

Children Discover the Spectral Skeletons in Their Native Language Before the Amplitude Envelopes

Susan Nittrouer and Joanna H. Lowenstein
Ohio State University

Robert R. Packer
Washington State University

Much of speech perception research has focused on brief spectro-temporal properties in the signal, but some studies have shown that adults can recover linguistic form when those properties are absent. In this experiment, 7-year-old English-speaking children demonstrated adultlike abilities to understand speech when only sine waves (SWs) replicating the 3 lowest resonances of the vocal tract were presented, but they failed to demonstrate comparable abilities when noise bands amplitude-modulated with envelopes derived from the same signals were presented. In contrast, adults who were not native English speakers but who were competent 2nd-language learners were worse at understanding both kinds of stimuli than native English-speaking adults. Results showed that children learn to extract linguistic form from signals that preserve some spectral structure, even if degraded, before they learn to do so for signals that preserve only amplitude structure. The authors hypothesize that children's early sensitivity to global spectral structure reflects the role that it may play in language learning.

Keywords: speech perception, sine wave speech, children's speech perception, perceptual object formation

One of the first ideas that students of human communication encounter is the acoustic theory of speech production. Largely attributed to Fant (1960), this theory states that the speech signal primarily arises from a glottal source, which provides broadband periodic sound to the vocal tract, and the vocal tract, which serves to filter that sound. Historically, the study of human speech perception has focused on the way that the vocal-tract filter shapes the speech signal to create linguistic form (primarily phonetic form). The acoustic properties arising from the actions of the vocal-tract filter that have been of most interest to investigators are those that are spectrally specific and temporally brief, properties that have come to be called *acoustic cues* (Repp, 1982). Scores of studies have been published documenting the relations between these kinds of properties and phonetic form, but the study reported here dealt with structure in the speech signal that is broader spectrally and longer temporally, structure that can be termed *global* in nature.

When one considers acoustic properties affiliated with long stretches of the speech signal (i.e., those that are the length of several segments or even words), properties arising from the glottal source generally come to mind. These properties fit under the heading of *prosody*. For example, fundamental frequency of

the source spectrum provides information useful in speech perception, as does the degree of breathiness, and the temporal pattern of when the source is active and when it is not (e.g., Bricker & Prusansky, 1976; Carrell, 1981; Jassem, 1971). These source properties inform mature listeners primarily about characteristics of the speech signal that are not phonetic in nature, such as speaker identity, but they also play an important role in language learning. *Prosodic bootstrapping* describes a process in which infants use prosodic cues to gain access to elementary syntactic structure, before lexical knowledge is well established (e.g., Gleitman & Wanner, 1982; Jusczyk, 1997; Morgan & Demuth, 1996; Soderstrom, Seidl, Kemler Nelson, & Jusczyk, 2003).

But the vocal-tract filter also imparts to the speech signal several kinds of structure that are best resolved from long signal stretches. These kinds of structure arise from general postural settings of the vocal tract, such as whether a language or speaker tends to use pharyngeal constrictions or heavy nasalization, from the continuously changing sizes and shapes of vocal tract cavities that determine formant frequencies, and even from modulations in overall amplitude due to the degree of vocal tract constriction. Likely because of the emphasis on acoustic cues, and their relation to explicitly phonetic form, the potential role of these more global kinds of structure went largely unexamined in early speech perception research. It was not until the early 1980s that a study breaking with traditional methods reported that adult listeners can both recognize a signal as speech and recover linguistic units when only the skeleton of spectral structure is presented (Remez, Rubín, Pisoni, & Carrell, 1981). In that study, the authors replaced the normally rich harmonic spectrum of speech with three time-varying sinusoids that traced the dynamic changes in the three lowest formants. The finding that adult listeners could perceptually integrate these impoverished signals into coherent linguistic forms could not be explained within the framework of speech perception theories existing at that time, theories that generally held that

Susan Nittrouer and Joanna H. Lowenstein, Department of Otolaryngology–Head & Neck Surgery, College of Medicine, Ohio State University; Robert R. Packer, Department of Psychology, Washington State University.

This work was supported by the National Institute on Deafness and Other Communication Disorders Grant R01 DC000633. We thank Vignesh Hariharan for help with programming.

Correspondence concerning this article should be addressed to Susan Nittrouer, Department of Otolaryngology–Head & Neck Surgery, College of Medicine, Ohio State University, Cramblett Hall, Room 4126, 456 West 10th Avenue, Columbus, OH 43210. E-mail: nittrouer.1@osu.edu

listeners gather acoustic cues from the signal and use them to recreate the string of phonetic segments intended by the speaker. So while the finding that listeners could make sense of signals consisting only of these spectral skeletons was interesting, the significance of the finding was not clear.

Later, another group of investigators used a different signal processing scheme to represent a different kind of global structure imparted to the signal by vocal tract actions: Shannon, Zeng, Kamath, Wygonski, and Ekelid (1995) described a study in which they divided the speech spectrum into two, three, or four bands, recovered the amplitude envelopes (AEs) of each of those bands, and modulated noise bands limited by the same frequencies used to divide the original speech spectrum using those envelopes. In contrast to SW speech, which preserves only the global spectral structure, these AEs largely eliminated the spectral structure of the signal. Nonetheless, Shannon et al. showed that adults were able both to recognize those signals as speech and to recover linguistic form. This finding had clinical significance because of its relation to the development of the cochlear implant, which had recently made its appearance. Cochlear implants provide little in the way of spectral structure, so if listeners do indeed understand speech by recovering brief, spectrally specific bits of the signal, then these devices should not work. By 1995, however, results were clearly showing that hearing-impaired listeners were able to understand the signals provided by their cochlear implants. The findings of Shannon et al. served both to corroborate what researchers already knew hearing-impaired listeners could do as well as to perturb researchers' collective attention away from their earlier held focus on specific acoustic cues and their relation to specific phonemes. It became apparent that researchers' theories of human speech perception needed some tweaking.

At the same time that investigators were discovering that adults could recover linguistic form from processed signals preserving only the global structure created by the supraglottal vocal tract, scientists examining the development of speech production and perception were finding that this kind of structure may play an important role in language acquisition. For example, Charles-Luce and Luce (1990, 1995) showed that children, unlike adults, do not perform detailed acoustic-phonetic analyses when recognizing spoken words; instead children use strategies based on more global acoustic properties. Boysson-Bardies, Sagart, Halle, and Durand (1986) showed that the long-term spectra of 10-month-old infants resembled those of adults in their native language community. So before they had uttered their first real word, these infants had discovered the postural vocal-tract settings that help to define their native language.

Another group of investigators asked the question of whether or not children show the same trading relations between the spectro-temporal properties used in phonetic judgments that adults show. Work with adults had already demonstrated that the settings of some relevant acoustic cues that evoked specific phonetic judgments from listeners could be influenced by the settings of other properties. For example, Mann and Repp (1981) showed that adults would label noises with lower spectral means as *s* when formant transitions in trailing vocalic portions were appropriate for an alveolar constriction; but if those formant transitions were appropriate for a palatal constriction, then the mean frequency of the noise had to be higher for listeners to label the syllable as starting with *s*. Many psycholinguists of the time considered the

spectral shape of obstruent noises to be a primary, context-independent feature supporting correct classification by human listeners; vocalic formant transitions were considered secondary (e.g., Stevens & Blumstein, 1978). According to this view, human listeners learn how the secondary properties vary with the primary properties through extensive experience, and so the child must surely use spectral shape initially to label obstruent noises (Stevens, 1975).

In 1987, Nittrouer and Studdert-Kennedy reported a study in which they asked if children showed the same trading of acoustic cues as adults. In other words, did children appreciate how the setting of a primary acoustic property could be affected by a secondary property? The expectation was that children would not show trading relations of the same magnitude as adults did. Contrary to that expectation, however, Nittrouer and Studdert-Kennedy observed that children's labeling responses actually indicated stronger effects of the following formant transitions on the spectral composition of noises associated with specific sibilant decisions. Although that finding seemed at odds with theoretical notions of the day, it was nonetheless largely supported by subsequent research (e.g., Mayo, Scobbie, Hewlett, & Waters, 2003; Nittrouer, 1992; Nittrouer & Miller, 1997a, 1997b; Watson, 1997). Within the context of findings showing that adults can comprehend speech with only global spectral structure (e.g., Remez et al., 1981) and that infants replicate in their own babbled productions the global properties of speech produced by adults in their language community, a new hypothesis evolved (e.g., Nittrouer, 2006). That hypothesis suggested that children use global spectral structure to gain initial access to linguistic form. Then, as they acquire experience with their native language, they discover the spectro-temporal details specific to phonetic form in that language.

The study reported here sought to compare the abilities of native English-speaking children and adults to recognize signals processed to preserve only the spectral skeletons or the AEs (two forms of global structure created by the vocal-tract filter). This was considered an appropriate and timely focus because results could extend researchers' understanding of the development of speech perception, just as their perspectives concerning initial perceptual capacities and subsequent learning are evolving. Traditionally, theories of language acquisition have assumed that phonetic units are available to infants very early in life (Kuhl's 1987 report provides a substantive review), but now the idea is emerging that infants must discover these units through protracted perceptual learning (e.g., Beckman & Edwards, 2000; Nittrouer, 2001; Werker & Yeung, 2005). So the question arises of how children ever gain a toehold into a signal as highly continuous and variable as the speech signal. In the current study, we explored the possibility that such a toehold may be provided by some form of global acoustic structure, most likely the spectral skeleton.

One study has presented children with AE stimuli created from isolated words and sentences with high syntactic and semantic context (Eisenberg, Shannon, Schaefer Martinez, Wygonski, & Boothroyd, 2000). Results showed that 5- to 7-year-old children were poorer at recognizing both isolated words and sentences than adults and older children when presented with the same number of channels, spurring the conclusion that "speech pattern recognition is still maturing in children 5 to 7 years of age" (p. 2,709). It is not clear whether the authors are suggesting that some aspect of auditory development is not completed by 5 to 7 years of age, or

that children at these ages have not learned to recover perceptually meaningful form from those impoverished signals. Trying to determine which of these alternative explanations accounts for their results was one goal of the present experiment.

It has been a long-held perspective that any measured difference in children's and adults' speech perception may be attributable to immature sensory pathways in children (e.g., Sussman, 1993), and so it might be that the 5- to 7-year-old listeners in the Eisenberg et al. (2000) study did not process the AE signals in the same way as did adults. Consequently, the maturity of the auditory systems in listeners in the current study was an issue and needed to be controlled. To do so, we included a group of listeners with mature auditory systems but with different linguistic experiences from those of our English-speaking adults: namely, second language learners of English. We could have used listeners with any native language who had learned English as a second language but chose to use native speakers of Mandarin Chinese because it has been demonstrated that adult Chinese listeners recognize Chinese sentences processed to preserve only AEs as well as adult English listeners recognize English sentences processed in the same way (Fu, Zeng, Shannon, & Soli, 1998).

The other possible explanation for the results of Eisenberg et al. (2000) is that the sensory information in the AE stimuli is processed in the auditory system similarly by the young and older listeners but that the young listeners are not able to organize this information in such a way as to recover perceptual form. Visual analogies of what must be accomplished in listening to these acoustically impoverished signals abound and include examples such as the classic Rubin's vase, which can be recognized as a vase or two portraits: The sensory information is sparse, and the image that is recovered depends on how the perceiver organizes that information. In the work described here, the question was asked if different groups of listeners, and so listeners with different amounts and types of experience with the linguistic form to be recovered, were affected differently by different kinds of signal processing. The answer to this question could inform researchers as to what kinds of signal structure aid children in language learning.

In sum, in the current experiment we compared speech recognition for signals processed to preserve either spectral skeletons or AEs in order to explore the roles of these forms of structure in speech perception and its development. We processed sentences to preserve only spectral skeletons by creating SW analogs of those sentences. We obtained global amplitude structure for sentences using the methods of Shannon and colleagues (1995). The focus of this study was on the question of whether children recover linguistic form from spectral skeletons or AEs as well as adults do in speech perception. The answer to this question could extend researchers' understanding of how children acquire language. We incorporated nonnative speakers of English into the study to serve as a control group with mature auditory pathways but without the native language experience of our native adult listeners.

Method

Participants

Data are reported for 120 participants: 40 adult native English speakers, 40 seven-year-old native English speakers, and 40 adult

native Mandarin speakers who were competent second-language speakers of English. All Mandarin speakers had started learning English between the ages of 12 and 14, had spoken it an average of 14 years, and were in the United States as either graduate students, postdoctoral fellows, or junior faculty. All adults were between 18 and 38 years of age, with mean ages of 25 years for the English-speaking adults and 27 years for the Mandarin-speaking adults. To participate, all participants had to pass a hearing screening and demonstrate appropriate abilities to comprehend and produce spoken English. Half of the participants in each group listened to the SW stimuli and half listened to the AE stimuli, making a total of six groups of listeners. In all but two of these groups the proportions of male and female listeners were exactly the same (10 in each group): For the SW stimuli, 12 English-speaking adult listeners were female; for the AE stimuli, 12 Mandarin-speaking listeners were female.

Equipment

All speech samples were recorded in a soundproof booth, directly onto the computer hard drive, via an AKG C535 EB microphone, a Shure M268 amplifier, and a Creative Labs Soundblaster analog-to-digital converter. Perceptual testing took place in a soundproof booth, with the computer that controlled the experiment in an adjacent room. The hearing screening was done with a Welch Allen TM262 audiometer and TDH-39 earphones. Stimuli were stored on a computer and presented through a Samson headphone amplifier and AKG-K141 headphones.

Stimuli

For the 36 four-word sentences we used, we either adopted them from Boothroyd and Nittrouer (1988) or developed them according to the same rules used for developing those sentences: All sentences consisted entirely of monosyllabic words, were syntactically appropriate for English, but were semantically anomalous. Most of these sentences had some version of a subject-verb structure (e.g., *Dumb shoes will sing. Knees talk with mice*), although five had a command structure (e.g., *Paint your belt warm*). These sentences can be found in the Appendix. Six sentences were used for practice, and 30 were used for testing. Sentences were recorded by an adult male speaker of American English at a 44.1-kHz sampling rate with 16-bit digitization. Figure 1 displays a natural speech sample of *Late forks hit low* (top panel), the SW replica (center panel), and the four-channel AE replica (bottom panel). We equalized all sentences for mean root-mean-square amplitude across sentences before any processing was done.

For creating the SW stimuli, we used a Praat routine written by C. Darwin and available at his Web site (http://www.lifesci.sussex.ac.uk/home/Chris_Darwin/Praatscripts/SWS) to extract the center frequency of formants. We generated stimuli from these formant frequencies and compared spectrograms to spectrograms of the original stimuli to ensure that the trajectories of the SWs matched those of the formant frequencies.

To create the AE stimuli, we wrote a MATLAB routine. We created both four- and eight-channel AE stimuli to present to listeners because we wanted to have AE stimuli that would produce recognition scores in the same ballpark as those we expected

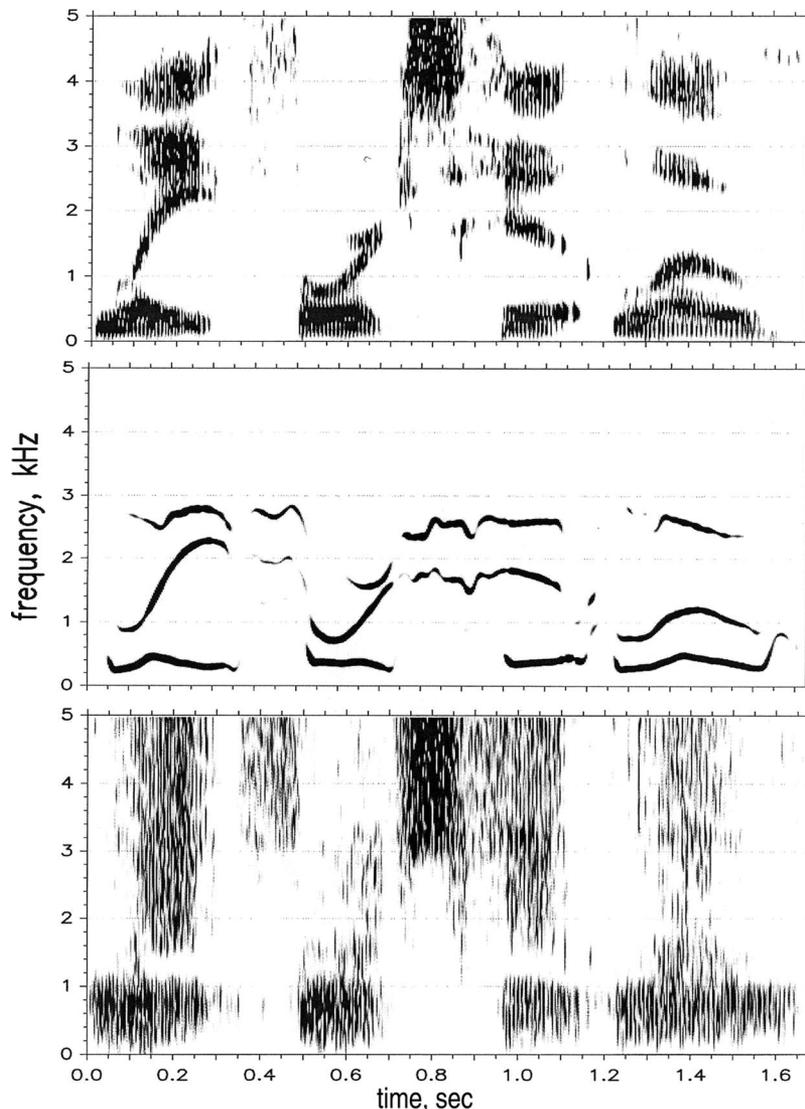


Figure 1. Spectrograms of the sentence *Late forks hit low*. Top: Natural speech sample from which sine wave and amplitude envelope stimuli were derived. Middle: Sine wave stimulus. Bottom: Four-channel amplitude envelope stimulus.

to obtain for the SW stimuli, and that was difficult to predict before the experiment. All signals were first low-pass filtered with an upper cutoff frequency of 8,000 Hz. For the four-channel stimuli, cutoff frequencies between bands were 800, 1600, and 3200 Hz; for the eight-channel stimuli, cutoff frequencies were 400, 800, 1200, 1800, 2400, 3000, and 4500 Hz. Each channel was half-wave rectified, and we used results to modulate white noise limited by the same band-pass filters as those used to divide the speech signal into channels. We used a 160-Hz high-frequency cutoff to low-pass filter resulting bands of modulated noise and combined them.

Procedure

All stimuli were presented under headphones in a soundproof booth at 68 dB sound pressure level. Practice was the same for

listeners hearing both the AE and SW stimuli, but listeners in the SW condition were given one additional task prior to practice because it can be difficult for some listeners to perceptually integrate SW stimuli into forms that are speechlike. A similar problem has not been reported for AE stimuli. To facilitate listeners' abilities to hear the SW stimuli as speechlike, we played a story before they heard the practice sentences. The story was presented first as natural speech, and then as SW replicas derived from that speech. A picture book without words accompanied the story.

Each of the six practice sentences was presented, either as an AE or SW signal. For listeners in the AE condition, three each were presented as four- and eight-channel AE stimuli. For each practice sentence, the natural version was played first, and the listener was instructed to listen and repeat it. Next the listener was told "Now

you will hear a robot say the sentence. Listen carefully so you can repeat it." The notion of a robot was invoked to help personify the stimuli. For both the AE and SW conditions, listeners were asked at the end of the six-sentence practice if they were able to recognize the robot's productions as speech. It was made clear that the question was whether they could hear it as if it were speech, rather than could they understand every word. If a listener either was unable to repeat any of the unprocessed sentences or failed to report hearing the processed signals as speech, they were dismissed. Thirteen individuals in addition to the 120 listeners whose data are reported here heard the practice sentences and were dismissed because they were unable to hear the SW stimuli as speech: 6 each from the children's and Mandarin-speakers' groups; 1 from the adult English-speakers' group. None of our listeners had difficulty hearing AE stimuli as speech. Listeners were paid for their participation, regardless of whether they could hear the stimuli as speech, so there was no reason for any listener to falsely state that the stimuli were heard as speech if they were not.

During testing, the order of presentation of the sentences was randomized independently for each listener. Each sentence was presented three times in a row. It is a common impression when listening to speech degraded in some way that recognition might improve if the sentence could just be heard again. Because of that impression, it was thought that recognition might improve with repeated trials, and the amount of improvement might vary across listener groups; that is, perhaps some listeners would benefit more than others from hearing each sentence again. If true, then optimal performance would not be obtained for one or more groups with just one trial.

The listeners presented with the AE stimuli heard half of them as four-channel and half of them as eight-channel stimuli. The selection of sentences to be presented as four- or eight-channel stimuli was randomly made for each listener by the software, and presentation of these two kinds of stimuli was intermingled during testing with the stipulation that no more than two four- or eight-channel stimuli could be presented in a row.

After listeners heard the sentences in their processed forms, these sentences were played to listeners in their unprocessed

forms. None of the listeners had difficulty comprehending those unprocessed sentences.

Results

The percentage of words recognized correctly across all sentences within any one condition served as the dependent measure. Table 1 shows mean percentage correct word recognition for each listener group on each trial. Mean recognition did not improve by more than 10% from the first to the third trial for any group, and all improvement was within one standard deviation of individual trial means. Consequently, mean results across the three trials were used in subsequent analyses.

Figure 2 shows mean correct word recognition for each listener group for each kind of stimulus. (For the Mandarin-speaking adults, there was no correlation between length of English experience and recognition scores.) Table 2 shows the results of one-way analyses of variance and post hoc comparisons performed on percentage correct word recognition scores. We computed all statistics using arc sine transforms because some listeners scored below 10% correct in some conditions.

Significant main effects of group were obtained for all three conditions; therefore post hoc comparisons between groups were examined. For the eight-channel AE stimuli, mean correct recognition for each group differed from that of every other group ($\eta^2 = .75$): English-speaking adults performed best, followed by English-speaking children, and lastly by Mandarin-speaking adults. For the four-channel AE stimuli, the children and Mandarin-speaking adults performed similarly, but both groups performed more poorly than the English-speaking adults ($\eta^2 = .41$). For the SW stimuli, the children and English-speaking adults performed the same, and both of those groups performed better than the Mandarin-speaking adults ($\eta^2 = .56$). So these between-groups results show that children and Mandarin-speaking adults performed more poorly than English-speaking adults with the AE stimuli, regardless of how many channels were used. From this result, we garnered support for the conclusion that amplitude

Table 1
Mean Percentage Correct Word Recognition for Each Trial in Each Condition, and Means Across Trials in Each Condition

Condition and trial	English adults	7-year-olds	Mandarin adults
Eight-channel amplitude envelope			
Trial 1	57.30 (8.55)	33.15 (9.18)	23.10 (9.41)
Trial 2	63.10 (6.34)	38.50 (8.32)	29.90 (10.25)
Trial 3	64.90 (5.64)	41.15 (8.39)	33.35 (10.42)
Mean	61.77 (6.39)	37.60 (7.94)	28.78 (9.58)
Four-channel amplitude envelope			
Trial 1	21.90 (8.66)	12.20 (8.18)	9.10 (4.46)
Trial 2	27.15 (9.65)	13.95 (8.19)	11.90 (6.37)
Trial 3	30.50 (10.25)	14.90 (8.81)	12.60 (7.21)
Mean	26.52 (9.20)	13.68 (8.20)	11.20 (5.77)
Sine wave			
Trial 1	24.60 (12.90)	22.05 (7.23)	6.75 (4.03)
Trial 2	31.45 (13.07)	26.80 (8.20)	9.25 (4.17)
Trial 3	33.45 (12.53)	28.10 (9.25)	10.35 (4.64)
Mean	29.83 (12.67)	25.65 (7.98)	8.78 (4.19)

Note. Standard deviations are shown in parentheses.

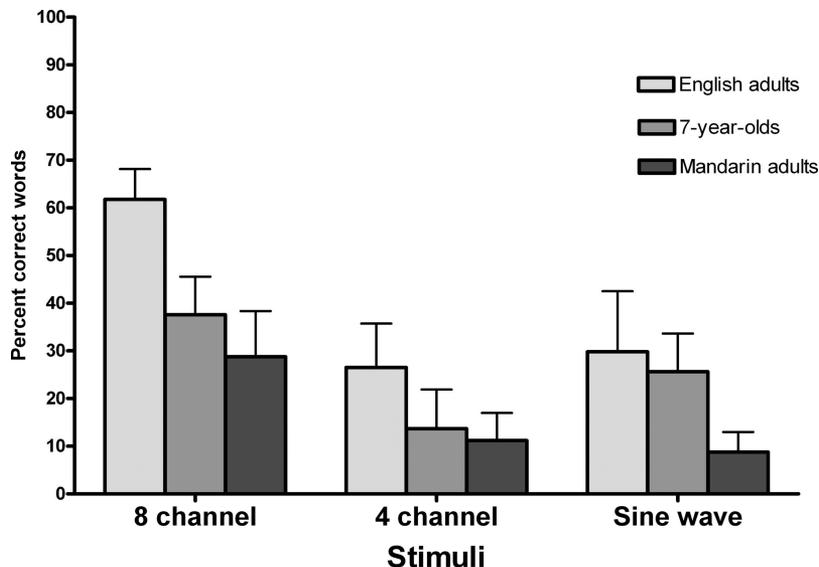


Figure 2. Percentage correct recognition for each of the three listener groups, by stimulus type. Error bars represent plus or minus one standard deviation.

structure differs across languages, and children must discover that structure in their native language. From results with the SW stimuli, we garnered support for the conclusion that global spectral structure also differs across languages, which suggests that children must discover that kind of structure as well. However, apparently they have already done so by 7 years of age.

Comparisons across stimulus types for each group were also examined, and recognition scores are replotted in Figure 3 to illustrate within-group comparisons for the three conditions. Table 3 shows results of *t* tests comparing performance across conditions for each listener group: For comparisons of four- and eight-channel AE stimuli, these were within-groups *t* tests; for comparisons of either

the four- or eight-channel AE stimuli with SW stimuli, these were between-groups *t* tests. Again, these analyses were performed using arc sine transforms. For listeners hearing the AE stimuli, performance was better with the eight-channel than with the four-channel stimuli, for all groups. When we look across AE and SW stimuli, we find for all three groups that listeners performed better with the eight-channel AE stimuli than with the SW stimuli. But for the comparison of four-channel AE and SW stimuli, children were the only group to show a difference in recognition scores: Children hearing the SW stimuli out-performed children hearing the four-channel AE stimuli.

Discussion

We designed this experiment to measure speech recognition for English-speaking adults, English-speaking children, and Mandarin-speaking adults, using sentences that were syntactically correct but semantically anomalous. Seven-year-old children learning English were found to recognize stimuli that preserved only the spectral skeleton of the speech signal as well as native English-speaking adults. Both the child and adult native speakers of English recognized these SW stimuli better than native Mandarin-speaking adults who had learned English as a second language. But when presented with signals that preserved only amplitude envelopes, the children performed similarly to the Mandarin-speaking adults, and both of these groups performed more poorly than the English-speaking adults. The finding that native Mandarin speakers who are fluent in English performed so poorly with both sorts of global signals suggests that variability in performance may be largely explained by differences in language experiences that would influence a listener’s ability to recover linguistic form from these signals. These results for native Mandarin speakers diminish the possibility that age-related differences observed by Eisenberg et al. (2000) and by us are explained by immature auditory systems in children. Of course, it is possible that Mandarin-speaking adults and English-speaking children

Table 2
Results of the Statistical Analyses for Group Effects, Using Arc Sine Transforms

Stimulus condition	<i>F</i> (2, 57) or <i>t</i> (57)	<i>p</i>
Eight-channel amplitude envelope	84.66	<.001
English-speaking adults vs. children	9.07	<.001
English- vs. Mandarin-speaking adults	12.61	<.001
Children vs. Mandarin-speaking adults	3.54	<.001
Four-channel amplitude envelope	20.07	<.001
English-speaking adults vs. children	4.95	<.001
English- vs. Mandarin-speaking adults	5.90	<.001
Children vs. Mandarin-speaking adults	0.94	<i>ns</i>
Sine wave	35.90	<.001
English-speaking adults vs. children	1.22	<i>ns</i>
English- vs. Mandarin-speaking adults	7.87	<.001
Children vs. Mandarin-speaking adults	6.65	<.001

Note. *F* ratios are for one-way analyses of variance, with group as the between-subjects factor; the *t* tests are for post hoc comparisons. Bonferroni corrections for three multiple comparisons require a *p* of less than .0033 for the comparison to be significant at the .01 level. All post hoc comparisons met that criterion.

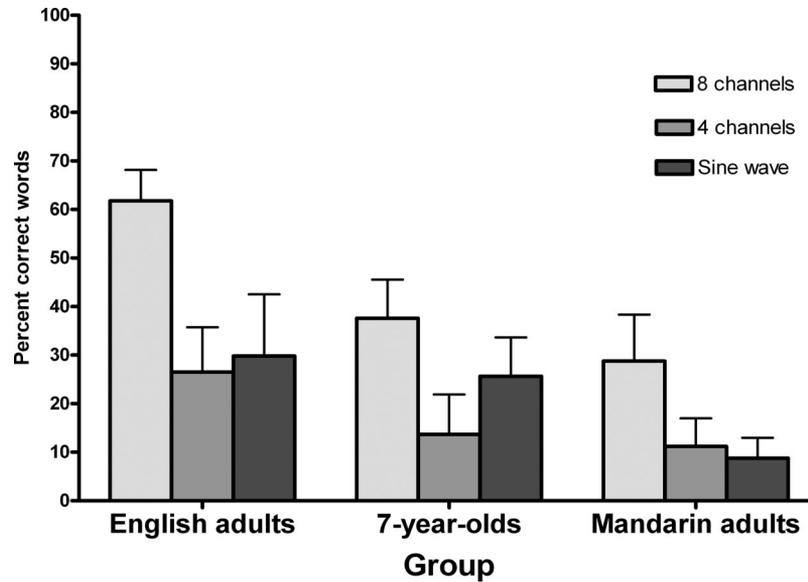


Figure 3. Percentage correct recognition for each of the three stimulus types, by listener group. Error bars represent plus or minus one standard deviation.

would have demonstrated similar results with the AE stimuli for different reasons (differences in language experience on the part of Mandarin-speaking adults and immature auditory systems on the part of children), but that seems unlikely. The principle of parsimony encourages us to adopt one common explanation for the results of both groups, and the finding that children performed similarly to adults for the SW stimuli indicates that children's auditory systems processed at least one kind of impoverished signal maturely.

The most significant implications of these results involve language development. By 7 years of age, these children demonstrated native skill in using spectral skeletons but not amplitude envelopes. Evidence is accumulating that children are not adultlike in their use of

the spectrally specific, temporally brief properties of the acoustic speech signal (i.e., acoustic cues) until much later, around the age of puberty (Hazan & Barrett, 2000). Taken together these results support the broader suggestion that global spectral structure plays a pivotal role in language acquisition. It may be that children are able to perceptually organize global spectral structure in the speech signal in order to start recovering linguistic form early in life, and the ability to do so facilitates the discovery of other sorts of structure by helping children divide the speech stream into separate linguistic units, such as words. Based on results of this study, we can only say that children acquired mature sensitivity to global spectral patterns in their native language by 7 years of age, but evidence from others suggests that children may learn about this spectral structure much earlier, possibly within the first year of life (Boysson-Bardies et al., 1986).

The focus of much speech perception research has traditionally been on discrete spectro-temporal properties, and methods of investigation have involved manipulating these properties in syllable-length signals to evaluate how those manipulations affect phonetic decisions. In accordance with that work, research on infant speech perception manipulated the same properties to see what infants could discriminate. (Labeling tasks cannot be used with infants.) This collective body of research has led to the perspective that infants are sensitive to the acoustic cues relevant in signaling phonetic segments, effectively from birth (e.g., Aslin, Pisoni, Hennessy, & Perey, 1981; Eimas, Siqueland, Jusczyk, & Vigorito, 1971; Kuhl, 1979). But infants rarely hear speech signals the length of syllables; instead, they hear speech in phrase or sentence length units. Because infants lack lexical knowledge, one early chore facing the child must surely be learning how to chisel out linguistically meaningful stretches of the signal. The question therefore arises of how children can identify stretches of the speech signal that form separate linguistic units, such as phrases within sentences or words within phrases. Regarding the first of these, several studies have shown that prosodic markers, such as pauses, are readily used by infants to parse the signal into clause-

Table 3

Results of the T Tests Done To Examine Stimulus Effects for Each Group, Using Arc Sine Transforms

Listener group	<i>t</i>	<i>p</i>
English-speaking adults		
Eight- vs. four-channel AE	18.60	<.001
Eight-channel AE vs. SW	9.44	<.001
Four-channel AE vs. SW	0.82	<i>ns</i>
Children		
Eight- vs. four-channel AE	10.01	<.001
Eight-channel AE vs. SW	4.61	<.001
Four-channel AE vs. SW	4.59	<.001
Mandarin-speaking adults		
Eight- vs. four-channel AE	8.70	<.001
Eight-channel AE vs. SW	9.21	<.001
Four-channel AE vs. SW	1.36	<i>ns</i>

Note. The test of eight- vs. four-channel AE stimuli involved within-group comparisons, and each had 19 degrees of freedom. The test of eight-channel AE vs. SW stimuli and of four-channel AE vs. SW stimuli involved two-group comparisons, and each had 38 degrees of freedom. AE = amplitude envelope; SW = sine wave.

length units (e.g., Hirsh-Pasek et al., 1987; Kemler Nelson, Hirsh-Pasek, Jusczyk, & Cassidy, 1989). However, speakers do not mark word boundaries with pauses, so the mystery remains of how infants learn to separate individual words from the ongoing speech signal. The suggestion made here is that infants could be using recurring patterns of global spectral structure to perceptually isolate and recognize the repeating strings. And although that suggestion is speculative at present, there are now two kinds of findings to motivate it: Young children rely on formant transitions more than adults for many phonetic decisions (e.g., Greenlee, 1980; Krause, 1982; Nittrouer, 1992; Nittrouer & Studdert-Kennedy, 1987; Wardrip-Fruin & Peach, 1984), and children recognize stimuli preserving only global spectral structure in an adultlike fashion sooner than they recognize stimuli preserving only global amplitude structure (this study) or discrete sequences of spectro-temporal structure (our current study compared with Hazan & Barrett, 2000). When combined, these findings suggest that global spectral structure may play an important role in early language learning. That structure could help infants begin to separate stretches of the ongoing speech signal that form linguistic units. There is already some evidence that infants use an acoustic property not related to individual phonetic segments to begin recovering separate words from the ongoing speech signal: Specifically, Mattys, Jusczyk, Luce, and Morgan (1999) found that 9-month-olds were more sensitive to a prosodic cue (i.e., stress) than to phonotactic constraints in recognizing word boundaries. Although prosody is a property arising largely from the laryngeal source of speech (while SW analogs preserve structure imposed by the filter), this finding nonetheless is one more piece of evidence that children can use global structure to begin recognizing word-length units. Once these units are extricated, children are better able to discover the details of those units (i.e., discrete spectro-temporal properties), forming more precise representations of each lexical item.

In conclusion, we undertook this experiment to examine children's abilities to recover linguistic form from acoustic speech signals processed in each of two ways that largely eliminate the spectro-temporal properties traditionally associated with phonetic structure. It is commonly known that processing algorithms such as those used in the current study preserve access to linguistic form for adult listeners who are native speakers of the language being examined. Results presented here indicate that children can recover linguistic form for processed signals that preserve global spectral structure as well as adult native speakers of the same language but cannot do so when only amplitude envelopes are preserved. Therefore, it is concluded that children discover the global spectral structure of their native language before the global amplitude structure. We presume that those spectral skeletons have heuristic value for language learning.

References

- Aslin, R. N., Pisoni, D. B., Hennessy, B. L., & Perey, A. J. (1981). Discrimination of voice onset time by human infants: New findings and implications for the effects of early experience. *Child Development*, *52*, 1135–1145.
- Beckman, M. E., & Edwards, J. (2000). The ontogeny of phonological categories and the primacy of lexical learning in linguistic development. *Child Development*, *71*, 240–249.
- Boothroyd, A., & Nittrouer, S. (1988). Mathematical treatment of context effects in phoneme and word recognition. *Journal of the Acoustical Society of America*, *84*, 101–114.
- Boysson-Bardies, B., Sagart, L., Halle, P., & Durand, C. (1986). Acoustic investigations of cross-linguistic variability in babbling. In B. Lindblom & R. Zetterstrom (Eds.), *Precursors of early speech* (pp. 113–126). New York: Stockton Press.
- Bricker, P. D., & Pruzansky, S. (1976). Speaker recognition. In N. J. Lass (Ed.), *Contemporary issues in experimental phonetics* (pp. 295–326). New York: Academic Press.
- Carrell, T. D. (1981). Effects of glottal waveform on the perception of talker sex. *Journal of the Acoustical Society of America*, *70*, S97.
- Charles-Luce, J., & Luce, P. A. (1990). Similarity neighbourhoods of words in young children's lexicons. *Journal of Child Language*, *17*, 205–215.
- Charles-Luce, J., & Luce, P. A. (1995). An examination of similarity neighbourhoods in young children's receptive vocabularies. *Journal of Child Language*, *22*, 727–735.
- Eimas, P. D., Siqueland, E. R., Jusczyk, P., & Vigorito, J. (1971, January 22). Speech perception in infants. *Science*, *171*, 303–306.
- Eisenberg, L. S., Shannon, R. V., Schaefer Martinez, A., Wygonski, J., & Boothroyd, A. (2000). Speech recognition with reduced spectral cues as a function of age. *Journal of the Acoustical Society of America*, *107*, 2704–2710.
- Fant, G. (1960). *Acoustic theory of speech production*. The Hague, Netherlands: Mouton.
- Fu, Q., Zeng, F. G., Shannon, R. V., & Soli, S. D. (1998). Importance of tonal envelope cues in Chinese speech recognition. *Journal of the Acoustical Society of America*, *104*, 505–510.
- Gleitman, L., & Wanner, E. (1982). Language acquisition: The state of the art. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 3–48). Cambridge, England: Cambridge University Press.
- Greenlee, M. (1980). Learning the phonetic cues to the voiced–voiceless distinction: A comparison of child and adult speech perception. *Journal of Child Language*, *7*, 459–468.
- Hazan, V., & Barrett, S. (2000). The development of phonemic categorization in children aged 6–12. *Journal of Phonetics*, *28*, 377–396.
- Hirsh-Pasek, K., Kemler Nelson, D. G., Jusczyk, P. W., Cassidy, K. W., Druss, B., & Kennedy, L. (1987). Clauses are perceptual units for young infants. *Cognition*, *26*, 269–286.
- Jassem, W. (1971). Pitch and compass of the speaking voice. *Journal of the International Phonetic Association*, *1*, 59–68.
- Jusczyk, P. W. (1997). *The discovery of spoken language*. Cambridge, MA: MIT Press.
- Kemler Nelson, D. G., Hirsh-Pasek, K., Jusczyk, P. W., & Cassidy, K. W. (1989). How the prosodic cues in motherese might assist language learning. *Journal of Child Language*, *16*, 55–68.
- Krause, S. E. (1982). Vowel duration as a perceptual cue to postvocalic consonant voicing in young children and adults. *Journal of the Acoustical Society of America*, *71*, 990–995.
- Kuhl, P. K. (1979). Speech perception in early infancy: Perceptual constancy for spectrally dissimilar vowel categories. *Journal of the Acoustical Society of America*, *66*, 1668–1679.
- Kuhl, P. K. (1987). Perception of speech and sound in early infancy. In P. Salapatek & L. Cohen (Eds.), *Handbook of infant perception: Vol. 2. From perception to cognition* (pp. 275–382). New York: Academic Press.
- Mann, V. A., & Repp, B. H. (1981). Influence of preceding fricative on stop consonant perception. *Journal of the Acoustical Society of America*, *69*, 548–558.
- Mattys, S. L., Jusczyk, P. W., Luce, P. A., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology*, *38*, 465–494.
- Mayo, C., Scobbie, J. M., Hewlett, N., & Waters, D. (2003). The influence of phonemic awareness development on acoustic cue weighting strategies in children's speech perception. *Journal of Speech Language and Hearing Research*, *46*, 1184–1196.
- Morgan, J. L., & Demuth, K. (1996). *Signal to syntax*. Mahwah, NJ: Erlbaum.
- Nittrouer, S. (1992). Age-related differences in perceptual effects of formant transitions within syllables and across syllable boundaries. *Journal of Phonetics*, *20*, 351–382.

- Nittrouer, S. (2001). Challenging the notion of innate phonetic boundaries. *Journal of the Acoustical Society of America*, *110*, 1581–1597.
- Nittrouer, S. (2006). Children hear the forest. *Journal of the Acoustical Society of America*, *120*, 1799–1802.
- Nittrouer, S., & Miller, M. E. (1997a). Developmental weighting shifts for noise components of fricative-vowel syllables. *Journal of the Acoustical Society of America*, *102*, 572–580.
- Nittrouer, S., & Miller, M. E. (1997b). Predicting developmental shifts in perceptual weighting schemes. *Journal of the Acoustical Society of America*, *101*, 2253–2266.
- Nittrouer, S., & Studdert-Kennedy, M. (1987). The role of coarticulatory effects in the perception of fricatives by children and adults. *Journal of Speech and Hearing Research*, *30*, 319–329.
- Remez, R. E., Rubin, P. E., Pisoni, D. B., & Carrell, T. D. (1981, May 22). Speech perception without traditional speech cues. *Science*, *212*, 947–949.
- Repp, B. H. (1982). Phonetic trading relations and context effects: New experimental evidence for a speech mode of perception. *Psychological Bulletin*, *92*, 81–110.
- Shannon, R. V., Zeng, F. G., Kamath, V., Wygonski, J., & Ekelid, M. (1995, October 13). Speech recognition with primarily temporal cues. *Science*, *270*, 303–304.
- Soderstrom, M., Seidl, A., Kemler Nelson, D. G., & Jusczyk, P. W. (2003). The prosodic bootstrapping of phrases: Evidence from prelinguistic infants. *Journal of Memory and Language*, *49*, 249–267.
- Stevens, K. N. (1975). The potential role of property detectors in the perception of consonants. In G. Fant & M. A. A. Tatham (Eds.), *Auditory analysis and perception of speech* (pp. 303–330). New York: Academic Press.
- Stevens, K. N., & Blumstein, S. E. (1978). Invariant cues for place of articulation in stop consonants. *Journal of the Acoustical Society of America*, *64*, 1358–1368.
- Sussman, J. E. (1993). Auditory processing in children's speech perception: Results of selective adaptation and discrimination tasks. *Journal of Speech and Hearing Research*, *36*, 380–395.
- Wardrip-Fruin, C., & Peach, S. (1984). Developmental aspects of the perception of acoustic cues in determining the voicing feature of final stop consonants. *Language and Speech*, *27*, 367–379.
- Watson, J. M. M. (1997). *Sibilant-vowel coarticulation in the perception of speech by children with phonological disorder*. Unpublished doctoral dissertation, Queen Margaret College, Edinburgh, United Kingdom.
- Werker, J. F., & Yeung, H. H. (2005). Infant speech perception bootstraps word learning. *Trends in Neurosciences*, *9*, 519–527.

Appendix

Practice Sentences

1. Cooks run in brooms.
2. Ducks teach sore camps.
3. Find girls these clouds.*
4. Gangs load near sweat.
5. Great shelf needs tape.
6. Throw his park hand.

Test Sentences

1. Blue chairs speak well.*
2. Cars jump from fish.
3. Cats get bad ground.
4. Cups kill fat leaves.*
5. Drive my throat late.
6. Dumb shoes will sing.*
7. Fan spells large toy.
8. Feet catch bright thieves.
9. Green hands don't sink.
10. Hard checks think tall.
11. Jobs get thick hay.
12. Knees talk with mice.*
13. Late forks hit low.*
14. Late fruit spins lakes.
15. Lead this coat home.*

16. Lend them less sleep.
17. Let their flood hear.
18. Paint your belt warm.
19. Pink chalk bakes phones.*
20. Sad cars want chills.*
21. Slow dice buy long.
22. Small lunch wipes sand.
23. Soap takes on dogs.
24. Socks pack out ropes.
25. Soft rocks taste red.
26. Suits burn fair trail.
27. Teeth sleep on doors.
28. Thin books look soft.
29. Trucks drop sweet dust.
30. Wide pens swim high.

Note. Asterisks indicate sentences from "Mathematical Treatment of Context Effects in Phoneme and Word Recognition," by A. Boothroyd & S. Nittrouer, 1988, *Journal of the Acoustical Society of America*, *84*, p. 113. Copyright 1988 by the Acoustical Society of America.

Received July 24, 2007
 Revision received July 3, 2008
 Accepted October 28, 2008 ■