Journal of Experimental Psychology:
Human Perception and Performance
1983, Vol. 9, No. 5, 739–752

Category-Boundary Effects and Speeded Sorting With a Harmonic Musical-Interval Continuum: Evidence for Dual Processing

Robert J. Zatorre
Brown University

In the first experiment, a continuum of 10 harmonic musical intervals was constructed from a minor to a major third. Four pairs of stimuli with constant physical distances were presented to seven musicians in a two-interval forced-choice discrimination task. Either silence, an interfering tone, or a noise burst was interposed between the two stimuli in a pair. Unbiased discriminability was found to be consistently higher for pairs straddling the boundary between two categories than for the endpoint pairs. The interfering tone lowered overall discrimination but left the shape of the function unchanged, whereas the noise burst had no effect. Experiment 2 used a similar paradigm, but the continuum consisted of the single tone that had cued the minor–major distinction for intervals. Discrimination of this series did not show consistent changes as a function of continuum position. In Experiment 3, triads that varied in either interval or overall pitch were presented to musicians for sorting according to one dimension or another. The result was that there were much longer latencies to sort according to interval when pitch varied irrelevantly than vice versa. These results demonstrate that there are changes in discriminability associated with learned categories and suggest that there may be two hierarchically organized stages. A dual-processing model is discussed in which the listener has available both auditory and categorical information.

In a now classic article, Miller (1956) pointed out that identification of unidimensional stimuli was closely linked to the capacity of short-term memory; however, in the realm of speech perception an exception to this formulation was reported. Liberman, Harris, Hoffman, and Griffith (1957) found that listeners were not able to perceive graded differences between stimuli but rather tended to hear only discrete changes between one phonetic category and another. This phenomenon became known as *categorical perception,* and was first believed to be unique to speech processing (cf. Eimas, 1963; Liberman, Harris, Kinney, & Lane, 1961; Mattingly, Liberman, Syrdal, & Halwes, 1971). Later versions (Fujisaki & Kawashima, 1969, 1970, 1971; Pisoni, 1973, 1975) added an auditory-processing stage to explain listeners' ability to discriminate within-category stimuli at above-chance levels.

The hypothesis of a special speech-processing mechanism has been criticized in recent years. One important critique is that of Macmillan, Kaplan, and Creelman (1977), who pointed out that (a) most research into categorical perception has not used bias-free measures of discriminability and (b) in most speech-perception studies, identification ability may be underestimated by providing too few categories. Thus, discrimination that is higher than predicted might be explained by the listener's ability to further categorize the stimuli rather than by the existence of an auditory-processing stage (cf. Liberman, Harris, Eimas, Lisker, & Bastian, 1961).

Some of these experiments formed part of a doctoral dissertation submitted to the Department of Psychology, Brown University, October 1980. Some of this material was presented at the 102nd meeting of the Acoustical Society of America, Ottawa, Canada, May 1981.

This research was supported in part by National Institute of Child Health and Human Development Grant HD 05331 to P. Eimas and National Science Foundation Grant BNS 75-08439 to R. Millward, to whom I am grateful for allowing me the use of laboratory facilities.

I wish to thank my advisor, Peter D. Eimas, for his assistance, guidance, and support and Francis Ganong for his patience and expertise. Helpful comments were also provided by Sheila Blumstein.

Requests for reprints should be sent to Robert J. Zatorre, who is now at the Montreal Neurological Institute, 3801 University Street, Montreal, Quebec, Canada H3A 2B4.

A second critique of categorical perception has emerged from work with nonspeech analogues of speech sounds. When certain nonspeech stimuli are presented under conditions comparable to those used in categorical-perception experiments, results very similar to those that had been labeled as "categorical" may be obtained (Cutting & Rosner, 1974; Cutting, Rosner, & Foard, 1976; Miller, Wier, Pastore, Kelly, & Dooling, 1976; Pastore, Ahroon, Baffuto, Friedman, Puleo, & Fink, 1977; Pisoni, 1977; Raz & Brandt, 1977).[1]

Yet another set of stimuli that has yielded categorical-like results is musical intervals, and they form the focus of the present investigations. Locke and Kellar (1973), Blechner (1977), and Siegel and Siegel (1977) reported that when a continuum is constructed between one musical category and another (e.g., minor and major thirds), musically trained subjects were able to parse the stimuli into discrete categories. Furthermore, discrimination data showed a peak near the boundary between these categories.

Other studies on musical-interval perception have extended the earlier findings. Burns and Ward (1978) used temporally successive (melodic) intervals whose absolute frequency varied over a small range so that only the relative frequency change between two notes of an interval could be used as a cue (and not the absolute frequencies of the individual notes). Categorical effects were clearly elicited under these conditions when using a two-interval forced-choice procedure. However, when an adaptive psychophysical paradigm was used for discrimination, only a small amount of training was found sufficient to eliminate the peaks in discrimination previously noted. On the other hand, when these same listeners were retested on fixed step-size discrimination, the peaks near the boundary were found once more. These results led Burns and Ward to conclude that categorical perception is related to the degree of stimulus uncertainty associated with the procedure used. With high-uncertainty situations, the conditions for occurrence of categorical perception are maximized.

Another study on musical-interval perception was that of Zatorre and Halpern (1979), who used simultaneous (harmonic) intervals. In standard identification and discrimination tasks the results once more showed musicians perceiving in a manner similar to that termed categorical for speech, whereas nonmusicians did not generally show these effects. In this study, discrimination was consistently much higher (by approximately 25%) than predicted on the basis of the identification function with the strict categorical prediction, unlike the functions obtained by Burns and Ward, which were generally well predicted by the same formula. Zatorre and Halpern argued that the good performance on discrimination was partly due to the availability of auditory information. It was possible that subjects based their discriminations partly on differences between the top notes of the interval pairs (in addition to interval size) because the bottom note was held constant. However, when a task was constructed in which the absolute pitch level varied randomly from stimulus to stimulus (paralleling the procedure used by Burns and Ward, 1978, with melodic intervals), Zatorre and Halpern's subjects were unable to consistently identify or discriminate the stimuli. This result was attributed to the lack of contextual cues for subjects to base their judgments on. The discrepancy between performance on melodic and harmonic intervals remains to be explained, however.

The experiment that follows was designed to investigate the existence of two sources of information in the discrimination of harmonic musical intervals. One source corresponds to simple pitch information, whereas the other corresponds to the availability of category coding, in this case major or minor. I postulated that if an auditory-processing stage exists, it should be possible to interfere selectively with it and thereby reduce overall discrimination without affecting the peak. A further issue concerns clarifying the relationship between identification and discrimination. A final aim of this experiment was to determine if categorical effects hold up when the data are analyzed in terms of signal-detection theory, an important issue in light of the criticisms made by Macmillan et al. (1977).

---

[1] Note, however, that the results using rise time have recently been called into question (Kewley-Port & Pisoni, 1982; Rosen & Howell, 1981).

Experiment 1

Method

Subjects. Seven musically trained persons from the Providence area were used as subjects. Their mean age was 24, and the mean number of years they spent studying and playing an instrument was 13.5. No subject reported possession of absolute pitch or any history of hearing disorders.

Stimulus materials. Stimuli were produced on a PDP-8 computer. A continuum of two simultaneously presented pure tones was constructed. The lower tone was constant and corresponded to the note G (392.0 Hz); the top tone of the continuum varied in 10 logarithmically equal steps (11.1 cents each) from B flat (466.2 Hz) to B natural (493.9 Hz). The continuum thus covered the range from minor third (GB flat, 300 cents) to major third (GB, 400 cents). For convenience, the stimuli in the continuum are numbered from 1 to 10; 1 corresponds to the 300-cent stimulus, 2 to the 311.1-cent stimulus, and so forth. Each two-note interval had a total duration of 500 msec. Amplitude was sent to the same nominal parameter for production of each stimulus; presentation to subjects was at 75 dB SPL, A scale, as measured at the headphones with a sound pressure meter.

In addition to the continuum just described, two types of interference stimuli were constructed. The tone interferers consisted of 500-msec pure tones at a frequency 133.3 cents above the midpoint frequency of the top tones of the two intervals being discriminated; thus it was at a constant distance for each pair of intervals. The distance of 133.3 cents was chosen on the basis of results by Deutsch (1972), which showed maximal interference to occur in this range in a tonal memory task. The other interference stimulus consisted of 500 msec of broad-band white noise. All stimuli were band-pass filtered between 100 and 1000 Hz with a Krohn-Hite 3550 analog filter. A comparison of the spectral characteristics of the noise burst and the tones showed that the overlap for the two types of stimuli was almost complete. However, the noise burst was about 10 dB louder than the tones when measured at the headphones.

Procedure. For the discrimination task, four pairs of stimuli were chosen from the 10-stimulus continuum: two within-category and two between-category. Within-category and between-category pairs were chosen simply on the basis of proximity to the endpoints (within) or proximity to the midpoint of the continuum (between). The chosen pairs were always a constant distance apart on the continuum (i.e., the ratios of the top notes of the two intervals were constant at 22.2 cents).[2] Pairs 1-3 and 8-10 were considered to be within-category; 4-6 and 5-7 were between-category pairs.

The discrimination task itself was the two-interval forced-choice (2IFC) method.[3] On each trial two different stimuli from one of the four pairs previously described were presented in one or another order; the subjects' task was to determine which of the two stimuli sounded "more minor," and to respond accordingly by pressing one of two keys on a teletype in front of them (Key 1 if the first stimulus was more minor, Key 2 if the second was more minor). Each pair was presented 15 times per block of trials for three blocks of trials, with short rests between blocks. Order of stimulus presentation was counterbalanced (so that Pairs 1-3 and 3-1 appeared 15 times each, etc.). All eight possible pairs were presented in a pseudorandom order of 120 trials per block, for total of 360 trials per experimental session.

In Condition 1 (no interference) there were 1.5 sec of silence interposed between the two stimuli in the discrimination pair. In Condition 2 (tone interference) the interference tone previously described was placed between the two intervals to be discriminated; this tone started 500 msec after the offset of the first interval and ended 500 msec before the onset of the second interval. Subjects were told to disregard the interfering tone as much as possible and to respond on the basis of the first and last signals presented. In Condition 3 the interfering stimulus was the noise burst previously described; subjects were told to disregard this interferer also.

Subjects were tested individually or in groups of three or fewer in a quiet room. The three conditions were run in the order described for all seven subjects, with at least 1 week between sessions. On each testing day listeners were given five trials per discrimination pair as practice before the actual test began.

Once the three sessions were finished, subjects returned for an identification task. Only six of the original seven subjects were available for this session. All 10 intervals in the continuum were presented in a single random order of 20 repetitions each. A 6-point rating scale was used for identification, and subjects were told to distribute their responses approximately equally among the 6 categories (where 1 was "most minor," 6 was "most major"), and the remaining categories corresponded to different degrees of major and minor. Again, 5 repetitions of each stimulus were given as practice before each session. Subjects were tested interactively with the PDP-8 in both identification and discrimination tasks; accordingly, the time interval between trials varied depending on the speed of response. However, there was always at least 1 sec between the last response and the onset of the next trial.

Results

Discrimination ability was measured by calculating an unbiased sensitivity index. As

---

[2] Minimum discriminability for pitch is difficult to evaluate because it varies significantly with the nature of the stimuli, tasks, and subjects' training. However, the average deviation for adjustment of simultaneous intervals has been reported to be in the vicinity of 20 cents, depending on the interval (Moran & Pratt, 1926; Rakowski, 1976; Ward, 1954). Because the fixed discrimination distance in this experiment was 22.2 cents, it follows that performance should be in the appropriate range to be sensitive to experimental manipulations.

[3] Note that this procedure differs from the more common discrimination tasks used in speech research; it was chosen on the basis of Macmillan et al.'s advice in that it is the most straightforward way of determining sensitivity and because it is relatively robust with respect to inequality of variance in the underlying distributions.

suggested by Macmillan et al. (1977), for a 2IFC task this index equals $d'_{2IFC}/\sqrt{2}$; this index corresponds to the d' that would be obtained in a yes–no detection paradigm and is referred to hereafter simply as d'.[4]

The principal result is presented in Figure 1, which shows mean d' scores calculated separately for each subject then averaged across subjects for each condition. The most striking effect is the large increase in discriminability for the two between-category pairs as compared to within-category pairs. An analysis of variance (ANOVA) on these scores shows a highly reliable effect of position on the continuum, $F(3, 18) = 19.70$, $p < .001$. Post hoc tests using Tukey's honestly significant difference (HSD) procedure indicated that only the differences between the two between-category and the two within-category pairs differed significantly ($p < .05$).

The second finding of interest is that Interference Condition 2 (tone) lowered overall performance but left the peak in the function unchanged relative to the troughs (see Figure 1); in other words, the functions for Conditions 1 and 2 remained parallel. Condition 3 (noise) had no noticeable effect on discrimination. The analysis showed a main effect of condition,

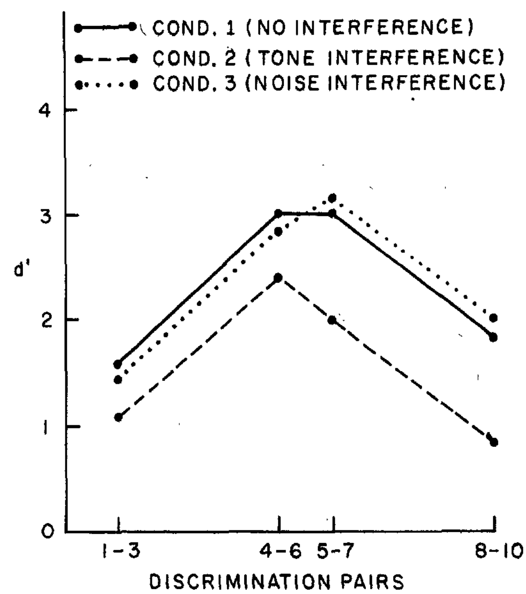

Figure 1. Mean d' scores for seven musicians in the discrimination task of Experiment 1 as a function of position on the continuum and interference conditions.

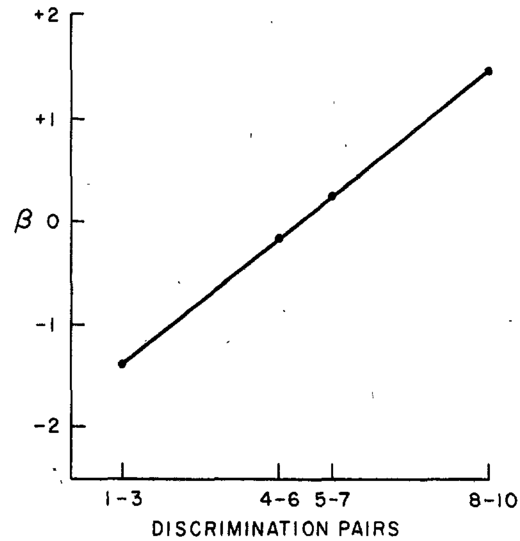

Figure 2. Mean response bias scores for discrimination data of Experiment 1, Condition 1 for seven listeners as a function of continuum position. (A negative value of beta corresponds to a bias to respond by pressing Key 2; a positive value corresponds to a bias to respond by pressing Key 1.)

$F(2, 12) = 12.70$, $p < .005$, but no interaction. Post hoc tests confirmed that only the difference between Condition 2 and the others was significant. The average decrement in d' between Conditions 1 and 2 was .76.

Because the signal-detection theory allows the separation of sensitivity from response bias, it is of interest to examine bias in this task. Figure 2 shows the mean beta scores for the seven subjects in Condition 1 of the discrimination task. A negative value of beta corresponds to a bias to respond by pressing Key 2 (i.e., that the second stimulus was more minor than the first), whereas the opposite is true of a positive bias score. There was a large change in beta as a function of position on the continuum, $F(3, 18) = 11.08$, $p < .01$.

[4] For several subjects, $p(H)$ and/or $p(FA)$ reached 1.00 and .00, respectively, where $p(H)$ and $p(FA)$ are the proportion of hits (H) and false alarms (FA). In such cases, d' is indeterminate; however, following the suggestion made by Macmillan et al., it was assumed that proportions of 1 and 0 are experimentally indistinguishable from .995 and .005, respectively, thereby giving a maximum d' score of 3.65. It should be kept in mind that these estimates of discriminatory ability are likely to be less reliable than others.

Post hoc tests showed all means to be significantly different from one another except for the values for the two between-category comparisons. Changes in response bias for Conditions 2 and 3 were essentially identical to those for Condition 1.

The next important result concerns the relation between identification and discrimination. The identification was monotonically increasing, but there were larger changes in ratings near the middle of the continuum than at the endpoint regions, $F(3, 15) = 16.03, p < .001$, and post hoc tests confirmed that the within-category differences were indeed smaller than the between-category differences.

Discrimination data were predicted from the identification data in two ways. First, two-category identification functions were used to predict discrimination on the basis of the strict categorical hypothesis that discrimination is totally limited by identification (for simplicity, this is referred to as the "Haskins prediction").[5] Two-category identification functions were obtained from the rating data by collapsing Responses 1, 2, and 3 into one category and Responses 4, 5, and 6 into the other. The second way the identification data were used to predict discrimination was by the method outlined by Braida and Durlach (1972) and advocated by Macmillan et al. (1977). With this procedure a d' score is obtained from pairs of stimuli according to how well the subject was able to use different ratings for the pair of stimuli in question. This is referred to as the signal-detection prediction.

Because the two predictions use different dependent measures, direct comparison is problematic. However, the predicted proportion correct from the Haskins formula may be converted to a d' score for comparison with the signal-detection prediction.[6] This function is shown in Figure 3 along with the obtained d' scores and the signal-detection prediction. When all three sets of values were analyzed it was found that the effects of obtained versus predicted, $F(2, 10) = 35.3, p < .001$, and continuum position, $F(3, 15) = 24.8, p < .001$, were significant, but there were no interaction effects. Post hoc tests showed that both predictions significantly underestimated obtained performance; however, the Haskins prediction underestimated true discrimination more than the signal-detection model did.

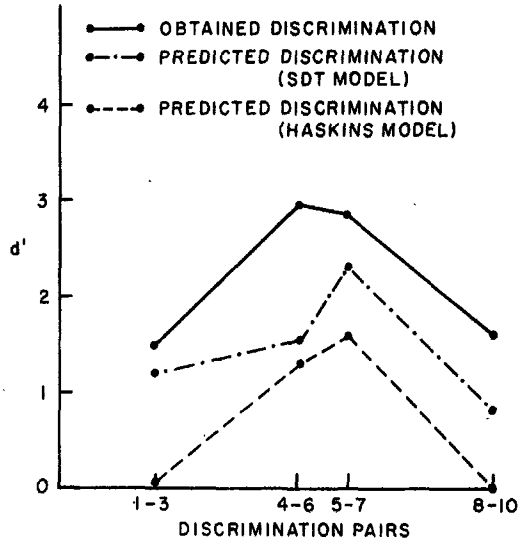

Figure 3. Obtained and predicted discrimination data for six subjects in Experiment 1, Condition 1. (Predicted scores are based on [a] the signal-detection model, which calculates the d' corresponding to a detection task from rating data in identification and [b] a d' score derived from the Haskins formula, which assumes discrimination to be totally limited by identification.)

Discussion

These results support a dual-processing model. In this view, when an interfering tone is presented between the two intervals to be discriminated, it interferes with the internal auditory representation of the stimulus but does not affect memory for the binary variable that results from categorization. The outcome is a decrement in performance but no drop in the peak relative to the trough because categorical information is still available to make decisions even after interference. When no interference is present, the listener makes use of both sources of information for the between-category comparisons but retains only auditory

[5] The formula used to calculate predicted discrimination was $p(D) = .5\,[p1(A)p2(A) + p1(B)p2(B)] + [p1(A)p2(B) + p1(B)p2(A)]$, where $p1(A)$ is the probability that Stimulus 1 is identified as A, $p2(A)$ is the probability that Stimulus 2 is identified as A, and so forth.

[6] Total proportion correct is equal to $[p(H) + p(CR)]/2$, where $p(CR)$ is the proportion of correct rejections (CR). Using the formula for response bias, beta equals $z[p(H)] + z[p(FA)]$; and the fact that $p(CR) + p(FA) = 1$ allows the calculation of d' if the value corresponding to beta is known.

information as the basis for within-category discrimination.

The fact that noise interference had essentially no effect is important in that it points to a specialization of the auditory-processing stage. The noise burst, although spectrally overlapping with the intervals and somewhat louder, was apparently not processed by the hypothetical auditory memory system. It can at least be said that the memory trace for a musical interval is not disrupted by all auditory stimuli. Such a conclusion is also in agreement with Deutsch's (1970) finding that memory for tones is not disrupted by concurrent memorization of words.[7]

This dual-processing model might still be criticized on the basis of the arguments outlined by Macmillan et al. (1977). In particular, it might be claimed that categorical perception did not occur because discrimination was not limited by identification; rather, it was consistently better than predicted by a more or less equal amount throughout the continuum, a result typical of continuous perception. Because identification functions did not predict discrimination very well, it might be wiser to call the effect of category-boundary effect (cf. Wood, 1976). However, accepting the definition of categorical perception given by Macmillan et al. (equivalence of identification and discrimination) would fail to distinguish the discrimination functions found here from functions without peaks at the boundary.

There appear to be two important effects to account for theoretically: (a) the peak in discrimination and (b) the higher than predicted results. Macmillan et al. noted the possibility of just this combination of results and claimed that it can be totally explained by a unidimensional model with inequality of stimulus spacing along a psychological dimension. They criticized the dual-processing explanation by the following analogy:

> The trouble with this model lies in its embarrassing lack of parsimony: It is equivalent to a simpler model in which the existence of the boundary is ignored. (The situation is analogous to deciding which of two weights is heavier by first comparing each with a 1-kg standard. If one weight is heavier than the standard and the second lighter, then choose the first; if not, put the two weights themselves on opposite arms of a pan balance. Clearly, the use of the standard weight does not affect the outcome of this two-step decision process.) (Macmillan et al., 1977, p. 467)

This analogy is invalid, however, because it does not reflect the task in question; it assumes performance is constrained only by sensory noise and not memory noise.

Let us consider the following modification of the example. Two weights are presented in succession, each of which is placed on a scale that records weight. Under these conditions the observer need only compare the two values to decide which is heavier; therefore, the observer will be constrained by the scale's resolution (sensory noise) and his or her memory for the two values to be compared (memory noise). Now let us further assume that our subject, a trained weightlifter, is familiar with the values along the scale that correspond to certain standard weights. There will now be two sources of information for him to use: his memory for the values on the scale as well as knowledge about the weights' proximity to a standard value. He should therefore perform much better under these circumstances, and a dual-processing theory would predict a peak in discrimination at values close to the midpoint between two standards.

The applicability of this example to the present results should be obvious. Note that this model involves a change from previous dual-processing models (e.g., Fujisaki & Kawashima, 1971; Pisoni, 1973), since in those models phonetic information only is used for between-category discrimination. In the current formulation, auditory information is available for all stimulus comparisons. If auditory information were not being used in between-category discrimination, then the peak should have remained unchanged after interference. Because the peak fell by about the same amount as the within-category points, the amount of auditory information available must have been approximately constant across the continuum.

This model also gives a plausible explanation for the observed changes in response bias.

[7] Massaro (1970) reported that noise was an effective interference stimulus in a tonal memory task, although tones provided more interference. The discrepancy between that finding and the current one can probably be accounted for by task and stimulus variables. It is of interest to speculate, however, that musicians may have learned to "tune out" irrelevant noise when making musical judgments, which would account for the lack of interference effect in the case of the noise stimulus.

In the cases where two stimuli drawn from the minor category are presented, the listener perceives both as minor but can still discriminate them to some extent. If we assume that the subject is more certain about the second stimulus presented than the first, he or she will tend to respond by pressing Key 2 more of the time when unsure because the subject knows the second stimulus was definitely a minor interval. Conversely, when two major stimuli are presented, the subject will be certain that the second stimulus, at least, was major and so will tend to say that the first must have been more minor, thereby biasing the responses toward pressing Key 1. When the stimuli come from separate categories there should be no consistent bias because on the average the second stimulus is major as often as minor.

Experiment 2

The results from Experiment 1 suggest that changes in discrimination are associated with learned categories. There is a confounding factor, however, in that the peak in discrimination always occurs near the midpoint of the continuum. Thus, it may be that range-related contextual effects are playing a role. The listeners might be simply bisecting the range, for example: such effects must be separated from any effects due to the presence of categories or boundaries.

One approach to this problem is to examine perception of the isolated minimal acoustic cue necessary for a distinction between two categories, in this case the pitch of the top note of the interval. If range-related effects are responsible for the discrimination peak, then it should not matter whether the tones are presented alone or as part of a musical interval. Also, a possible criticism of Experiment 1 is that no evidence was given that the task actually compelled the subjects to the use the categories they had presumably learned. Thus, it is theoretically possible that the results obtained in the first experiment would have been obtained even in the absence of interval information. A secondary purpose of Experiment 2, therefore, was to answer this potential criticism.

Blechner (1977) performed an experiment in which chords or single tones were presented to musicians for discrimination. The basic result was that single tones did not produce reliable peaks; however, in one condition a slight peak was found, though it was not in the place predicted by the identification function. Musical chords, on the other hand, always produced large peaks near the boundary. Blechner's experiment, however, suffers from the sorts of difficulties described by Macmillan et al., and it therefore appeared useful to attempt a study similar to Experiment 1 using a continuum of single tones. Thus, the acoustic cues are held constant across conditions (interval or single tone), but in the case of single tones the sounds are not categorizable as a major or minor.

Method

Subjects. Seven musically trained listeners participated, two of whom had already taken part in Experiment 1. Other subject characteristics were the same as in Experiment 1.

Stimulus materials. The stimulus continuum was constructed in precisely the same manner as in Experiment 1 except that the lower tone was omitted. This resulted in a single-tone continuum from B flat to B natural in 10 logarithmically equal steps. The same interference tones as in Experiment 1 were also used in one condition of this experiment.

Procedure. To provide comparison data and for purposes of replication, I first tested the five new subjects on Condition 1 from Experiment 1. Next, the discrimination procedure from Experiment 1 was repeated but using only the single-tone continuum without interference. Finally, single-tone discrimination proceeded with a tone interference as in Condition 2 of Experiment 1. This latter condition was included because pilot results indicated a ceiling effect with the single-tone continuum. In all other respects the discrimination procedure was exactly as in Experiment 1.

Results

The results are shown in Figure 4 as mean d', which was calculated separately for each subject and then averaged. The first result of note is the replication of the discrimination peak for the interval continuum. For the single-tone continuum it is notable that, without interference, this discrimination is much easier than for the comparable interval continuum. As for the presence of peaks, the situation is less clear. Without interference there is certainly no peak. In the interference condition there was a good deal of individual variation, with some subjects showing distinct peaks and others a totally flat function.

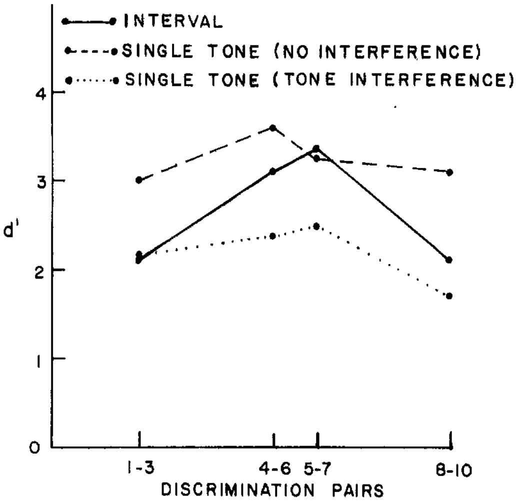

Figure 4. Mean d' scores for seven musicians for the three conditions of Experiment 2 as a function of position on the continuum.

An ANOVA was performed on the d' values with position on the continuum as one factor and condition (interval, single note, single note with interference) as the second factor. Significant results were observed for both main effects as well as the interaction: condition, $F(3, 18) = 12.53$, $p < .005$; continuum, $F(2, 12) = 12.06$, $p < .005$; interaction, $F(6, 36) = 4.66$, $p < .005$. Post hoc tests indicated significantly higher discrimination for the two middle pairs in the case of the interval continuum, as in Experiment 1. No significant differences were found for the single-tone continuum without interference, but this is not surprising as many scores were near ceiling on this task. For the single-tone continuum with interference it was found that the d' value from the high end of the continuum (Pair 8-10) was significantly lower ($p < .05$) than for Pair 5-7; the other three points did not differ from one another.

Discussion

The single-tone discrimination data clearly differ from the interval discrimination results, and in this respect the findings parallel Blechner's (1977) results. The results of a study by Kopp and Livermore (1973) are also relevant here. They gave nonmusicians a single-tone continuum for identification and discrimination. Analysis via signal detection showed large changes in response bias near the boundary but no peak in true discriminability.

There is one troubling feature of the present results, however: Although there is no statistically reliable peak as such, in the single-tone-with-interference condition there is one data point at the end of the continuum that differs from one near the middle. Inspection of the individual results suggests that for a few listeners, at least, discrimination was indeed better for stimuli from the midpoint of the continuum. Yet every subject demonstrated substantially larger differences between the middle stimuli and the extremes for the intervals than for the single tones (i.e., the functions were always more peaked for the interval continuum). Thus, it appears that range-related effects may contribute a small component to the peak in discrimination performance. Nevertheless, the peaks observed with intervals are of greater magnitude as well as greater reliability than those seen with single tones. Therefore, it is reasonable to conclude that discrimination peaks with intervals are not due simply to bisecting the range or other contextual effects but rather are more likely primarily related to the existence of categories and boundaries.

Another interesting aspect of these results is that single-tone discrimination without interference appears to be much better than the equivalent interval discrimination, even though the same frequency ratios were used in both cases (e.g., one subject did not miss a single trial out of 450 in this condition). It is possible that the generally poorer discrimination when simultaneous intervals were presented may have been due to some sort of masking of the upper tone by the lower one.

Experiment 3

This experiment was designed to gather further support for a distinction between two levels of processing in the perception of musical intervals by using the speeded-classification paradigm. Wood (1975) has shown that two levels of processing can be distinguished by such a task. The assumption is that variation of an irrelevant stimulus dimension will interfere with speeded sorting only if the irrelevant dimension must first be processed by a lower order mechanism. Alternatively, if the two stages are totally independent of one another, then no interference would be expected.

Wood (1975) studied fundamental frequency and place of articulation of consonant-vowel syllables. When subjects were to classify just two of the four possible stimuli on one of the dimensions (control conditions), reaction times (RTs) were equivalent. When all four possible stimuli were combined in a block of trials (orthogonal condition), classification according to pitch was just as quick as in the control condition, but classification of place of articulation was slower by about 50 msec. Evoked potentials recorded concurrently with the sorting task also supported these results. These results were interpreted as evidence that a specialized mechanism was brought into play when phonetic analysis was required and that lower order information (fundamental frequency) was extracted prior to the phonetic identification. Wood also reported that this pattern of results (asymmetric interference) did not occur with two auditory dimensions, such as pitch and intensity.

Further experiments with various other dimensions such as place and manner of articulation, vowel quality, and intensity (Eimas, Tartter, Miller, & Keuthen, 1978; Miller, 1978; Miller, Eimas, & Zatorre, 1979) have shown various complex interactions, which appear to reflect the different degrees of dependency of various levels of processing (for a review see Eimas, Tartter, & Miller, 1981).

As might be expected, the so-called phonetic-processing results have also been found for certain nonspeech stimuli. Blechner, Day, and Cutting (1976) found that variation in intensity interfered with classification of rise time but not vice versa. Pastore, Ahroon, Puleo, Crimmins, Golowner, and Berger (1976) reported a similar pattern for stimuli that varied in tone-buzz relative onset time. Once more, then, the auditory-phonetic distinction seems to rely not so much on the presence of phonetic processing per se but rather, as Blechner et al. (1976) put it, "on the coding of sounds within a hierarchically organized system, or on the interaction of acoustic properties within such a system" (p. 265).

The application of this paradigm to musical intervals should prove useful in supporting the evidence found in the first two experiments for a dual-processing hypothesis. Additionally, musical intervals appear particularly well-suited to this sort of technique because their relatively simple acoustic structure permits comparability between different classification dimensions, unlike many speech experiments. The auditory dimension in this experiment is pitch (i.e., overall pitch of a chord, corresponding to key in musical terminology). The major–minor distinction is the symbolic higher order dimension. A correlated condition (i.e., where both dimensions vary in a redundant fashion) was not used because of the difficulty in interpreting redundancy gains that may be observed (cf. Eimas et al., 1981).

In a pilot study, the following four pure-tone stimuli were constructed: low minor, low major, high minor, and high major. The pitch (low–high) dimension was cued by a semitone change in both notes of the two-note chord, whereas the major–minor dimension was cued by the interval between the bottom and top notes, as before. When this experiment was attempted, only one listener could be found that was able to classify the stimuli accurately in the orthogonal condition; the rest were inconsistent in their classification, although they could perform correctly in the control conditions. This result was surprising but corroborates the finding of Zatorre and Halpern (1979) that identification of harmonic intervals is extremely difficult in the absence of a stable pitch context.

To increase the musical nature of the stimuli as well as to make the task easier it was decided to use three-note chords as stimuli; this is the experiment to be reported here.

Method

Subjects. A total of nine subjects was used, one of whom performed at chance levels and was excluded. The average musical training of the group was 10.5 years spent studying or playing an instrument. Other characteristics of the group were the same as in the previous experiments.

Stimulus materials. Stimuli were four pure-tone triads generated by the PDP-8. Two low (key of F) and two high (F-sharp) triads were constructed, one minor and one major within each key, thus giving four chords: F minor (F, A flat, C), F major (F, A, C), F-sharp minor (F sharp, A, C sharp) and F-sharp major (F sharp, A sharp, C sharp). Each stimulus was 250 msec in duration.

Procedure. The eight subjects were divided into two groups of four; each received two control and two orthogonal conditions in an order given by a balanced Latin square design. In the orthogonal conditions, all four possible stimuli were presented in a pseudorandom order, and subjects were instructed to press one of two buttons on the teletype with the index finger of the preferred hand. In the orthogonal pitch condition the buttons were labeled *high* and *low*, and subjects were instructed to ignore whether the triads were minor or major. In the orthogonal interval condition the buttons were labeled *minor* and *major*, and subjects were asked to classify the stimuli accordingly, regardless of key.

There were also two control conditions: In the interval control condition two triads were chosen that shared the same key (either F minor and major or F-sharp minor and major), and only these two were presented in a pseudorandom order for classification according to interval. Subjects in the first group received one combination, whereas those in the second received the other possible combination of stimuli (e.g., if the first group heard the two F triads, the other heard both F-sharp triads). The same procedure was followed for the pitch control condition, so that one group received the low and high (F and F-sharp) major stimuli, whereas the oher received the two minor triads. In all cases, right–left positioning of the response buttons was counterbalanced across groups.

Before participating in the experiment, listeners were familiarized with the four stimuli by indicating to them which chords would be presented in musical notation as well as by permitting them to listen to all four stimuli several times. Before each of the four conditions the subjects were told (and shown) which stimuli would be presented and 20 practice trials were given with feedback on the number of errors. The actual experimental conditions consisted of 80 presentations of each stimulus, thus making a total of 160 trials for each control condition and 320 trials for each orthogonal condition.

Results

The data files from each subject in each condition were edited by removing all the RTs from error trials, keeping track of the number of errors. Mean and median RTs were calculated for each block of trials that constituted a condition for each subject. These means and medians were entered into separate ANOVAs with two factors: condition (control vs orthogonal) and dimension (pitch vs. interval).[8] For simplicity, only median RT scores are reported here; means showed essentially the same pattern. The principal results are shown in Table 1. Although there were significant main effects for both condition and dimension, these were subsumed under the interaction effect, which was highly reliable, $F(1, 7) = 14.03$, $p < .01$. Post hoc Tukey HSD tests showed that the orthogonal interval condition was significantly slower that the rest; the other three medians did not differ ($p > .05$) from one another.

Errors accounted for a relatively large proportion of trials in the orthogonal interval condition (6% average errors compared to 1% or 2% for the other conditions). Nevertheless, an ANOVA for the error proportions showed no significant main effects or interactions. In any case, the greater number of errors for the orthogonal interval condition corroborates the slower latencies observed in that condition and would speak against the occurrence of a speed–accuracy trade-off.

Table 1
Average Median Reaction Times and Standard Deviations for Four Conditions of Experiment 3

	Condition			
	Control		Orthogonal	
Dimension	*M*	*SD*	*M*	*SD*
Pitch	376.6	27.8	436.9	71.9
Interval	408.1	42.7	654.5	103.2

An item analysis was also carried out on the data from the orthogonal interval condition to examine the possibility that some of the stimuli may have been more responsible for the observed effect than others. In particular, it was thought that perhaps the low-minor and high-major (F minor and F-sharp major) chords were easier to respond to because they contained the lowest and highest pitches of the entire set. The result of this item analysis revealed no significant effects for either median RTs or error proportions as a function of stimulus. Therefore, there is no evidence that subjects responded faster to some stimuli than to others in this condition.

Discussion

The pattern of results obtained in this experiment suggests that processing of interval information depends on the prior processing of pitch, whereas the pitch decision can be made independently of the intervals formed by the stimuli. This result supports and extends the data from the first two experiments with respect to the existence of separate processing levels.

In previous studies using this paradigm (Blechner et al., 1976; Eimas et al., 1978; Pas-

[8] A previous analysis showed no differences between the two groups of subjects who received different control conditions, so the two groups were combined for subsequent analyses.

tore et al., 1976; Wood, 1975) there has always been the problematic issue of comparability of the dimensions used; it is often difficult to specify equivalent levels of discriminability when such disparate dimensions as intensity and place of articulation are used. Similar problems arise when the two cues differ markedly in duration (cf. Pastore et al., 1976). These problems should be much attenuated in this study because the underlying cues in both pitch and interval conditions involved semitone changes in the component tones of the stimuli. There may still be a slight problem in the control conditions in that only a single tone change cues the minor–major distinction, whereas in the low–high distinction all three tones differ by a semitone. The availability of these multiple cues may have been responsible for the slight increase in RT between the control pitch and interval conditions (see Table 1), although this increase was not statistically significant.

Note that in the orthogonal conditions, which are of greatest interest, the two dimensions are more nearly comparable. When all four stimuli are presented and a decision on the pitch dimension is called for, the cues are given by the pitches of the top and bottom notes (the central note is not a reliable cue). A subject might still respond on the basis of all three tones in the two extreme stimuli (F minor and F-sharp major) thereby increasing speed, but the item analysis previously described indicates that this did not occur. In the case of the orthogonal interval condition there are also two cues: the interval between the bottom and central notes and the interval between the central and top notes (the bottom and top notes form a constant perfect fifth at all times).

It is of interest to note that the increase between control and orthogonal conditions for the interval dimension (246 msec) is much larger than usually found in such studies, even though the control latencies are roughly comparable. On the other hand, comparison of the relative increase in interval classification time between control and orthogonal conditions shows that this increase is approximately four times larger than the corresponding increase for pitch. When evaluated in this fashion other studies have found similar results: Eimas et al. (1981) reported a relative increase of about 8:1 for voicing versus pitch dimensions and 3:1 for place and loudness. The data of Pastore et al. (1976) are in the same range; however, Wood's (1975) data show much larger relative increases, on the order of 25:1. The amount of interference in Experiment 3, when measured in this manner, does not appear to be as large as it seems at first glance, then, but is instead comparable to some previous studies.

Another factor that is relevant to the substantial interference observed for interval classification is the lack of a stable pitch context against which intervals may be judged. Because the pitch varied randomly from trial to trial, subjects may have found it difficult to establish a stable contextual memory (cf. Zatorre & Halpern, 1979). Studies of melodic recognition (Dowling, 1982) have come to similar conclusions regarding the role of tonal context.

In spite of these considerations, the fact remains that listeners were able to accurately classify the stimuli in the orthogonal interval condition and that the interference effect was very strong. This result, along with the data from Experiments 1 and 2, upholds the idea of a hierarchically organized processing system in which pitch is extracted first and intervals analyzed only after the pitch information has been extracted.

General Discussion

It appears from the first two experiments that, at least under certain conditions, regions exist where discrimination of musical intervals is heightened, and these regions are associated with the presence of categories along the continuum. The results from Experiment 2 qualify this conclusion in that the discrimination changes may be influenced by range-related variables; however, these changes are of a much smaller magnitude and are less reliable than those associated with musical-interval categories. The interference results from Experiment 1 provide evidence for a dual-processing model in which auditory memory and categorical-coding mechanisms can be dissociated. The data from Experiment 3 support this distinction and provide evidence that the stages are hierarchically organized, with categorical coding being a higher order process.

These results are in some ways similar to those previously reported for speech, but it would be unwise to overstate the similarities.

Perhaps the most important difference between the two classes of stimuli is that speech categories appear to be operative from or near the time of birth (Eimas, Siqueland, Jusczyk, & Vigorito, 1971) and therefore do not depend on explicit training. There is at present no compelling evidence that musical-interval categories have any such innate component, even though a number of theorists have posited that such categories exist and correspond to small-integer frequency ratios that have either a neural (Boomsliter & Creel, 1961; Roederer, 1973) or environmental (Terhardt, 1974, 1978) basis. The fact that in discrimination tasks nonmusicians do not show the effects associated with the presence of musical-category boundaries (Blechner, 1977; Burns & Ward, 1978; Zatorre & Halpern, 1979) casts doubt on the idea that there may be natural musical categories (although under certain circumstances pitch categories can be used by nonmusicians in production; cf. Attneave & Olson, 1971). It might still be argued that nonmusicians have lost an ability present earlier, much like adult Japanese speakers are unable to discriminate /r/ from /l/ (Miyawaki, Strange, Verbrugge, Liberman, Jenkins, & Fujimura, 1975) even though infants can (Eimas, 1975). This possibility is a real one, but it must be kept in mind that most Western musical intervals are quite different from those used in other cultures and that there are many different interval systems in the world. Thus the situation is not like that in speech perception where the world's languages tend to place phoneme boundaries in similar regions (e.g., Abramson & Lisker, 1970). If musical-interval categories are indeed learned and if speech categories do not really depend on explicit learning, then the categorical effects associated with musical intervals and speech sounds may in fact arise from different sources.

The data from these experiments also shed some light on various theoretical descriptions of categorical effects. Miller et al. (1976) proposed an explanation of discrimination peaks at category boundaries by reference to the existence of a threshold along the continuum and to the operation of Weber's law and masking effects. This explanation, although perfectly plausible for some continua (such as voice onset time), cannot account for the results using musical intervals. Pastore et al. (1976) proposed a slightly different hypothesis in which categorical perception is due to a limitation along the continuum. Although this hypothesis is not completely incompatible with results of interval-perception studies, it is not altogether clear what would constitute the limitation in this case. The existence of many interval categories along the continuum would also be difficult to explain, as would the cross-cultural variations in the spacing of intervals along the pitch continuum.

As for the dual-processing explanation of discrimination peaks as opposed to a continuous perception model, the final word must clearly await further experimentation. The dual-processing formulation provides a cogent explanation of discrimination changes along a continuum and is able to deal with the effects of interfering stimuli on discrimination as well as response bias changes along a continuum. It is further supported by the pattern of interference in the speeded-sorting task. It is parsimonious as well, with respect to the assumptions made about the role of auditory memory. The continuous model can fit the uneven spacing of stimuli along the psychological continuum but cannot explain it; the dual-processing model provides a ready explanation. A continuous model has no ready explanation for the asymmetric interference effects observed in Experment 3 either.

A related question is whether these results can be taken to apply more generally. The arguments presented have stressed the role of coding mechanisms and auditory memory in category-boundary and interference effects. It may be expected, therefore, that similar effects would be present whenever the details of stimulus physical structure are not stored in memory but instead are reduced to a few discrete categories. It is not claimed, then, that the use of categorical coding is due to some basic sensory limitation; rather, the degree to which auditory or categorical information is used will vary depending on the stimulus characteristics, task demands, and so forth (cf. Burns & Ward, 1978).

There are some theoretical models of music perception that bear some relationship to this discussion. Deutsch (1969) presented a music-recognition model based on a physiological analogy. She proposed that lower order units detect pitch and pass on the information to higher order units that are activated only by certain combinations of pitches; in turn, these units feed on to still higher level (more abstract) units that determine frequency ratios between tones. This hierarchical system fits in well with the model described in this article, particularly in the distinction between simple pitch and intervals at a given pitch (which correspond to the proposed auditory and categorical levels, respectively). However, as previously mentioned, the role of contextual factors is extremely important in real musical situations. Dowling (1982) has presented a model in which interval judgments are referred to an underlying context rather than being made directly, as in Deutsch (1969). In his view, tonally organized music gives rise to a schema of a scale, which in turn is used for the internal organization and representation of intervals. This model makes a great deal of sense because of the difficulty encountered by musically trained subjects in making interval judgments (at least for harmonic intervals) when there is no stable central tonal structure or key.

It may be fruitful, in closing, to view the results from the present experiments as elucidating the primary steps in a process that, under normal circumstances, leads to the recognition of chords and intervals and eventually melodies and songs. What remains unknown at this point is how these primary stages are integrated with abstract schemata such as scales and how these levels interact with one another.

References

Abramson, A. S., & Lisker, L. Discriminability along the voicing dimension: Cross-language tests. *Proceedings of the 6th International Congress of Phonetic Sciences.* Prague, Czechoslovakia: Academic Press, 1970.

Attneave, F., & Olson, R. K. Pitch as a medium: A new approach to psychophysical scaling. *American Journal of Psychology,* 1971, *84,* 147-166.

Blechner, M. J. Musical skill and the categorical perception of harmonic mode. *Haskins Laboratories Status Report on Speech Perception,* 1977, SR-51/52, 139-174.

Blechner, M. J., Day, R. S., & Cutting, J. E. Processing two dimensions of nonspeech stimuli: The auditory-phonetic distinction reconsidered. *Journal of Experimental Psychology: Human Perception and Performance,* 1976, *2,* 257-266.

Boomsliter, P., & Creel, W. The long pattern hypothesis in harmony and hearing. *Journal of Music Theory,* 1961, *5,* 2-31.

Braida, L. D., & Durlach, N. I. Intensity perception: II. Resolution in one-interval paradigms. *Journal of the Acoustical Society of America,* 1972, *51,* 483-502.

Burns, E. M., & Ward, W. D. Categorical perception—phenomenon or epiphenomenon: Evidence from experiments in the perception of melodic musical intervals. *Journal of the Acoustical Society of America,* 1978, *63,* 456-468.

Cutting, J. E., & Rosner, B. S. Categories and boundaries in speech and music. *Perception & Psychophysics,* 1974, *16,* 564-570.

Cutting, J. E., Rosner, B. S., & Foard, C. F. Perceptual categories for musiclike sounds: Implications for theories of speech perception. *Quarterly Journal of Experimental Psychology,* 1976, *28,* 361-378.

Deutsch, D. Music recognition. *Psychological Review,* 1969, *76,* 300-317.

Deutsch, D. Tones and numbers: Specificity of interference in short-term memory. *Science,* 1970, *168,* 1604-1605.

Deutsch, D. Mapping of interactions in the pitch memory store. *Science,* 1972, *175,* 1020-1022.

Dowling, W. J. Musical scales and psychophysical scales: Their psychological reality. In T. Rice & R. Falck (Eds.), *Cross-cultural approaches to music: Essays in honor of Mieczyslaw Kolinski.* Toronto, Canada: University of Toronto Press, 1982.

Eimas, P. D. The relation between identification and discrimination along speech and nonspeech continua. *Language and Speech,* 1963, *6,* 206-217.

Eimas, P. D. Auditory and phonetic coding of the cues for speech: Discrimination of the [r-l] distinction by young infants. *Perception & Psychophysics,* 1975, *18,* 341-347.

Eimas, P. D., Siqueland, E. R., Jusczyk, P., & Vigorito, J. Speech perception in infants. *Science,* 1971, *171,* 303-306.

Eimas, P. D., Tartter, V. C., Miller, J. L., & Keuthen, N. J. Asymmetric dependencies in processing phonetic features. *Perception & Psychophysics,* 1978, *23,* 12-20.

Eimas, P. D., Tartter, V. C., & Miller, J. L. Dependency relations during the processing of speech. In P. D. Eimas & J. L. Miller (Eds.), *Perspectives on the study of speech.* Hillsdale, N.J.: Erlbaum, 1981.

Fujisaki, H., & Kawashima, T. On the modes and mechanisms of speech perception. *Annual Report of the Engineering Research Institute,* 1969, *28,* 67-73.

Fujisaki, H., & Kawashima, T. Some experiments on speech perception and a model for the perceptual mechanism. *Annual Report of the Engineering Research Institute,* 1970, *29,* 207-214.

Fujisaki, H., & Kawashima, T. A model of the mechanisms for speech perception—quantitative analysis of categorical effects in discrimination. *Annual Report of the Engineering Research Institute,* 1971, *30,* 59-68.

Kewley-Port, D., & Pisoni, D. B. Discrimination of rise time in nonspeech signals: Is it categorical or noncategorical? *Journal of the Acoustical Society of America,* 1982, *71,* S36. (Abstract)

Kopp, J., & Livermore, J. Differential discriminability or response bias? A signal detection analysis of categorical perception. *Journal of Experimental Psychology,* 1973, *101,* 179-182.

Liberman, A. M., Harris, K. S., Eimas, P. D., Lisker, L., & Bastian, J. An effect of learning speech perception: The discrimination of durations of silence with and without phonemic significance. *Language and Speech,* 1961, 4, 175-195.

Liberman, A. M., Harris, K. S., Hoffman, H. S., & Griffith, B. S. The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology,* 1957, *54,* 358–368.

Liberman, A. M., Harris, K. S., Kinney, J. A., & Lane, H. The discrimination of relative onset time of the components of certain speech and nonspeech patterns. *Journal of Experimental Psychology,* 1961, *61,* 379–388.

Locke, S., & Kellar, L. Categorical perception in a nonlinguistic mode. *Cortex,* 1973, *9,* 355–369.

Macmillan, N. A., Kaplan, H. I., & Creelman, C. D. The psychophysics of categorical perception. *Psychological Review,* 1977, *84,* 452–471.

Massaro, D. W. Retroactive interference in short-term recognition memory for pitch. *Journal of Experimental Psychology,* 1970, *83,* 32–39.

Mattingly, I. G., Liberman, A. M., Syrdal, A., & Halwes, T. Discrimination in speech and nonspeech modes. *Cognitive Psychology,* 1971, *2,* 131–157.

Miller, G. A. The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review,* 1956, *63,* 81–97.

Miller, J. D., Wier, C. C., Pastore, R. E., Kelly, W. J., & Dooling, R. J. Discrimination and labeling of noise-buzz sequences with varying noise-lead times. *Journal of the Acoustical Society of America,* 1976, *60,* 410–417.

Miller, J. L. Interactions in processing segmental and suprasegmental features of speech. *Perception & Psychophysics,* 1978, *24,* 175–180.

Miller, J. L., Eimas, P. D., & Zatorre, R. J. Studies of place and manner of articulation in syllable-final position. *Journal of the Acoustical Society of America,* 1979, *66,* 1207–1210.

Miyawaki, K., Strange, W., Verbrugge, R., Liberman, A. M., Jenkins, J. J., & Fujimura, O. An effect of linguistic experience: The discrimination of [r] and [l] by native speakers of Japanese and English. *Perception & Psychophysics,* 1975, *18,* 331–340.

Moran, H., & Pratt, C. C. Variability of judgments of musical intervals. *Journal of Experimental Psychology,* 1926, *9,* 492–500.

Pastore, R. E., Ahroon, W. A., Baffuto, K. J., Friedman, C., Puleo, J. S., & Fink, E. A. Common factor model of categorical perception. *Journal of Experimental Psychology: Human Perception and Performance,* 1977, *3,* 686–696.

Pastore, R. E., Ahroon, W. A., Puleo, J. S., Crimmins, D. B., Golowner, L., & Berger, R. S. Processing interaction between two dimensions of nonphonetic auditory signals. *Journal of Experimental Psychology: Human Perception and Performance,* 1976, *2,* 267–276.

Pisoni, D. B. Auditory and phonetic memory codes in the discrimination of consonants and vowels. *Perception & Psychophysics,* 1973, *13,* 253–260.

Pisoni, D. B. Auditory short term memory and vowel perception. *Memory & Cognition,* 1975, *3,* 7–18.

Pisoni, D. B. Identification and discrimination of the relative onset time of two component tones: Implications for voicing perception in stops. *Journal of the Acoustical Society of America,* 1977, *61,* 1352–1361.

Rakowski, A. Tuning of isolated musical intervals. *Journal of the Acoustical Society of America,* 1976, *59,* S50. (Abstract)

Raz, I., & Brandt, J. F. Categorical perception of nonspeech stimuli by musicians and nonmusicians. *Journal of the Acoustical Society of America,* 1977, *62,* S60. (Abstract)

Roederer, J. *Introduction to the physics and psychophysics of music.* New York: Springer-Verlag, 1973.

Rosen, S. M., & Howell, P. Plucks and bows are not categorically perceived. *Perception & Psychophysics,* 1981, *30,* 156–168.

Siegel, J. A., & Siegel, W. Absolute identification of notes and intervals by musicians. *Perception & Psychophysics,* 1977, *21,* 143–152.

Terhardt, E. Pitch, consonance and harmony. *Journal of the Acoustical Society of America,* 1974, *55,* 1061–1069.

Terhardt, E. Psychoacoustic evaluation of musical sounds. *Perception & Psychophysics,* 1978, *23,* 483–492.

Ward, W. D. Subjective musical pitch. *Journal of the Acoustical Society of America,* 1954, *26,* 369–380.

Wood, C. C. Auditory and phonetic levels of processing in speech perception: Neurophysiological and information-processing analyses. *Journal of Experimental Psychology: Human Perception and Performance.* 1975, *1,* 1–33.

Wood, C. C. Discriminability, response bias and phoneme categories in discrimination of voice onset time. *Journal of the Acoustical Society of America,* 1976, *60,* 1381–1389.

Zatorre, R. J., & Halpern, A. R. Identification, discrimination and selective adaptation of simultaneous musical intervals. *Perception & Psychophysics,* 1979, *26,* 384–395.

Received June 17, 1982
Revision received March 21, 1983 ■