity and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10 (1), 104-14.

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. Journal of Experimental Psychology: General, 115, 39-57.

Pisoni, D.B., Aslin,R.N., Perey, A.J., & Hennessy, B.L. (1982). Some effects of laboratory training on identification and discrimination of voicing contrasts in stop consonants. Journal of Experimental Psychology: Human Perception & Performance, 8, 297-314.

Sussman, J.E. & Lauckner-Moreno, 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, 539-52.

Trehub, S. (1976). The discrimination of foreign speech contrasts by infants and adults. Child Development, 47, 466-72.

Trehub, S. (1996). Perceptual anchors and magnets in infancy. Journal of the Acoustical Society of America, 99, 2521.

Werker, J.F. & Tees, R.C. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. Infant Behavior & Development, 7, 49-63.