

Generalization in perceptual learning for speech

TANYA KRALJIC and ARTHUR G. SAMUEL

State University of New York, Stony Brook, New York

Lexical context strongly influences listeners' identification of ambiguous sounds. For example, a sound midway between /f/ and /s/ is reported as /f/ in "sheri_" but as /s/ in "Pari_." Norris, McQueen, and Cutler (2003) have demonstrated that after hearing such lexically determined phonemes, listeners expand their phonemic categories to include more ambiguous tokens than before. We tested whether listeners adjust their phonemic categories for a specific speaker: Do listeners learn a particular speaker's "accent"? Similarly, we examined whether perceptual learning is specific to the particular ambiguous phonemes that listeners hear, or whether the adjustments generalize to related sounds. Participants heard ambiguous /d/ or /t/ phonemes during a lexical decision task. They then categorized sounds on /d/-/t/ and /b/-/p/ continua, either in the same voice that they had heard for lexical decision, or in a different voice. Perceptual learning generalized across both speaker and test continua: Changes in perceptual representations are robust and broadly tuned.

Speech is inherently variable both within and across speakers. Despite such variation, listeners are able to accurately perceive the speech signal across a wide range of speakers and contexts. Most theorists propose that listeners achieve such perceptual constancy by somehow compensating for, or *normalizing*, the variations in the acoustic signal. The assumption is that the perceptual system maintains abstract representations that consist of broadly defined phonemic categories, enabling many phonetic variations to be mapped onto relatively few perceptual categories.

However, a growing body of research suggests that listeners maintain representations that incorporate context-dependent information about specific speakers and situations (see, e.g., Nygaard & Pisoni, 1998). If listeners' perceptual categories are narrowly defined for particular speakers and contexts (as opposed to consisting of abstractions), these representations must somehow be adjusted to reflect experience. In fact, a number of studies show that with some exposure, speech that is initially difficult to understand becomes more intelligible (see, e.g., Bradlow & Bent, 2003): Listeners learn something about the speech they are hearing, and such learning leads to improved speech perception.

Norris, McQueen, and Cutler (2003) have termed such effects *perceptual learning*; Vroomen and colleagues (Bertelson, Vroomen, & de Gelder, 2003; Vroomen, van Linden, Keetels, de Gelder, & Bertelson, 2004) have termed similar cases *perceptual recalibration*. However, it is not clear *how* listeners are able to adjust to variations

in the speech signal: What information guides perceptual learning? What are the parameters of such learning? These questions are central to understanding perceptual learning. They are also critical for understanding the nature of perceptual representations more generally, and for understanding how listeners achieve such successful language comprehension in the face of rampant variation.

Listeners' ability to adjust preexisting phonemic categories appears to depend on lexical context, as Norris et al. (2003) have recently demonstrated. Dutch listeners were exposed to an ambiguous sound midway between /f/ and /s/; this sound occurred at the end of Dutch words that normally end in /f/, at the end of words that normally end in /s/, or at the end of nonwords. Listeners who heard the ambiguous sound in the context of /f/-final words later categorized more items on an /f/-/s/ continuum as /f/, whereas listeners who heard this sound in /s/-final words categorized more items on the same /f/-/s/ continuum as /s/. Hearing the ambiguous sound in nonwords produced no such shift. Thus, listeners use lexical knowledge to dynamically "tune" their phonemic representations to reflect the incoming speech signal.

Norris et al. (2003) suggested that the function of perceptual learning was to adapt to a particular speaker's accent. But perceptual learning could reflect adaptation to an accent or dialect *regardless* of speaker, or even a more general modification. In other words, perceptual learning may reflect very specific adjustments that are applied only to the particular words, phonemes, or speakers listeners are exposed to, or a more abstract adjustment that generalizes to new words, phonemes, and speakers.

Eisner and McQueen (2005) have recently addressed the specificity of perceptual learning. Participants were first exposed to a particular speaker's voice while they performed a lexical decision task. The participants then identified items from an /ɛf/-/ɛs/ continuum. For half of the participants, the continuum was presented in the

This project was supported by a National Science Foundation Graduate Fellowship, by NIMH Grant R0151663, and by NSF Grant 0325188. We thank Steve Goldinger and an anonymous reviewer for their constructive suggestions. Correspondence should be addressed to T. Kraljic, University of California, San Diego, Center for Research in Language, 9500 Gilman Drive, La Jolla, CA 92093-0526 (e-mail: tkraljic@crl.ucsd.edu).

same female voice that they had heard during the exposure phase. For the other half, the vowel portion (/ε/) of each item was produced in a different, but similar, female voice; the fricative portion was produced by the same speaker as in the exposure phase. Although a posttest questionnaire indicated that participants in the second group believed that the test items were produced by a new speaker, the perceptual learning effect was obtained.

In subsequent experiments, new participants performed the lexical decision task and then categorized items on a continuum that was created either with the original female speaker's fricatives but with a vowel produced by a male speaker (Experiment 2), or on one created entirely from the male speaker (Experiment 3). When only the vowel was changed, listeners showed significant perceptual learning, ruling out the possibility that the generalization to a new speaker was due to an insufficient acoustic difference between the new speaker's voice and the original speaker's voice. However, when both the vowel and the fricative were produced by the male speaker, listeners did not show perceptual learning. Eisner and McQueen (2005) concluded that (1) perceptual learning is speaker specific, and (2) because the effect generalized to "new" speakers in Experiments 1 and 2 (when the critical phoneme actually was produced by the original speaker, but the vowel was changed), learning occurs at a phonemic level.

In another recent perceptual learning study, Maye, Aslin, and Tanenhaus (2005) exposed listeners to a novel accent (using a synthetically created voice) and examined how they subsequently mapped new phonetic forms onto lexical items spoken by the same speaker. On separate days, participants heard the same 20-min story read in standard American English (Day 1) and in an accented English in which front vowels were systematically lowered (e.g., *wetch* instead of *witch*; Day 2). On both days, participants completed identical auditory lexical decision tasks immediately after listening to the story. Some items in the lexical decision task had standard American English front vowels; others had accented English (lowered) front vowels. Half of these items had occurred in the stories, and half were new. Participants more often and more quickly indicated that items with lowered front vowels (*wetch*) were words after hearing the story read in accented English than after hearing it in standard English. The increase in "word" responses to accented English items on Day 2 was significant even for new items, indicating that the participants had learned something general about the accent ("front vowels are lowered"), rather than something specific to the items that they had heard in the story. In a follow-up experiment, Maye et al. found that participants were no more likely to accept *raised* front vowel pronunciations (e.g., *weech*) as words after being exposed to lowered front vowel accented English. Thus, whatever perceptual change occurred after exposure to accented words did not result in phonemic categories that simply allowed for more noise in pronunciation; rather, the change was specific to the direction of the accent.

Many fundamental questions about perceptual learning remain unanswered. The results in hand do not clearly

specify the level(s) of representation that may be subject to perceptual retuning. For example, in the Norris et al. (2003) and the Eisner and McQueen (2005) studies, exposure to odd versions of /f/ or of /s/ increased report of those sounds. But the data do not determine that the perceptual system has been retuned at the phonemic level: A shift in the definition of place of articulation more generally (e.g., "accept a wider range of sounds as medial"), or even some mix of the phonemic and the featural levels (e.g., "accept a wider range of voiceless fricatives as medial") could also produce the observed shifts.

The available data also provide only preliminary suggestions about the extent to which perceptual learning is speaker specific; additional data are clearly needed. All of the demonstrations in this literature have been conducted with phonemes (vowels and fricatives) that depend on spectral cues that provide a fair amount of speaker-specific information. In the present study, the perceptual learning is based on manipulation of the stop consonants /d/ and /t/. These consonants differ in voicing, with this distinction based on more categorically perceived cues (for a review, see Repp, 1984) that do not provide as much speaker-specific information. The /d/-/t/ contrast has an additional property that is at least as important: The voicing distinction in such stops can be cued by a set of parameters (e.g., voice onset time and duration of aspiration) that are well controlled and that are applicable in essentially the same form to other stop contrasts (e.g., /b/-/p/). This property provides a very clean test of whether perceptual learning takes place at the phonemic level, or at a more abstract featural level (see below).

The present experiment addressed how general perceptual learning is with respect to both speakers and phonemes. Listeners may adjust a particular phonemic category only for the particular speaker to whom they have been exposed (e.g., *this speaker produces odd /d/s*). This relatively specific pattern of adjustment is consistent with what Eisner and McQueen (2005) found using /f/ and /s/ as the critical phonemes. Because the voicing distinction between the stop consonants in the present study was based on cues that provided less speaker specific information, it remains to be seen whether perceptual learning with these stop consonants would be speaker specific.

We used the paradigm developed by Norris et al. (2003). Listeners initially performed an auditory lexical decision task so that they could be exposed to an ambiguous sound midway between /d/ and /t/ (?dt). For half of the listeners, the ambiguous ?dt sound replaced the /d/ in 20 words such as *secondary* and *crocodile*; for the other half, the ?dt sound replaced the /t/ in words such as *cafeteria* and *frontier*. The 20 critical words were randomly interspersed among 180 other words or nonwords. Each participant heard the list read by either a female or a male speaker.

If this exposure leads to perceptual learning, listeners who heard the ?dt embedded in /d/-words should subsequently categorize more items on a /d/-/t/ continuum as /d/; listeners who heard the ?dt embedded in /t/-words should categorize fewer items as /d/. To test whether this perceptual learning is specific to a particular speaker, all listeners categorized items on /d/-/t/ continua produced

by two speakers—the same speaker they had been exposed to in the lexical decision task, and a new speaker.

To test whether perceptual learning occurs at a phonemic level or at a more abstract featural level of representation, all listeners also were required to label items on a /b-/p/ continuum (which shares the same featural voicing relationship as /d-/t/ does). The /b-/p/ items were presented in the same two voices as were the /d-/t/ items. If listeners retune their phonemic categories in a general way, we should see the same categorization shifts for new-voice continua and for the /b-/p/ continua as for the /d-/t/ continuum presented in the same voice used in the lexical decision task. If retuning is more specific, then any categorization shift should be restricted to test items that match the original training conditions.

METHOD

Participants

Seventy-two students from the State University of New York at Stony Brook received research credit in a psychology course for their participation. All were 18 years of age or older, and all identified themselves as native English speakers with normal hearing.

Materials

Phase 1: Exposure (lexical decision). Two experimental lists and two control lists were created for the auditory lexical decision task, each with 100 words and 100 nonwords. The lists were identical except for 40 critical items, which were words in the experimental conditions, and nonwords in the control conditions.

Stimulus selection. The 40 critical words ranged in length from two to five syllables and contained no /b/ or /p/ (see Table 1). Twenty of the words also contained no /t/, but each had a single instance of the critical phoneme /d/. The other 20 critical words contained no /d/; these each had a single instance of the phoneme /t/. Each /d/ or /t/ occurred in the initial position of a syllable that was relatively late in the word, and that had primary or secondary stress, to ensure

that these critical phonemes were well articulated and preceded by enough of the word to generate strong lexical activation. The two sets of critical words were matched in mean syllable length and frequency of occurrence (Zeno, Ivens, Millard, & Duvvuri, 1995). We selected 100 filler words that had no occurrences of /d/, /t/, /b/, or /p/. The fillers were matched to the critical words in terms of stress pattern, number of syllables, and word frequency; 60 of these words were used in the lexical decision task. One hundred filler nonwords were created by changing one phoneme per syllable of each filler word; phonemes were changed to another phoneme with the same manner of articulation (i.e., glides changed to glides, stops to stops, etc.). As with the filler words, no /b/, /p/, /d/, or /t/ appeared in any position. Applying this method to the critical words, we made critical nonwords with either a /d/ or a /t/ in the same position as in the words. Each participant heard 100 words and 100 nonwords: In the ?D condition, there were 20 /d/-words (containing ?dt), 20 normal /t/-words, and 60 filler words; for the ?T group, the /t/-words contained the ?dt, and the /d/-words were normal.

Stimulus construction. Each of the 40 critical words, 40 critical nonwords, 60 filler words, and 100 filler nonwords was recorded by both a male and a female speaker. Each speaker produced a second version of every critical word and nonword in which the critical phoneme (/d/ or /t/) was replaced with the other one (e.g., *crocodile* and *crocodile*; *cafeteria* and *cafeteria*). By recording both a /d/ and a /t/ version of the same word, we were able to create a unique ambiguous (?dt) mixture for each word pair for use in the exposure phase. To create the ambiguous ?dt sound, we varied three cues: the relative amplitude weightings of /t/ and /d/ in the mixture (100% /t/ plus 0% /d/, down to 0% /t/ plus 100% /d/, in 5% increments), the length of ?dt after the voicing burst onset (longer aspiration favors /t/; range = 100% of the original /t/ length, down to 0% of the original /t/, in 5% increments), and the length of silence before the burst onset (longer favors /t/; range = 40 to 0 msec, in 2-msec steps). We made 21 mixtures of the /d/ and /t/ phonemes for each critical item. A single mixture for each item (the most ambiguous mixture, as determined by the authors and several phonetically untrained raters) was selected for use in the experiment.

Eight lexical decision lists were created, varying the critical items in which participants would be exposed to the ?dt mixtures. The

Table 1
Critical Stimuli Presented During the Lexical Decision Exposure Phase

/d/ Words	Syllables	Frequency	/t/ Words	Syllables	Frequency
handy	2	8.7	frontier	2	25.9
kingdom	2	45.7	magnetism	3	12.6
crocodile	3	7.8	lunatic	3	1.38
hazardous	3	4.9	relative	3	37.8
agenda	3	2.4	casualty	3	1.4
merchandise	3	66.7	novelty	3	3.0
iodine	3	7.7	royalty	3	2.7
melody	3	9.3	infantile	3	1.2
residence	3	8.4	overtime	3	5.5
confidence	3	27.1	authentic	3	1.4
evidence	3	81.8	romantic	3	16.8
accordion	4	1.0	warranty	3	5.7
remedial	4	1.5	scientific	4	76.8
armadillo	4	1.6	voluntary	4	11.6
comedian	4	1.4	cemetery	4	7.0
legendary	4	2.9	military	4	81.4
coincidence	4	2.7	momentary	4	3.0
avocado	4	0.7	hesitation	4	5.4
secondary	4	28.5	consultation	4	2.6
academic	4	13.4	cafeteria	5	5.7
Average	3.35	16.3		3.4	15.5

Note—Sixty filler words (with no /d/, /t/, /b/, or /p/) were an average of 3.3 syllables in length and had an average word frequency of 16.45.

experimental groups heard either words with intact /t/s and ambiguous /d/s (?D) or the reverse (intact /d/-words and ambiguous /t/s) (?T). The control groups heard the same ambiguous sounds but in the context of nonwords (?CD and ?CT). These four conditions were crossed with the voice in which the lists were presented (male or female).

Phase II: Category identification. In the second phase of the experiment, participants heard six items on four separate vowel-consonant-vowel (VCV) continua, each presented in two voices. Each of the endpoints for the four continua (/ada/-/ata/, /aba/-/apa/, /ɪdɪ/-/ɪtɪ/, /ɪbɪ/-/ɪpɪ/) was recorded by the same male and female speakers who produced the lexical decision stimuli. Twenty-one mixtures of each continuum were created, and raters chose six consecutive tokens for each continuum. These stimuli ranged from relatively /d/- or /b/-like to relatively /t/- or /p/-like, with four ambiguous points in between.

Procedure

Participants were randomly assigned to one of the eight lexical decision conditions. Because the lack of perceptual learning with nonwords is well established (see Eisner & McQueen, 2005; Norris et al., 2003), the control conditions were simply a check on our manipulations. Accordingly, we assigned twice as many participants to the experimental groups as to the control groups.

Up to 3 participants were tested simultaneously in a soundproof booth. In the lexical decision task, participants were instructed to respond "Word" or "Nonword" to each item by pressing the corresponding button on a response panel. The participants were not told that some of the items might have ambiguous sounds.

After the lexical decision phase, all participants categorized sounds on /b/-/p/ and /d/-/t/ continua, presented in both the original voice to which they had been exposed as well as in a previously unheard voice. The continua were blocked by phoneme (/b/-/p/ vs. /d/-/t/) and by speaker (same vs. different). The participants always completed both the /b/-/p/ and then the /d/-/t/ continua for a single voice first; each continuum was then tested in the other voice. We presented /b/-/p/ first to prevent any possible carryover effects from categorizing /d/-/t/ (the trained sounds) onto /b/-/p/. The order of presentation voice was counterbalanced. Ten randomizations of each continuum were presented.

RESULTS

Lexical Decision

Any participant whose accuracy on the lexical decision task was below 70% was replaced. Six of the 72 participants were replaced for this reason.

Table 2 provides the accuracy and response time (RT) data for each type of critical item (ambiguous ?d or ?t vs. natural /d/ or /t/). Listeners performed very well on the lexical decision task; mean accuracy was 94.7%. Accuracy was higher for the natural versions of the critical items (98.1%) than for the ambiguous versions (91.4%)

[$F_1(1,47) = 16.88, p < .001$; $F_2(1,19) = 18.96, p < .001$]. People correctly categorized ambiguous items (838 msec) as words more quickly than they categorized the natural versions (874 msec) [$F_1(1,47) = 6.177, p = .017$; $F_2(1,19) = 5.69, p = .028$]. Overall, these data suggest that our ?dt mixtures were relatively natural sounding.

Category Identification

For each participant, we calculated the average percentage of test syllables identified as voiced (/d/ or /b/). There was a clear effect of lexical decision exposure condition on phonemic categorization performance. Listeners who were exposed to the ambiguous ?dt phone in words that normally have a /d/ categorized more items on our continua as /d/ or /b/ (35.9%) than did those who heard ?dt in words that normally have a /t/ (32.2%) [$F(1,40) = 5.19, p = .028$]. In contrast, there was no such training effect for the control groups, in which people were exposed to the ambiguous ?dt phoneme in the context of nonwords [$F(1,16) < 1, p = .513$]. Figure 1A shows the identification functions for the experimental groups, and Figure 1B shows the results for the controls. The significant training effect in the experimental groups, and the lack of this effect in our control groups, confirms that people do use lexical knowledge to adjust their perceptual representations, even for categorically perceived stop consonants.

The remaining analyses examine whether perceptual learning generalizes, and if so, in what ways. In these analyses, we collapse over the two factors of order of voice (same first, same second) and vowel context (/a/ or /ɪ/), because there was no interaction between order and training condition [$F(1,40) < 1, p = .689$] or between vowel context and training condition [$F(1,40) = 1.720, p = .197$].

We first consider whether the retuning of perceptual representations generalized to a new speaker. Figure 2 shows that this was indeed the case. There was no interaction between speaker (same, different) and training (?d, ?t) [$F(1,40) < 1, p = .647$]. The perceptual learning effect was reliable for a different voice [$F(1,190) = 5.075, p = .025$], just as it was for the same voice [$F(1,40) = 6.340, p = .013$] heard during the exposure phase.

Perceptual learning also generalized to new phonemes that shared the same featural relationship as did the phonemes to which listeners were exposed. There was no interaction between consonant at test (same vs. different) and training (?d vs. ?t) [$F(1,40) = 2.375, p = .131$]. The

Table 2
Mean Accuracy and Response Times (RTs) for Correct Items (in Milliseconds)
for Natural and Ambiguous Critical Words and Nonword Controls

	Natural				Ambiguous			
	/t/		/d/		?t		?d	
	% Correct	RT	% Correct	RT	% Correct	RT	% Correct	RT
Critical words	99.4	893	96.9	854	98.1	851	84.6	824
Nonword controls	92.9	929	91.3	939	91.7	976	94.2	891




Figure 1. (A) Overall training effect, experimental groups. (B) Control groups, no training effect.

perceptual learning effect was reliable for the /b/–/p/ phonemes [$F(1,190) = 8.859, p = .003$]. In fact, the effect for /b/–/p/ was actually more stable than the effect for /d/–/t/ [for the latter, $F(1,190) = 2.45, p = .12$]. Figure 3 shows the identification functions for both cases. Recall that we always tested /b/–/p/ before /d/–/t/, to avoid any contamination of the transfer test. The stronger effect for /b/–/p/ presumably demonstrates the erosion of the perceptual retuning caused by the dozens of test syllables. The average size of the perceptual learning effect for the first two

continua tested was almost double that for the second two tested (5% vs. 2.6%). This supports the idea that perceptual learning for /d/–/t/ fades as a function of order of testing (although the difference is not statistically significant).¹

DISCUSSION

Listeners adjust their perception of phonemic categories even for stop consonants, and this perceptual learning generalizes both to new speakers and to new pho-




Figure 2. The training effect generalizes to a new speaker: (A) Training effect when participants are tested on the voice to which they were exposed in the lexical decision task. (B) Training effect when participants are tested on a voice different from the one to which they were exposed.




Figure 3. The training effect also generalizes to new phonemes: (A) Training and test on same consonant (/d/-/t/). (B) Training and test on different consonant (/b/-/p/)

nemes. The shifts for these categorically perceived stops are about one fourth the size of those for fricatives we have tested (Kraljic & Samuel, 2005), but they are reliably present.

Norris et al. (2003) suggested that perceptual learning might be used to adjust to a particular person's accent. Testing this suggestion requires information about what kind of phonemic input is sufficient to make adjustments, and whether such adjustments affect processing of speech from subsequent speakers. We found that listeners are able to apply learning from one phoneme and speaker to new phonemes and new speakers, indicating that perceptual learning occurs at the featural level.

This result is contrary to Eisner and McQueen's (2005) conclusion that perceptual learning occurs at the phonemic level and is speaker specific. However, the present results are not necessarily incompatible with theirs. Rather, taken together, they imply something quite interesting about the way in which the perceptual system may be organized, and about how perceptual learning can be used to improve perception: The perceptual system could actually use relatively simple procedures to achieve the impressive generality of learning that we have observed. Recall that we chose stop consonants in order to explore our questions primarily because they (1) contrast on a cue that does not provide very systematic speaker information (as opposed to fricatives and vowels), and (2) offer a very clean test of feature generalization (because of the comparability of the /d/-/t/ contrast with the /b/-/p/ contrast; in both cases, voiceless sounds have longer pre-release silence and longer aspiration than voiced sounds do). Thus, any learning about /d/ or /t/ in one voice could be applied to /b/-/p/, and to the other voice. The fricative contrasts used by

Norris et al. (2003) and Eisner and McQueen (2005) are spectrally cued, and they systematically vary across male versus female voices. This could lead to speaker-specific, and possibly phoneme-specific, perceptual learning.

This analysis highlights the flexibility of the speech perception system. It can make use of different levels or types of information, and it can generalize (or not) accordingly. Such flexibility is consistent with what has been observed more generally in other levels of language processing: Resolving syntactic or lexical ambiguities, comprehending referring expressions, and many other types of processing often occur with respect to contextual and speaker-related information, when that information is available. This kind of cognitive flexibility appears to play an important role in listeners' ability to recognize words in the face of the enormous variability of the speech signal.

REFERENCES

- BERTELSON, P., VROOMEN, J., & DE GELDER, B. (2003). Visual recalibration of auditory speech identification: A McGurk aftereffect. *Psychological Science*, *14*, 592-597.
- BRADLOW, A. R., & BENT, T. (2003). Listener adaptation to foreign accented English. In M. J. Sole, D. Recasens, & J. Romero (Eds.), *Proceedings of the XVth International Congress of Phonetic Sciences* (pp. 2881-2884). Barcelona: Futurgraphic.
- EISNER, F., & MCQUEEN, J. M. (2005). The specificity of perceptual learning in speech processing. *Perception & Psychophysics*, *67*, 224-238.
- KRALJIC, T., & SAMUEL, A. G. (2005). Perceptual learning for speech: Is there a return to normal? *Cognitive Psychology*, *51*, 141-178.
- MAYE, J., ASLIN, R., & TANENHAUS, M. (2005). *The weckud wetch of the West: Rapid adaptation to a novel accent*. Manuscript submitted for publication.
- NORRIS, D., MCQUEEN, J. M., & CUTLER, A. (2003). Perceptual learning in speech. *Cognitive Psychology*, *47*, 204-238.

- NYGAARD, L. C., & PISONI, D. B. (1998). Talker-specific learning in speech perception. *Perception & Psychophysics*, **60**, 355-376.
- REPP, B. H. (1984). Categorical perception: Issues, methods, findings. In N. J. Lass (Ed.), *Speech and language: Advances in basic research and practice* (pp. 243-335). New York: Academic Press.
- VROOMEN, J., VAN LINDEN, S., KEETELS, M., DE GELDER, B., & BERTELSON, P. (2004). Selective adaptation and recalibration of auditory speech by lipread information: Dissipation. *Speech Communication*, **44**, 55-61.
- ZENO, S., IVENS, S., MILLARD, R., & DUVVURI, R. (1995). *The educator's word frequency guide*. Brewster, NY: Touchstone Applied Science Associates.

NOTE

1. In an ongoing follow-up project, we have used a single test continuum (/idi/-/iti/) immediately after training with only 10 (rather than 20) critical items in the lexical decision task. The size of perceptual learning effect (4%-5%) was similar to what we have reported here and was statistically reliable [$F(1,120) = 5.33, p = .023$].

(Manuscript received December 23, 2004;
revision accepted for publication July 14, 2005.)