Music discriminations by carp (*Cyprinus carpio*)

AVA R. CHASE
Rowland Institute for Science, Cambridge, Massachusetts

Studies using three koi (*Cyprinus carpio*) investigated discrimination of musical stimuli. The common protocol used a single manipulandum and a multiple continuous reinforcement-extinction schedule signaled by music of the S+ and S− types in 30-sec presentations separated by a silent 15-sec intertrial interval. In a categorization study, the fish learned to discriminate blues recordings from classical, generalizing from John Lee Hooker (guitar and vocals) and Bach (oboe concertos) to multiple artists and ensembles. A control-by-reversal test developed into a demonstration of progressive improvement in iterated reversal learning. The subjects next learned to discriminate single-timbre synthesized versions of similar music. In the final study, which used melodies with the same order of note-duration values, but with mirror-image orders of pitch values, one fish discriminated melodies with no timbre cues, in contrast to results reported in rats.

Fish live in a medium that is an especially good conductor of sound, and the myriad extant species—some 25,000 in the class Osteichthyes (bony fishes) alone (Nelson, 1984)—have evolved diverse mechanisms for transducing acoustic stimuli. It should not be surprising if sophisticated functionality in the auditory domain turns out to be a major component of the perceptual and cognitive capabilities of fish.

This paper's point of departure is the elegant study by Porter and Neuringer (1984), which demonstrated the hitherto unsuspected capability of pigeons to classify music into stylistic categories, just as humans do. My interest in testing koi on a similar task evolved into the present series of experiments, which range over the areas of discrimination learning, categorization, reversal learning, and music perception and which indicate the existence of hitherto unsuspected capabilities in fish.

Like Porter and Neuringer (1984), I prefer to optimize the chance of discovering currently unknown capabilities in an animal by seeking a positive result in a top-down fashion, first challenging the subject with a task complex enough to require such capabilities, and then simplifying only as necessary. The current understanding of fish psychoacoustics does nothing to discourage high expectations of fishes' higher perceptual functions. In the case of perhaps the most extensively studied species, "absolute detection thresholds, frequency discrimination, intensity discrim-
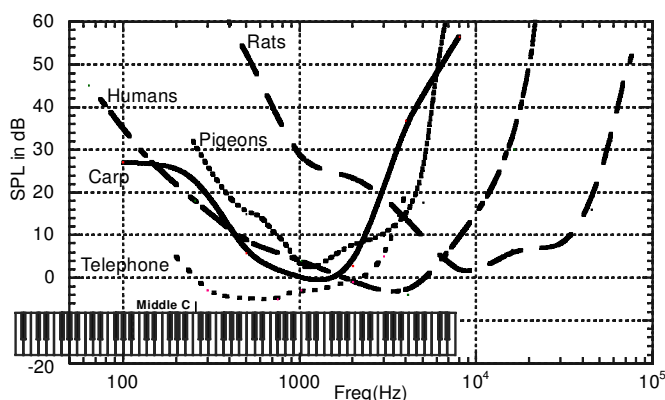
ination, temporal summation, complex spectrum discrimination, temporal discrimination and resolution, and various measures of 'auditory filter' characteristics show that the goldfish falls solidly within a general vertebrate pattern of auditory processing" (Fay & Ream, 1986, p. 1883). Furthermore, goldfish are known to perceive acoustic dimensions equivalent to what humans perceive as pitch and timbre (Fay, 1995). Indeed, for more than a decade, it has seemed clear that humans and animals share the basic psychoacoustic functions and that there are probably no fundamental qualitative differences in auditory perception between humans and at least those fish whose anatomy allows them to hear reasonably well (Fay, 1992, 1994, 1995, 1998; Jacobs & Tavolga, 1968; Popper & Fay, 1993; Popper, Platt, & Saidel, 1982).

The present subjects were koi, which are members of the carp family and close relatives of the goldfish. Neither goldfish nor koi are known to communicate by making sounds (Fay, 1995), but both are classified as otophysans, a group that has bony structures (the Weberian ossicles) coupling the swim bladder to the inner ears, in which there are hair cells with a specialized orientation pattern. Because of this anatomy, carp are among the species referred to as *hearing specialists*, which have the greatest auditory bandwidth and sensitivity of any fish (Popper & Fay, 1993). Figure 1 shows the frequency sensitivity range of koi in comparison with the spectral ranges of other animals and of music. Their relatively low upper limit of hearing means that koi would hear music roughly as it would sound to us over an ordinary voice-grade telephone line.

In one of the more complex experiments in fish psychoacoustics, Fay (1992) conditioned goldfish to a two-tone chord and then tested them on a range of individual pitches, obtaining essentially a bimodal generalization curve whose peaks corresponded to the component pitches of the training chord. Thereby, Fay was the first to demonstrate analytic listening (Hartman, 1988) in a nonhuman animal—in this case, the ability to discriminate the individual components

Figure 1. Hearing sensitivity threshold curve averages for various animals (Fay, 1988), sound pressure level (SPL) versus frequency. The SPL in water is adjusted by −33 dB with respect to air. A high SPL level indicates reduced sensitivity at that frequency. The animals are carp (Sigmond & Wolff, 1973, in Fay, 1988), rats (Kelly & Masterton, 1980, in Fay, 1988), humans (Sivian & White, 1933, in Fay, 1988), and pigeons (Goerdel-Leich & Schwartzkopff, 1984, in Fay, 1988). The relative attenuation with frequency of typical analog phone lines in Houston, TX (Garfield, Houston Area League of PC Users, http://www. houston.tx.us/ internet/dialup.shtml) is shown for comparison, as is a piano keyboard (middle C = 256 Hz).

of a compound stimulus consisting of simultaneous tones or pulse trains (Fay, 1992, 1998). Although perhaps complex by psychophysical standards, however, even these stimuli did not embody any large-scale pattern or structure that would challenge or demonstrate a listener's higher mental functions.

Porter and Neuringer (1984) pioneered a dramatic way to begin exploring the connections between psychoacoustics and cognition. Music has become an important kind of stimulus in comparative psychology because of its excitation of fundamental cognitive functions, the accessibility of its perceptual dimensions (pitch, timbre, tempo, etc.), and its analytically tractable internal structures (Hulse & Page, 1988). Little is known about the correspondence between the elements of music, the units of perception, and the stimulus-feature *attending hierarchies* (Baron, 1965; Segal & Harrison, 1978)—that is, the various stimulus dimensions to which different species preferentially attend. I felt that a demonstration that fish could learn to categorize stimuli as complex as music, according to a criterion as unnatural as musical genre, not only would show that they extract meaningful signals, but also might suggest a capability for higher order auditory processing comparable with that of pigeons and people.

GENERAL METHOD

Subjects

The subjects in all the experiments were 3 pet store koi (*Cyprinus carpio*) named Beauty, Oro, and Pepi. By the start of the work reported here, they had been living together and serving as experimental subjects for 5 years, and they ranged in age from 7 to 11 years and in size from 29 to 36 cm (snout to base of tail). In a prior study of sensory reinforcement, all 3 had had extensive exposure to sev-

eral kinds of music, including the blues and classical genres, and even to some of the actual recordings that they would learn to actively discriminate in the present series of experiments. Beauty had also learned a visual categorization task in a two-operandum procedure.

On average, the fish received a collective maintenance ration consisting primarily of about 20 grams of Hikari pellets (Kyorin Food Ind. Ltd.) every weekday. On experiment days, each animal also typically consumed a gram of 20-mg reinforcement pellets during the session; the maintenance ration was given only after the day's experimentation.

Apparatus

The fish apparatus was novel and was developed in my laboratory. Its rationale, design, and construction have been described in detail elsewhere (Chase & Hill, 1999). During experimental sessions, the home aquarium was divided by a clear partition into an experimental chamber and a holding area, and the noisy aeration system was turned off. There were no special provisions for visual or acoustic isolation; indeed, visual isolation from the shoal was found to impede performance. Auditory stimuli were fed from either of two computer-controlled SCSI (small computer system interface) compact disc (CD) players to an underwater speaker. The S+ CD was loaded into one drive, and the S− CD into the other. Compression circuitry (DBX Professional Products, Model dbx 163X) made stimuli from the two sources sound equally loud. A 200-Hz, eighth-order high-pass filter was available to provide the option of reducing possible lateral-line stimulation in fish by removing music frequencies lower than 200 Hz, in which this sensory system is most responsive (Coombs, Jannsen, & Montgomery, 1992). The tank setup is illustrated in Figure 2. The operandum was a horizontal button positioned at the bottom of the tank in front of the speaker. A circular puck made of plastic and buoyant foam was the button's only moving part. Designed to be pressed downward by a fish and to be restored by buoyancy alone, the puck floated horizontally beneath the overhanging rim of the cylindrical housing. To respond, a fish had to swim down into the cylinder and push the puck with its snout. The puck was translucent, and through it a responding fish could see an internal light source that provided feedback that a response had

been registered. The unconditioned reinforcer was food, in the form of 20 mg-sinking pellets (Noyes, formula J/G). The pellets were dropped by an automatic feeder (Gerbrands G5120 or MED [Georgia, VT] Model ENV 203-20) into a plastic tube that terminated underwater in an infant's Stage One nipple (Johnson & Johnson, Skillman, NJ) that had been modified to let fish suck the pellets out. The instantaneous conditioned reinforcer was an external light that backlit the nipple. A computer controlled the experiment, and sessions were videotaped.

Preliminary Phase

The processes of magazine training and of familiarization with the response button, the stimulus presentation system, and the rest of the apparatus are described in detail in Chase and Hill (1999). Before conducting the genre discrimination experiment, it was necessary to test the novel apparatus and the intended discrimination protocol. For example, one objective was to verify that the new horizontal response button did not artificially inflate response rates, either by being intrinsically reinforcing or by eliciting instinctive mouthing that could be mistaken for a discriminative response.

The subjects were given a simple discrimination task—namely, sound versus silence—which is described in detail in Chase and Hill (1999). The procedure presented randomly timed intervals (typically between 40 and 90 sec) of music (S+) and silence (S−) in alternation, and reinforced each response that was made in the presence of music with food. To ensure that the onset of music would never reinforce a response made in silence, any response made within 20 sec before the scheduled end of a silent interval prolonged the silence by resetting the countdown timer to 20 sec. The reinforcement schedule was multiple continuous reinforcement-extinction (CRF-extinction)—that is, while music was playing every response was reinforced with a food pellet, and during silence no response was reinforced. A session ended with the S+ trial in which the 50th reinforcement was dispensed.

Because CRF provides informational feedback on every response, a subject could "fake" a sound discrimination, at least in the sense of exhibiting an appropriate response rate differential, merely by sampling and discriminating the prevailing reinforcing contingency. However, if a fish suppressed responding substantially more in the presence of one kind of stimulus than in the presence of the other, it would be clear that the fish was not routinely sampling the button and, thus, that it was under the control of the sound. As is shown in Figure 3, Beauty and Oro met this standard within the first few sessions. The characteristically hyperactive Pepi went 13 sessions before his first S− trial, during which he did not respond at all, but otherwise from the start he was much like the other subjects in exhibiting downward trends in the average number of responses per S− trial and in the percentage of his total responses that occurred during S− (Chase & Hill, 1999).

In Figure 3 and in all other graphs of results, the session axis denotes experimental workdays, in each of which a single set of stimuli was used with all subjects that ran that day. Every workday is indicated on every subject's session axis, but an individual subject's statistics are based on his actual runs only. Skipped sessions were rare and were caused by anomalous conditions that either prevented an animal's participation outright or else threatened to render meaningless any data that might be taken. The fact that the results were analyzed in terms of within-subjects comparisons renders immaterial such fine-grained differences between the experiences of different subjects.

The final function of the sound-versus-silence task was to accustom the subjects to responding in the presence of the music that they would later hear as the first S+ in the genre discrimination task, and therefore all of the sound stimuli were taken from John Lee Hooker's *Blues Before Sunrise* CD. The hope was that this preparation would facilitate the acquisition of the music–music discrimination.

EXPERIMENT 1
Discrimination and Categorization Based on Musical Genre

In this experiment, it was asked whether fish can learn to discriminate between two stylistically different musical
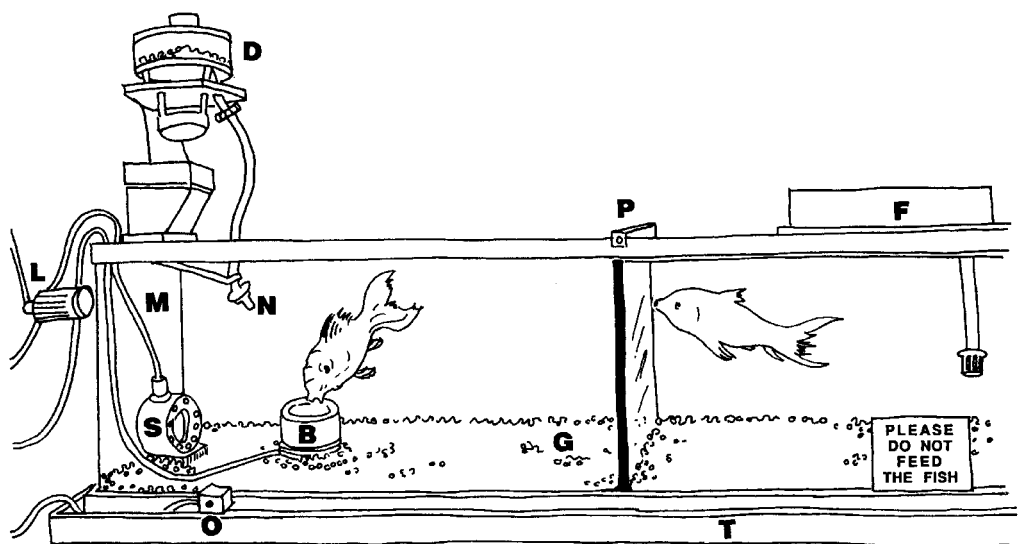

Figure 2. Working end of the experimental tank, as seen by the video camera (not shown), includes a partition (P, at the silicone joints), a filter (F), a tank tray (T), a gravel bed (G), a speaker (S), a response button (B), a light (L, the conditioned reinforcer), which shines through a light-diffusing mylar (M) onto a nipple (N), a pellet dispenser (D), and an LED (O), which shows the camera that the computer has registered a response. From "Reliable Operant Apparatus for Fish: Audio Stimulus Generator, Response Button, and Pellet-Dispensing Nipple," by A. R. Chase and W. Hill, 1999, *Behavior Research Methods, Instruments, & Computers*, 31, p. 471. Copyright 1999 by the Psychonomic Society. Reprinted with permission.

stimuli and, if so, whether they can generalize to unfamiliar music from the same stylistic categories, capabilities previously demonstrated in pigeons by Porter and Neuringer (1984).

A standard demonstration of categorization behavior begins by training an animal to classify the members of two nonoverlapping sets of stimuli by responding only in the presence of the members of one particular set. Each set is an experimenter-defined category, and each stimulus is an exemplar of one category. The animal is said to have learned the two categories if it then responds appropriately to novel exemplars of each and if there is reason to believe that it can also discriminate exemplars within each set, as well as between sets. The latter criterion ensures that the animal was perceiving categories, rather than exhibiting categorical perception; that is, the animal's treating certain exemplars alike must not be due merely to an inability to tell them apart (Herrnstein, 1992; Sturdy, Phillmore, Price, & Weisman, 1999).

In visual categorization studies, pigeons have been trained to classify photographs according to whether they do or do not contain people, water, fish, aerial views of a particular location, tree leaves damaged by a particular species of caterpillar, and other such complex criteria (Herrnstein, 1992). If intermittent reinforcement is used, particular exemplars might be repeated yet never associated with reinforcement; if there is improvement in the accuracy of classification of even these stimuli, it is reasonable to conclude that the animal is responding to them on the basis of features that imply membership in a specific category instead of as isolated stimuli. The training sets can range from closed and repetitive stimuli, with testing done by using novel stimuli as probes, to streams of continually novel stimuli, in which case classification accuracy measures categorization behavior on an ongoing basis—even under CRF—because rote memory (at least of whole stimuli, as opposed to features) can play no role (Malott & Siddall, 1972).

Categories can also be defined with no systematic relationship to stimulus properties, so that the only way to distinguish them is through rote memory of the reinforcement contingency associated with each exemplar. Rote memory capacity in animals can be considerable. Vaughan and Greene (1984) trained pigeons to sort visual images into rote categories and found that they could easily memorize hundreds of pictures and remember them for at least 2 years. In foraging and food-hiding birds, such as the Hawaiian honeycreeper and Clark's nutcracker, estimates of the memory capacity for locations range well into the thousands (Herrnstein, 1992).

The protocol in the present study used closed training sets and CRF, but within-session repetition was almost nonexistent. The first presentation of a pair of novel CDs constituted a stream of novel stimuli and served as a measurement probe.

Because of the prevailing skepticism that koi could demonstrate the capabilities I was hoping to find, I tried to make every task as easy as possible for the animals, a constraint that dictated some procedural choices that might be regarded as unconventional. One example is the use of a free operant and CRF, rather than discrete trials or a lean reinforcement schedule. Another is the use of reversals instead of extinction trials in order to demonstrate control by stimuli, rather than by prevailing reinforcement contingency. The rationale for CRF was to minimize stress and maximize feedback, given that the exertional cost of buttonpressing appeared relatively high and that, consequently, the maximum possible response rate was expected to be relatively low. Unlike a discrete-trials procedure, the free-operant approach fixed each trial's stimulus presentation time, while letting the numbers of responses and reinforcements vary. My intention was to minimize interference with the formation of associations between the positive stimulus and the instantaneous conditioned reinforcement or the actual consumption of food. With respect to elucidating the stimulus control, I was concerned that a switch to total extinction might cause excessive emotional side effects and reduce the subjects' long-term viability. Therefore, in my experiments, the usual role of extinction trials was played by reversals. In a study with so few sub-
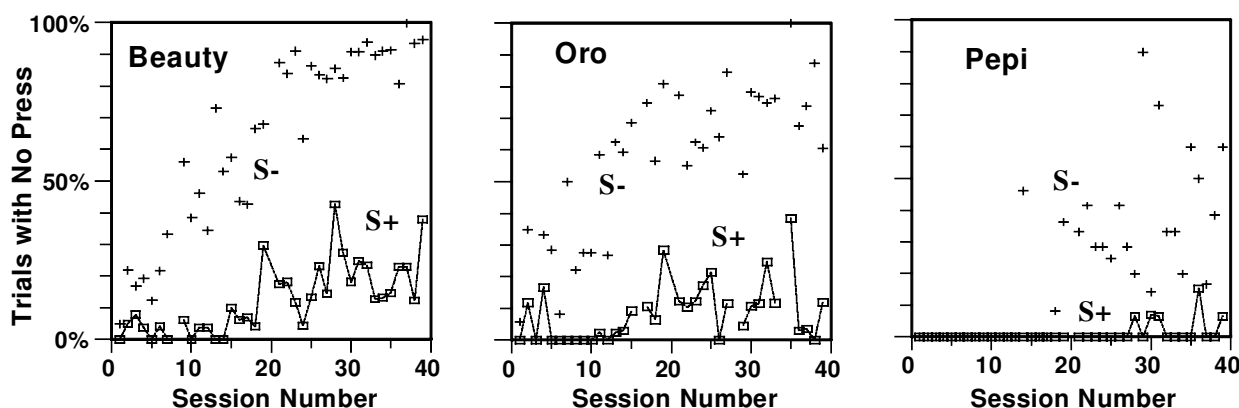

Figure 3. Preliminary phase (sound vs. silence): percentages of S+ and S− intervals in which the fish did not press the button at all. A comparatively high percentage of nonresponding for S− indicates successful discrimination. Crosses denote silence.

jects, reversals can also act like counterbalancing in controlling for possible differences in familiarity and saliency with respect to the stimuli designated as S+ or S−. Reversals have the further advantage of offering glimpses of additional phenomena—in particular, reversal learning per se.

Method

The protocol was an extension of that used in the preliminary sound-versus-silence discrimination as described above. However, now there were three kinds of intervals—namely, S+ trials (blues music), S− trials (classical music), and a silent intertrial interval (ITI). Also, the intervals were now of fixed duration, 30 sec for S+ and S− sound trials and 15 sec for the ITI. During the ITI, a correction procedure was programmed so that any response that occurred less than 10 sec before the silence was scheduled to end prolonged the ITI by resetting the countdown timer to 10 sec. The sequence of S+ and S− trials was randomized in real time, but a session always began with two successive S+ trials, and there could be no more than three consecutive trials with any one stimulus type. A session ended at the end of the full S+ interval during which the 50th reinforcement was delivered.

An important aspect of the procedure was that the computer did not restart the piece of music at the beginning of each trial but, instead, resumed the CD where it had left off (unless there were fewer than 30 sec of music remaining on that track, in which case the stimulus would be the first 30 sec of the following track, the intent being to try to ensure that every stimulus would be a full 30 sec of music). Each of the S+ and S− CDs contained roughly 120 distinct excerpts, and a session ended with the S+ excerpt during which the 50th response was made, so it was rare for a session to run long enough for any CD to start over and, thus, there was relatively little within-session repetition of stimuli. Each excerpt was effectively a separate mini-piece, in one of the two styles represented by the two CDs. Thus, even within a single piece of music or CD track, what the fish confronted was essentially a categorization problem, in contrast to a discrimination between two brief musical fragments that would be repeated identically many times in each session. Because the computer consistently divided a CD into the same sequence of fixed-length excerpts each time, a musical CD was somewhat analogous to a tray containing 120 photographic slides as a source of exemplars in a visual categorization experiment.

Figure 4 summarizes the sequence of experimental conditions and the recordings used, starting with the preliminary sound-versus-silence phase. The computer software allowed individual tracks to be marked for exclusion, and, as is noted in the CD reference list, in two of the blues CDs certain tracks were excluded because they contained speech or silence.

In the beginning of Experiment 1 (Sessions 40–55) the stimulus sources were two entire CDs: S+ was John Lee Hooker's *Blues Before Sunrise* (carried over from the sound-versus-silence discrimination), and S− was Bach oboe concertos (familiar from a prior study). Over the course of the first 6 sessions, 2 fish showed no evidence of learning the discrimination, whereas the 3rd achieved a success that proved short-lived. Therefore, I simplified the task by restricting the stimuli to a single track from each CD. The first such pair of tracks was used for 30 sessions, and over the next 15 sessions three new pairs of tracks were used, each pair for 5 consecutive sessions. After the animals had learned to discriminate a single classical track from a single blues track these four times, I again allowed the stimuli to come from the entire CD of each genre.

Because CRF was always in effect during S+, providing feedback as well as reinforcement, the best way to reveal whether the fish were classifying the music on the basis of a categorical principle, rather than on the basis of rote memory, was to present novel stimuli. Therefore, scattered throughout the experiment were nine *probe* sessions in

which both the S+ and the S− CDs were novel. These were Sessions 115, 120, 131, 138, 152, 157, 163, 168, and 177. The music was selected for diversity within the genre: S+ included four different performers' CDs, plus five compilation CDs with more than 20 additional blues performers. S− consisted of six baroque-style CDs and three non-baroque composers—specifically, Mozart, Beethoven, and Schubert.

The buttonpressing response involves a fish's entire body, since the fish must position itself over the button and essentially ram it. Therefore, such a response is inherently somewhat slow—far slower than, for example, a bird's peck at a key or a rat's press of a lever. Because the fish typically would not hover over the button and because they appeared to commit to a response by swimming in from a distance, after each S+ interval the silent ITI began with a grace period of 1 sec, during which a buttonpress would still be counted and reinforced, on the assumption that the response had begun during S+.

During Sessions 147–149, the 200-Hz high-pass filter was turned on in order to minimize stimulation of the lateral line. The idea was to try to ensure that the discrimination between blues and classical would be made on the basis of input from the fish's inner ears, the sensory system presumably most capable of detecting the fine-grained acoustic features of music that form the basis for stylistic distinctions by humans.

The possibility of acoustic artifacts owing to hardware was tested by a control procedure in which the blues and the classical CDs were each loaded into the drive programmed for the other (Sessions 182–238). The effect was to reverse the definitions of S+ and S− (i.e., to reinforce responses to classical music, but not to blues). This was expected to produce a dramatic decrement in accuracy, which would demonstrate that responding had indeed been under the control of the musical content of the S+ and S− channels.

Results and Discussion

The use of CRF imposed the obligation to verify that the subjects were discriminating the stimuli, and not merely the current state of the multiple CRF-extinction schedule. Proof that a subject was not being cued by feedback from routine test responses would be a statistically significant difference between the percentages of S− and S+ trials in which the subject did not press the button at all. Such a finding would allow rejection of the hypothesis that the subject had been testing the button equally during both stimulus conditions. The alternative hypothesis is directional in that, beyond exhibiting a difference per se, an auditorily discriminating subject would be expected to ignore the response button more during S− than during S+. Therefore, unless otherwise specified, throughout this paper significance is reported in terms of the p value of a one-tailed, paired t test (Microsoft Excel 97 SR-1) on the nonresponse percentages associated with the two stimulus types in each of n sessions; the alpha level used was .05.

As is indicated graphically in Figures 5 and 6, all 3 koi discriminated blues from classical music, and this behavior generalized into an ability to classify unfamiliar exemplars by genre [for each of the 3 subjects, 9 probe sessions, $t(8) > 10$, $p < .001$]. To my knowledge, this is the first demonstration in fish of *open-ended* categorization (Herrnstein, 1992) and of any kind of categorization with auditory stimuli.

Figure 5 shows the percentages of S+ and S− trials in which there was no responding. In Session 115, the first probe, Beauty responded only during S+ and not at all dur-

SESSIONS	S+	S-	REMARKS
1–39	•••	•••	Prelim. Phase: sound/silence
40–55	•	•••	Experiment 1. Entire CDs.
56–85	8	5	Single tracks, repeated.
86–90	6	8	"
91–95	1	6	"
96–100	5	7	"
101–107	•	•	Back to entire CDs.
108–114	•	•••	Novel S–
115–119	•••	•••	Probe #1 = session 115
120–126	•••	•••	Probe #2 = session 120
127	•	•	
128–130	•	•	
131–137	•••	•••	Probe #3 = session 131
138–144	•••	•••	Probe #4 = session 138
145–146	•	•	
147–149	•	•	High-pass filter ON
150–151	•	•	" " " OFF
152–156	•••	•••	Probe #5 = session 152
157–162	•••	•••	Probe #6 = session 157
163–167	•••	•••	Probe #7 = session 163
168–171	•••	•••	Probe #8 = session 168
172–175	•	•	
177–181	•••	•••	Probe #9 = session 177
	S–	S+	Reversal #1. Experiment 2.
182–192	•	•	
193–202	•	•	
203–212	•	•	
213–223	•	•	
224–238	•••	•••	Probe = session 224
	S+	S–	Reversal #2.
239–247	•	•	
248–263	•	•	
	S–	S+	Reversal #3.
264–282	•	•	
	S+	S–	Reversal #4.
283–300	•	•	

Stimulus CDs (S+ column, bottom of figure, left to right):
J. L Hooker, *Blues Before Sunrise*
J. L. Hooker, *Real Folks Blues*
J. L. Hooker, *Boom Boom*
Muddy Waters
Koko Taylor
Chess Blues compilation
Another Chess Blues compilation
Blending Blues compilation
Legendary Blues compilation
Sonny Boy Williamson
House of Blues compilation
The Blues Story compilation

Stimulus CDs (S– column, bottom of figure, left to right):
Silence
Bach Oboe Concertos
Handel Trio Sonatas
Mozart Oboe Concertos
Beethoven Cello Sonatas
Vivaldi Cello Concertos
Vivaldi Guitar Concertos
Schubert: Trout Quintet
Baroque compilation
Brandenburg Concertos #1-3
Bach & Vivaldi, for 2 violins
Baroque horn compilation
Bach Trio Sonatas

Figure 4. Experiments 1 and 2 (categorization and iterated reversals): stimulus chart. Each column represents a single stimulus CD, which is identified by text that starts in that column at the bottom of the figure. Each row represents a group of sessions. Bullets indicate the CDs presented in their entirety as the S+ and S− in each session group, whereas the numbers associated with Sessions 56–100 identify the specific CD tracks used as stimuli. A triple bullet marks the first presentation of a given CD; a single bullet denotes a stimulus heard previously. Rows in which both stimuli are represented by triple bullets denote groups of sessions in which the first one exposed the subject to all-novel stimuli and thus served as a probe session.

ing S− or ITI. In Probe Session 138, the novel S− was Vivaldi guitar music, which proved to raise problems. Figure 6 shows all the responses for each session; note the higher S− response rate at Session 138. Session 182 was the first control-by-reversal session. Sessions 138 and 182 are further discussed below.

The data from the probe sessions are plotted in Figure 7 and are shown in historical context in Figure 5. Figure 7

also shows means for the nonprobe and the probe sessions in the session range 115–181. The fish classified music with which they had had no training about as accurately as they classified familiar training stimuli. This generalization to unfamiliar exemplars demonstrates that the fish were not relying on rote memory but were, instead, exhibiting open-ended categorization (Herrnstein, 1992).

To assess whether the fish were exhibiting true perception of categories, preclusion of categorical perception (i.e., the inability to discriminate exemplars within a category) was necessary. The great complexity of musical stimuli renders likely the existence of a mix of feature dimensions, on the basis of which a sample only a few seconds long suffices to distinguish any two pieces, or even different performances of the same piece. The blues stimuli ranged from John Lee Hooker to Sonny Boy Williamson, and the classical ranged from baroque to Schubert; both categories featured a variety of ensembles. It is implausible that recordings that sound so very different to humans would be completely indistinguishable to these fish. The result of Experiment 4, below, provides strong additional evidence: Koi can discriminate generically similar melodies, which implies that they can discriminate generically similar recordings that differ along a multiplicity of dimensions, of which melody is only one.

In Figures 5 and 6, the data for Sessions 182–192 show the disruption of performance following the reversal of the reinforcement contingencies. This reversal served as a control procedure, the result of which confirmed that responding was indeed under the control of the musical content, as opposed to any electronic artifact in the S+ and S− audio channels.

In Sessions 138–144, the fish had unexpected difficulty discriminating Vivaldi guitar concertos from a blues compilation. Although it is plausible that the fish could have learned that guitar sounds were characteristic of blues, the guitar concertos arguably did not sound to humans as much like blues as like the Vivaldi cello concerto, which was the previous S− and which the fish had had no trouble distinguishing from blues. Another possible source of confusion for the fish was the fact that this S+ was the first compilation CD to be presented as a stimulus. Thus, the diversity of blues styles distinguished this S+ CD from any previous one. Given the fishes' problem with this discrimination, I paired the Vivaldi guitar concertos with some very familiar John Lee Hooker (Sessions 145–146). When discrimination did not improve, I left these stimuli in place but turned on the 200-Hz high-pass filter for three sessions (147–149). As can be seen from Figure 5, the result of turning on the filter was an immediate improvement in performance, which persisted after the filter was turned off. Apparently, these particular stimuli were salient and had similar features in the lower frequencies, which confused the fish until the high-pass filter forced them to attend to higher frequency features, on the basis of which the music was more readily discriminable. It would also appear that the lateral line was not necessary for the discrimination of familiar music. As D'Amato and Salmon

(1982), noted, "Individuals do not always rely on the same property of the stimulus and may, according to circumstances, shift their reliance from one feature to another" (p. 130). (In any case, Vivaldi evidently can impress listeners in odd ways. Porter & Neuringer, 1984, reported that a Vivaldi violin concerto sounded more like Stravinsky than like Bach to their pigeons, as it did also to some humans.)

Even a convincing demonstration of categorization can fail to identify the stimulus features that exert control at any given time, especially if the stimuli are complex. In particular, there can be uncertainty as to whether classification behavior had been under the stimulus control of the features in terms of which the experimenter had defined the categories or whether the subjects had discovered an effective discriminant of which the experimenter was unaware. The diversity of S+/S− pairings (see Figure 4) presumably rules out the possibility that the fish could have been relying on only a single discriminant, such as timbre, but a constant concern is the possible existence of a simple attribute that would have allowed the subjects merely to discriminate instead of categorizing. In the present experiment, the choices of stimuli and control procedures were intended to ensure that all the available discriminants were sufficiently complex. For example, even if a fish used, as the discriminant, the sound of a human voice, which was present in virtually all of the blues but in none of the classical, the diversity of singing voices would imply open-ended categorization nonetheless.

The possibility that the fish memorized individual or recurring features cannot be ruled out. An animal probably does exploit a collection of memorized features, along with various abstracted properties. It would be surprising if no features at all were salient enough to be remembered, and one wonders whether it is even possible to abstract a categorical discriminant without having particular features in memory, at least temporarily.

The notion that koi might classify blues and classical music on the basis of deeper generic attributes is supported by the results of experiments reported below, which indicate that timbre cues are not required. Porter and Neuringer (1984) trained pigeons to discriminate baroque from modern music, using initial exemplars that could have been alternatively classified as organ and orchestral music, but the effective discriminant turned out to be something other than instrumental timbre. In a recent report (Watanabe & Sato, 1999), Java sparrows learned to discriminate piano music by Schönberg from piano music by Bach, and they then correctly classified orchestral music as Schönberg-like or Bach-like, whether written by those two composers or by Carter and (even) Vivaldi. Studies that analyze the perceptual world of pigeons in terms of multidimensional scaling spaces reveal that the birds learn enough about nonexemplars to form a distinct category for them, as well as for exemplars. In other words, pigeons appear to evaluate a stimulus not only as possibly a good exemplar, but also as a good nonexemplar, which is more work than binary classification logically requires (Herrnstein, 1994). Given that the more features of the environment to which one is
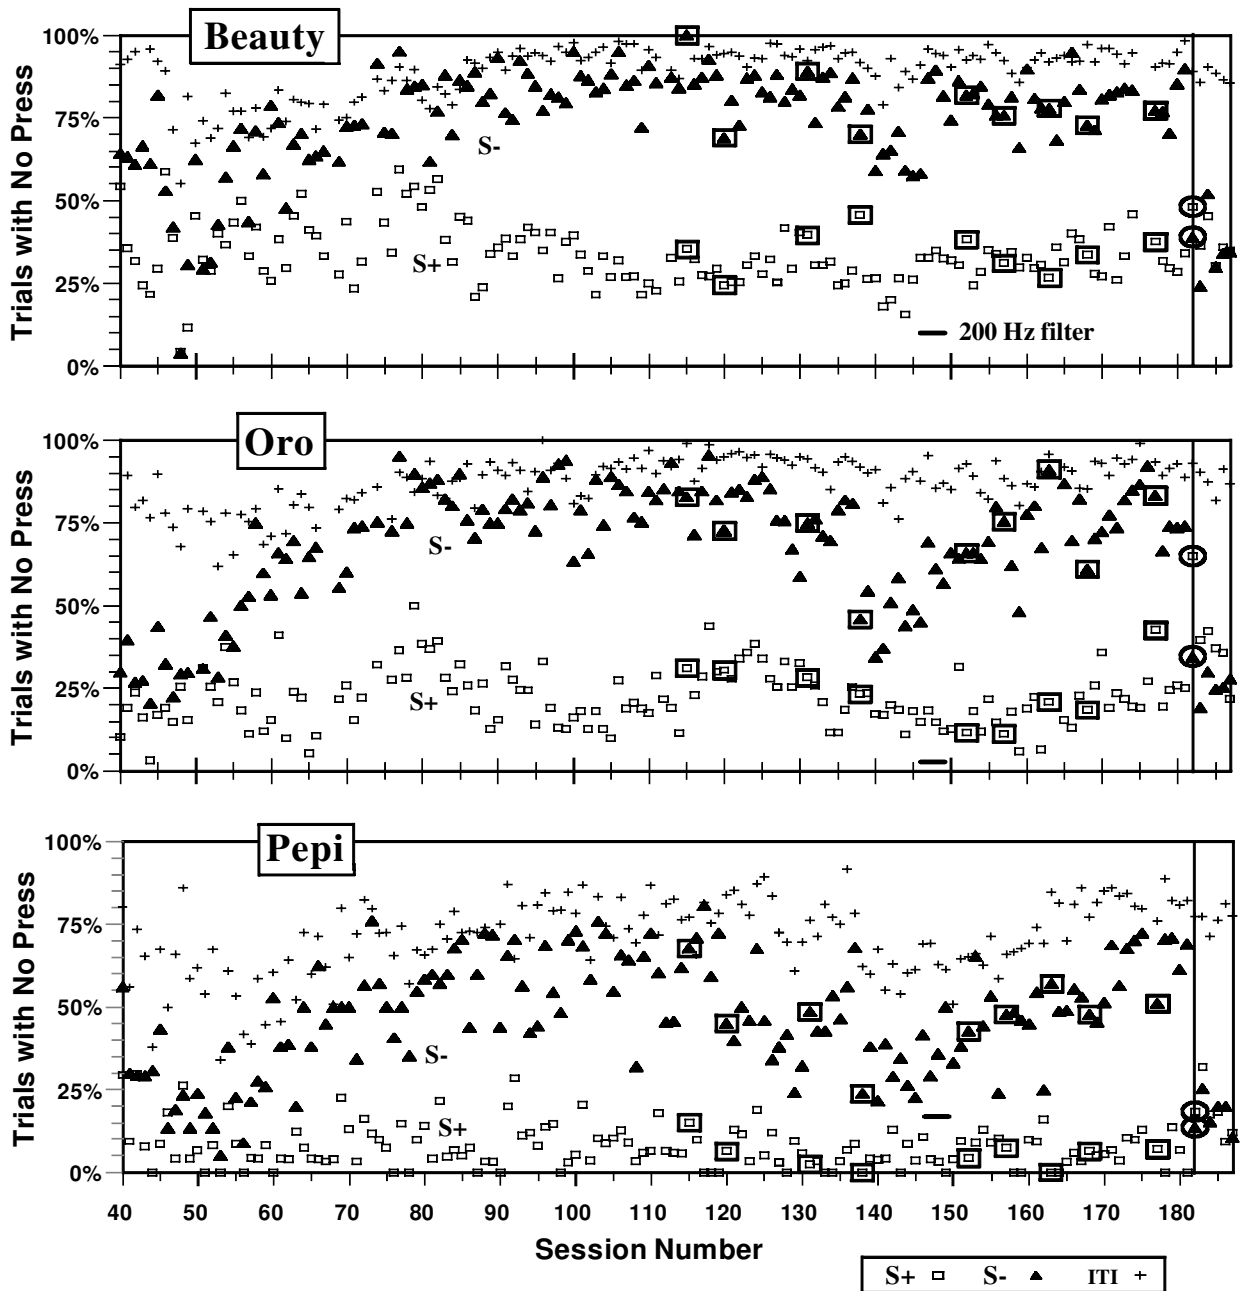
Figure 5. Experiment 1 (categorization): percentages of S+ (blues) and S− (classical) trials and intertrial intervals (ITIs, silence) during which the fish did not press the button at all, plotted for each session. The greater the separation between a session's S+ and S− nonresponse percentages, the more clearly the subject was discriminating the stimuli, as opposed to responding indiscriminately and receiving feedback from the reinforcement contingency. The boxed data points indicate probe sessions. The bar from Sessions 147–149 shows where the 200-Hz low-frequency cutoff filter was used. The circled data points (Session 182) indicate the first control-by-reversal session. In the reversal sessions (182–192), S+ (still represented by open squares) was now classical music, whereas S− (still represented by filled triangles) was blues. Crosses denote silence.

sensitive, the more information is available from which to discern patterns and changes, indications that animals learn more than the minimal solution to classification problems involving complex stimuli suggest that, as compared with trying to ignore as much as possible of the world, evolution favors embracing its complexity.

EXPERIMENT 2
Iterated Reversals

Method

The first categorization session in which the contingencies were reversed had served as a control procedure. This experimental con-
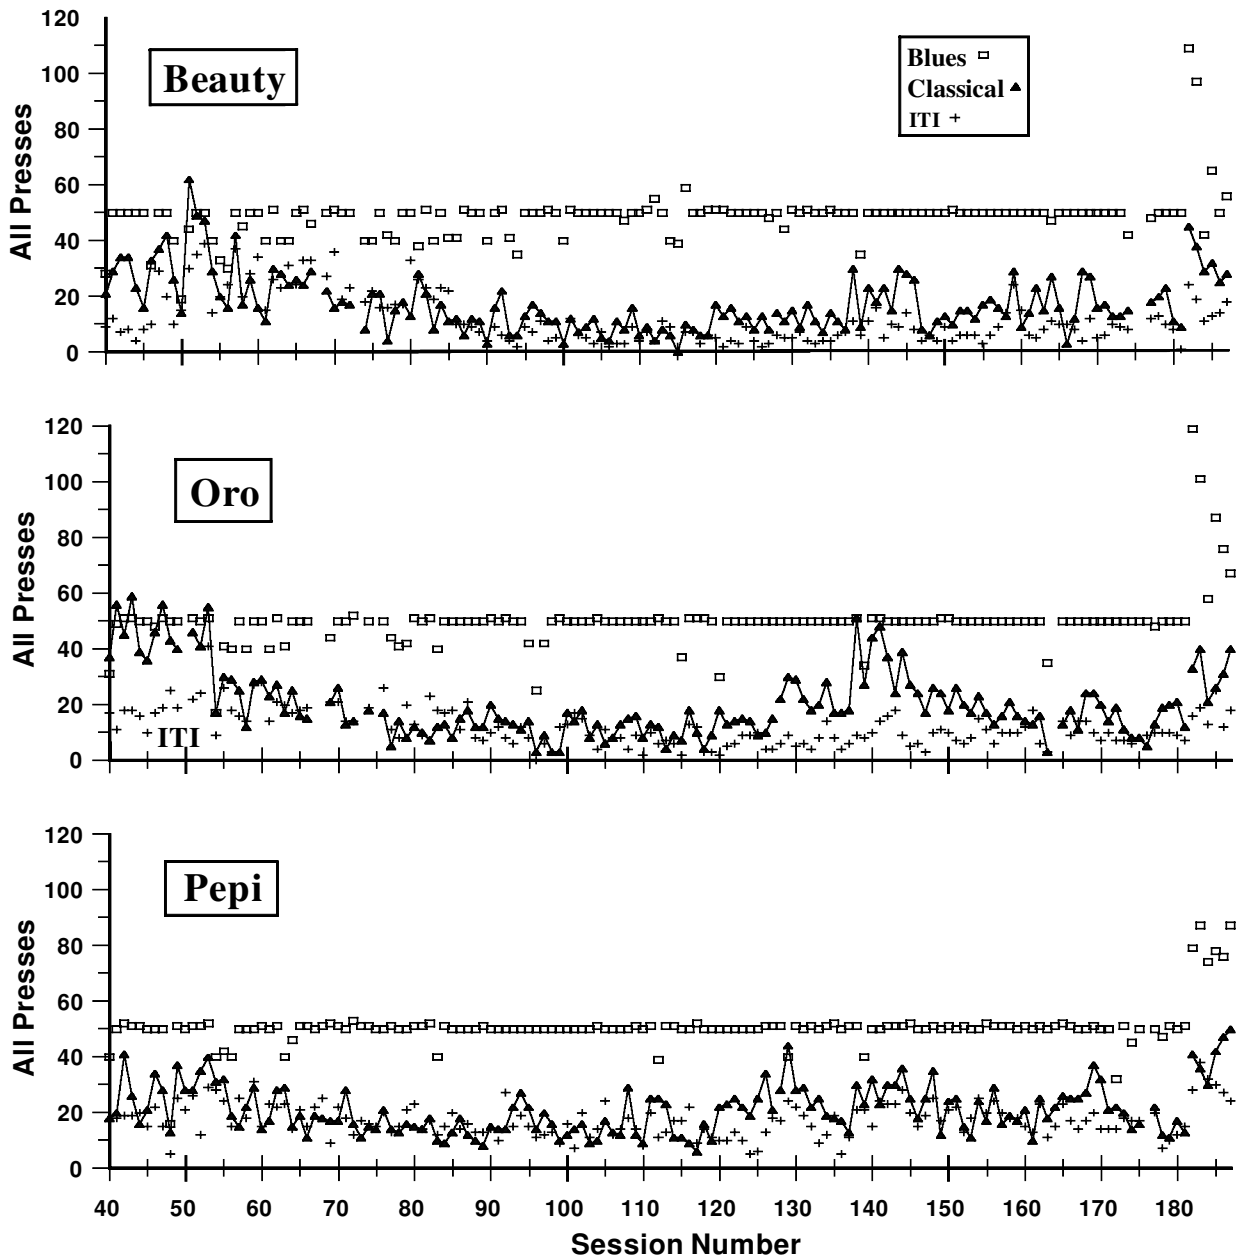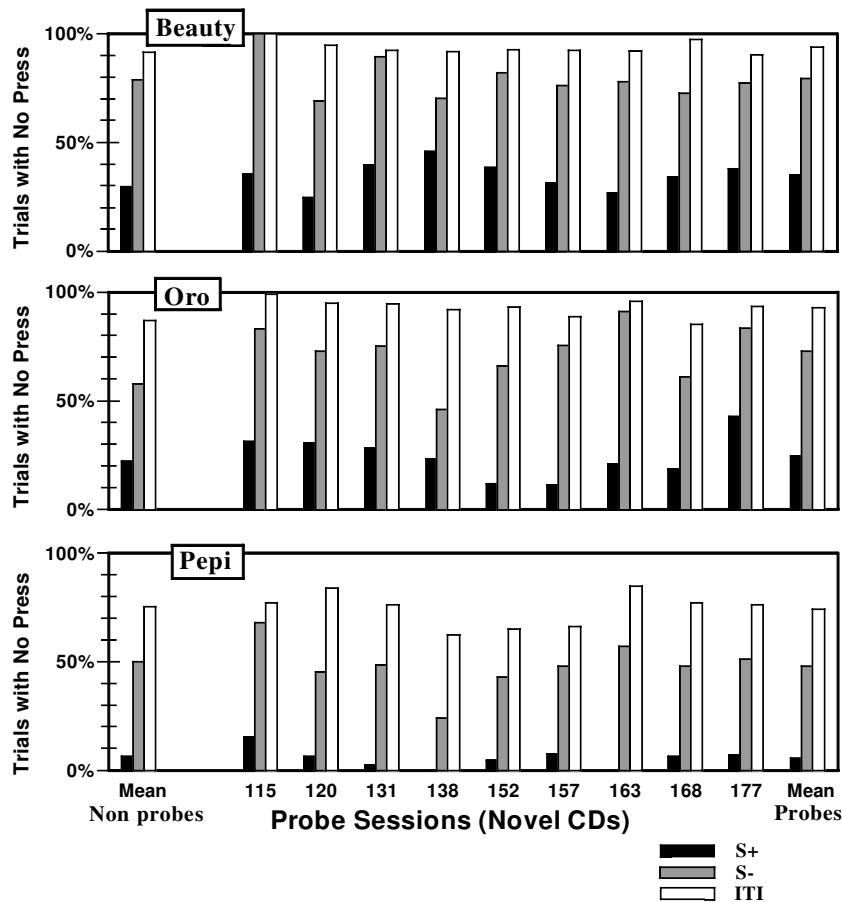
Figure 6. Experiment 1 (categorization): all responses. Number of responses in each session during blues trials, classical trials, and intertrial intervals (ITIs). Under the reversal condition that began with Session 182, blues became S− and classical became S+. Blues response counts above 50 reflect the onset of extinction. Reversal sessions in which the classical response count was less than 50 were terminated at 90 min. Crosses denote silence.

dition (i.e., with blues now the S− and classical now the S+) was then extended into an iterated-reversal experiment. The primary goal was to test koi in reversal learning with complex auditory stimuli, after having found nothing in the literature on the subject of auditory reversal learning in fish.

A secondary goal was to prepare for a planned experiment with a classification procedure involving two visually distinct buttons, with responses on a particular button to be reinforced only in the presence of a particular type of music. The rim of the original button had been machined from black Delrin. A new rim was made from an ivory-

colored blank of the same plastic and installed after six sessions of the unsignaled reversal (i.e., no change in the button). Thereafter, rim color alternated in synchrony with contingency reversals, with black and ivory signaling that S+ was blues and classical, respectively. The fish would experience a total of four color changes, the first of which did not signal a contingency change.

The sessions were terminated by either the trial in which the 50th pellet was dispensed or an elapsed time of 90 min, whichever came first. Altogether, there were 119 sessions spanning four reversals (see Figure 4).

Figure 7. Experiment 1 (categorization): categorization performance during probe sessions (i.e., all-novel stimuli); percentages of S+ trials, intertrial intervals (ITI), and S− trials in which the fish did not respond. Black bars indicate S+, gray bars indicate S−, and white bars indicate ITI. Low nonresponse percentages for S+ indicate a high likelihood of responding during S+, whereas high nonresponse percentages for S− indicate that the fish tended to ignore the button during S−. The leftmost "Mean" bars show the averages of all the nonprobe sessions in the session range 115–181, and the rightmost "Mean" bars show the averages of the probe sessions.

Just before the first reversal, the stimuli had been a *House of Blues* compilation (S+) and a baroque horn compilation (S−). When the contingencies were reversed, making blues S− and classical S+, both stimuli were also replaced. The first reversal stimuli were Bach oboe concertos (S+) and John Lee Hooker's *Real Folks Blues* (S−), both very familiar to the fish. In the next three phases of the first reversal, three additional pairs of familiar CDs were used. In the fifth and final phase of the first reversal, the fish were exposed to entirely novel stimuli (S−, *The Blues Story* compilation; S+, Bach trio sonatas). As in the categorization study, the first session with novel stimuli (Session 224) served as a probe, testing for open-ended categorization. Blues and classical successively exchanged roles as S+ and S− in three additional reversal phases, all using familiar stimuli. As can be seen in Figure 4, when the first and third reversals occurred, at least one stimulus CD changed as well. The final reversal brought the contingencies back to where they had been before the first reversal—namely, with blues as S+ and classical as S−.

Results

The first reversal (Session 182) produced considerable behavioral disruption and erroneous responding (see Fig-

ures 8 and 9). In the first 6–9 sessions, the increase in responding to the genre newly become S+ was insufficient for any fish to earn 50 reinforcements within the 90-min time limit (Figure 8). Beauty and Oro steadily improved and were reliably discriminating by the 16th and 23rd sessions, respectively; Pepi achieved significance by Session 9, but then his performance was substandard until Session 27 [5 sessions; $ts(4) > 2.29$; $p < .010$, $p < .009$, and $p < .042$, respectively]. Near the end of the first reversal, the probe session (224) revealed that the classification behavior exhibited by Beauty and Oro was clearly open-ended categorization. With each successive reversal, it took fewer sessions for Beauty and Oro to demonstrate clear discrimination (Figures 9 and 10). Even Pepi eventually learned each reversal—albeit, while continuing to call attention to the existence of considerable and systematic individual differences among fish. In Figure 9, least-squares regression lines are plotted for S− nonresponse data, and as is shown in Figure 10, their slopes indicate improvement rates of
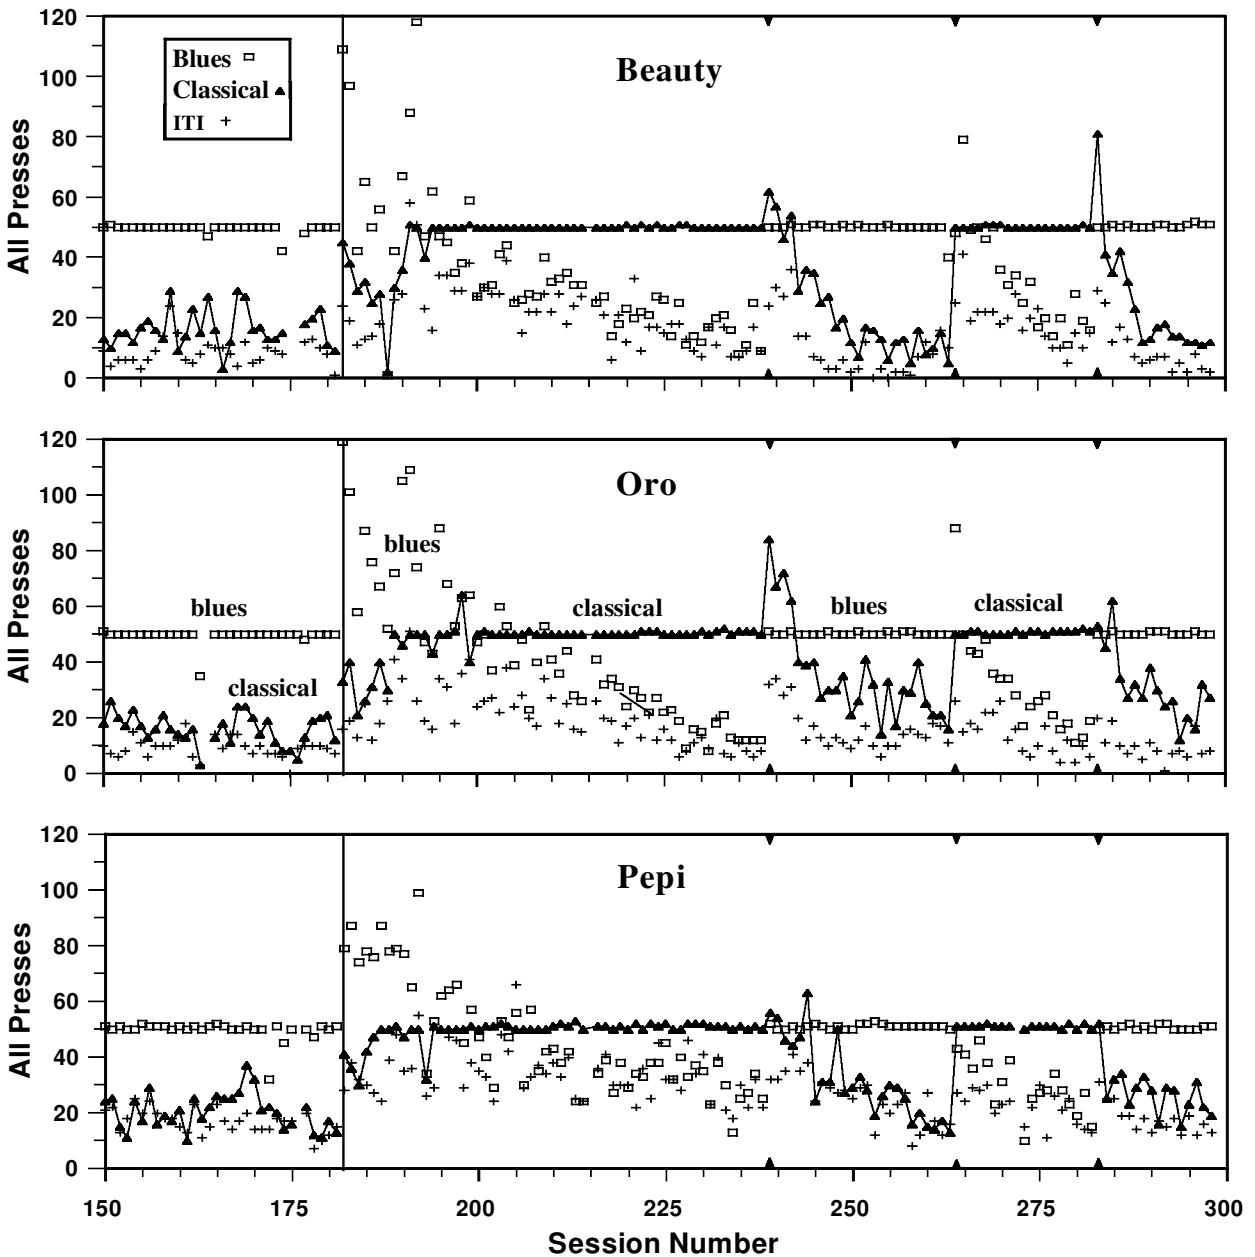
Figure 8. Experiment 2 (iterated reversals): counts of all responses, by session. Note that in graphs of reversal data, the symbols consistently indicate genre (open squares, blues; filled triangles, classical), not stimulus semantics (S+ or S−), which of course change at each reversal. This convention highlights the fluctuating response pattern associated with each genre longitudinally as successive reversals of reinforcement contingency produce disruption and then learning with increasing efficiency. The primary session termination trigger was the trial on which the 50th pellet was delivered, but in the early sessions of the first reversal, no subject made as many as 50 "correct" responses to the former S− within the session time limit, even while responding many more than 50 times, without reinforcement, to the former S+. Note that the extinction-induced response frenzy associated with a genre's suddenly becoming S− decreases with each reversal. Crosses denote silence.

0.8% and 1.1% per session for Beauty and Oro, respectively, during the first reversal (Sessions 182–238). Over the final 5 sessions of each reversal phase, discrimination by Beauty, Oro, and Pepi was statistically significant at levels always better than $p < .010$, $p < .016$, and $p < .042$, respectively [$ts(4) > 2.29$]. It can be seen in Figure 8 that with each reversal, the increase in responding to the new S+ was generally faster than the extinction to the former S+.
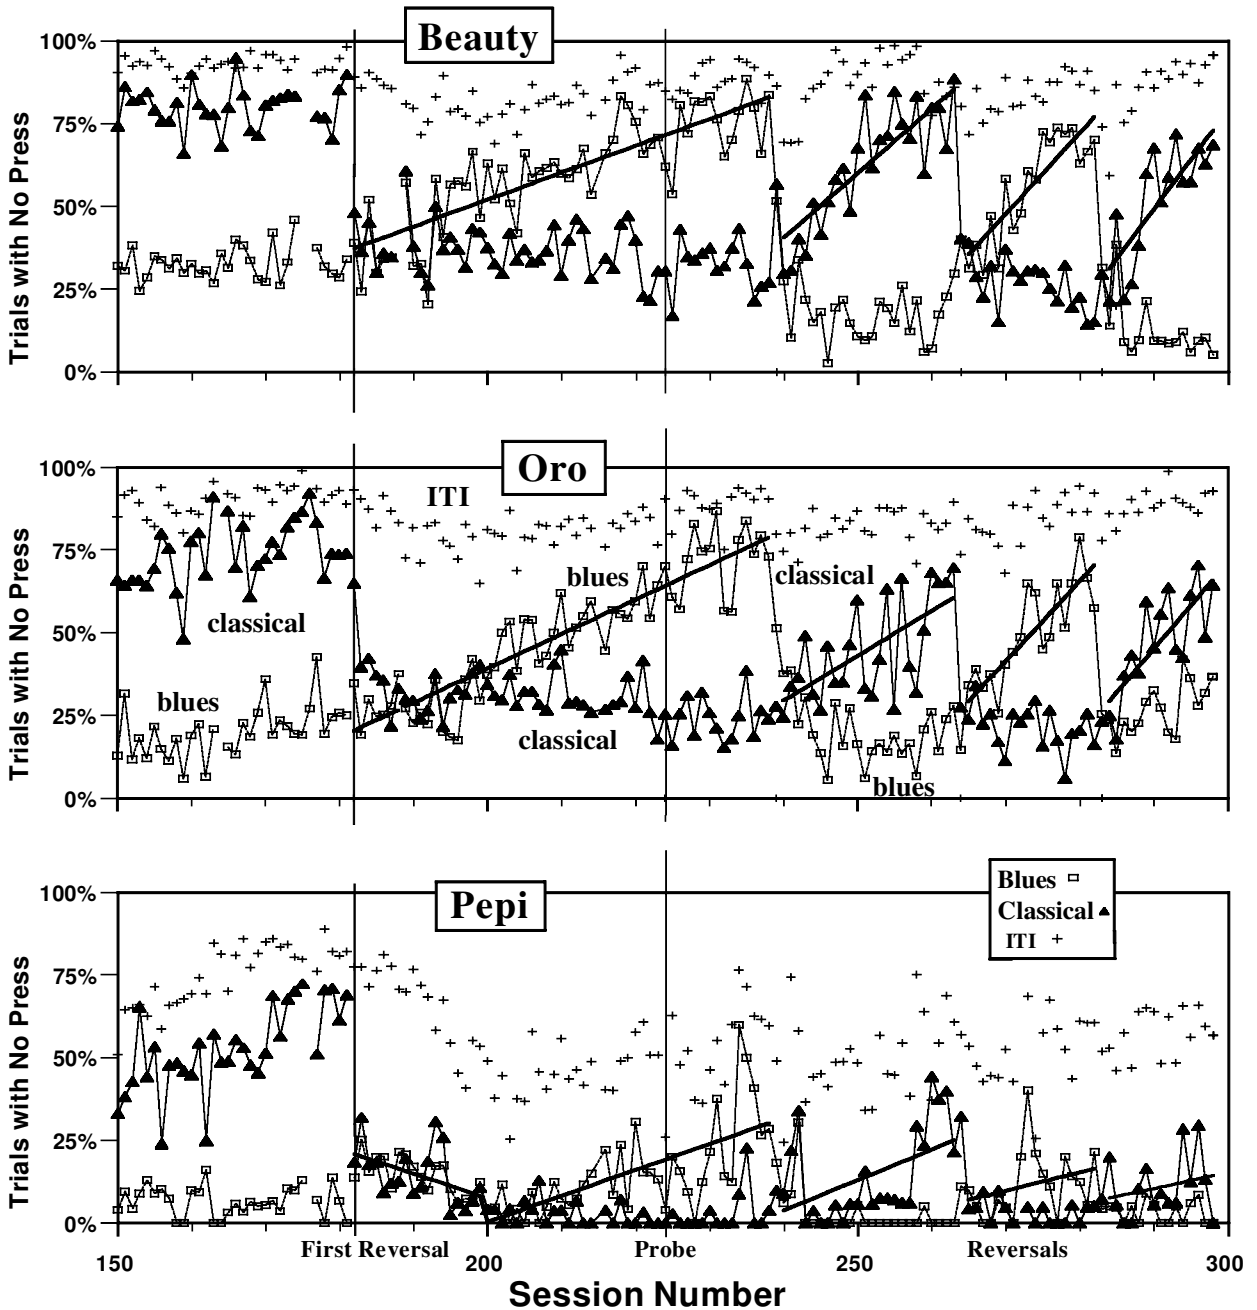
Figure 9. Experiment 2 (iterated reversals): percentages of trials with no responses, by session. Note the increasingly efficient learning of nonresponding to each former S+, as shown by the increasing slopes of the least-squares regression lines. Session 224 (dashed line) was a probe; the accurate discrimination of entirely novel stimuli indicates that the fish were performing open-ended categorization. Crosses denote silence.
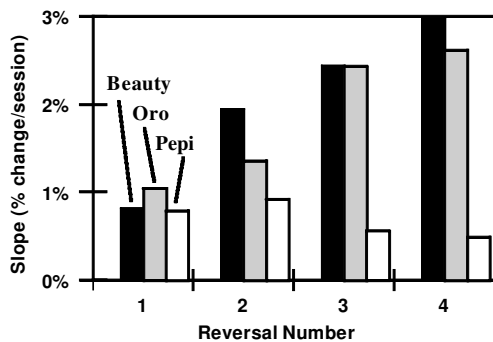
EXPERIMENT 3
Transition to Synthetic Keyboard Music and Polyphonic Discrimination with Identical Timbre

The next step was to try to extend the learned classification of natural music to MIDI (Musical Instrument Digital Interface) synthesized music in which S+ and S− were not distinguishable by timbre.

Method

Pretraining. After the iterated-reversal experiment, the fish remained in the home/experimental tank, from which all the apparatus had been removed, and they were not exposed to music or to any training or experimental procedures for a period of 10 months. The

Figure 10. Experiment 2 (iterated reversals): improvement trends. For Beauty and Oro, the improvement rate (Figure 9 regression line slopes, percentage of nonresponse S− trials per session) increases with each reversal.

work resumed with a reintroduction to the apparatus and protocol as left off at the end of the fourth and final reversal—that is, with S+ and S− being blues and classical, respectively. However, the fish were also being prepared to encounter the synthetic stimuli that were being created at that time, which consisted of a single short piece from each genre, rather than a collection comparable in duration and variety with the commercial CDs used previously. Therefore, in pretraining, the stimuli were the same two CDs the subjects had last heard (see Figure 4, Session 300), but only a single track from each. S− was an excerpt from a Bach oboe concerto. S+ was the John Lee Hooker song "One Bourbon, One Scotch, and One Beer," of which a synthesized version was in the works.

Initially, Beauty and Oro seemed to manifest effects of the long hiatus, in the form of low response rates (not enough to earn 50 pellets in a session) and some transient difficulties operating the button and taking pellets from the nipple. However, even allowing for some relearning and the reduced size of the new stimuli, the immediacy with which all 3 subjects exhibited highly accurate discrimination suggests that the original discrimination training had been retained (compare the leftmost sessions in Figure 11 with the rightmost in Figure 9).

Procedure. The new stimuli were musically simplified MIDI keyboard renditions of music of the types the fish had already learned to classify. S+ was the familiar blues song described above, and S− was the unfamiliar Fugue #2 from Book I of Bach's *Well-Tempered Clavier*. To minimize the presence of extraneous cues, all the performances were done by the same musician, with the same equipment (a Kurzweil Model k2000vp keyboard, Young Chang Research and Development Institute, Waltham, MA), the same sequencing software (Studio Vision Pro Version 4.0, Opcode Systems, Nashville, TN), and the same settings for volume and for timbre (acoustic piano). Because the new stimuli were so short that they would repeat after only a few trials, rote memory could play a role in their discrimination, and thus the operative level of stimulus control was being allowed to potentially regress from categorization to discrimination.

The initial synthetic stimuli turned out to trigger a deterioration in classification accuracy (Figure 11, Session 13). It seemed possible that this was an artifact of the musical translation process, because even humans had some difficulty recognizing the John Lee Hooker song in synthetic form: It had gone from bluesy to bouncy, its original bass inflections were absent, and there was no melodic voice comparable with that of the singer. Therefore, at Session 24, I introduced a second generation of stimuli. S+ remained the same piece, but the performance was somewhat more faithful to the model. The old S− (the Bach fugue) had been much shorter than S+ and thus had been repeating more often, so it was replaced by a syn-

thesized rendition of the first movement of Vivaldi's Concerto Grosso in A minor, Opus 3, No. 6, whose duration more closely matched that of S+. In addition, the timbre of the blues was changed from acoustic piano to rock piano, thus adding timbre as an available discriminant. However, from Session 47 until the end of the experiment, S+ and S− again had timbre in common, specifically rock piano (starting in Session 47), chiffloot (starting in Session 55), and guitar (starting in Session 62).

Results

With the first presentations in a synthetic timbre, discrimination accuracy deteriorated (Figure 11, Sessions 13–23), perhaps because the first synthetic S+ was not a good enough exemplar of the blues genre. Pepi essentially ceased responding during this time and had to be dropped from the study. Starting at Session 24, with improved stimuli and a difference in timbre between S+ and S−, the behavior of both remaining subjects immediately changed, and thereafter their discriminative accuracy improved continually.

When S+ and S− were again made identical in timbre (Session 47), accurate classification persisted, despite a performance decrement probably related to the loss of a salient cue. When the same pair of MIDI sequences was presented in additional timbres, generalization of accurate classification to the new sounds appeared seamless. The data from Session 47 and onward demonstrate [Sessions 18; ts(17) > 16; ps < .001] that koi are capable of discriminating polyphonic music on the basis of features other than the timbre of individual tone sources.

EXPERIMENT 4
Discrimination of Melodies Without Timbre Cues

After observing that instrumental timbre seemed not to be a necessary discriminant in polyphonic music, I investigated whether timbre was necessary for koi to discriminate single-voice music—that is, melodies.

Method

In order to familiarize the subjects with the S+ melody, and also in hopes of rehabilitating Pepi, this experiment began with several sessions of collective exposure to a sound-versus-silence contingency. CRF, signaled by the music that would later be the S+ in the sound-versus-sound discrimination, alternated with extinction, signaled by silence. In this group pretraining, Pepi was observed both pressing the button and taking pellets from the nipple, but when subsequently isolated for melody discrimination sessions, he consistently shied away from the nipple and again had to be dropped as a subject.

The stimuli for the melody discrimination experiment (Figure 12) were all derived from the theme from Paganini's 24th Caprice for solo violin, which is perhaps best known from Rachmaninoff's "Rhapsody on a Theme by Paganini." S+ was the original melody, as performed on the MIDI keyboard. S− was derived from S+ by using the transformation described by Poli and Previde (1991). By using a computer to modify the MIDI score, S+'s order of pitches was reversed, whereas its original order of note durations was retained. In other words, the first note of S− got the duration of S+'s first note but the pitch of S+'s last note, the second note of S− got the duration of S+'s second note but the pitch of S+'s next-to-last note, and so on. This procedure can be thought of as retaining the rhythm track while
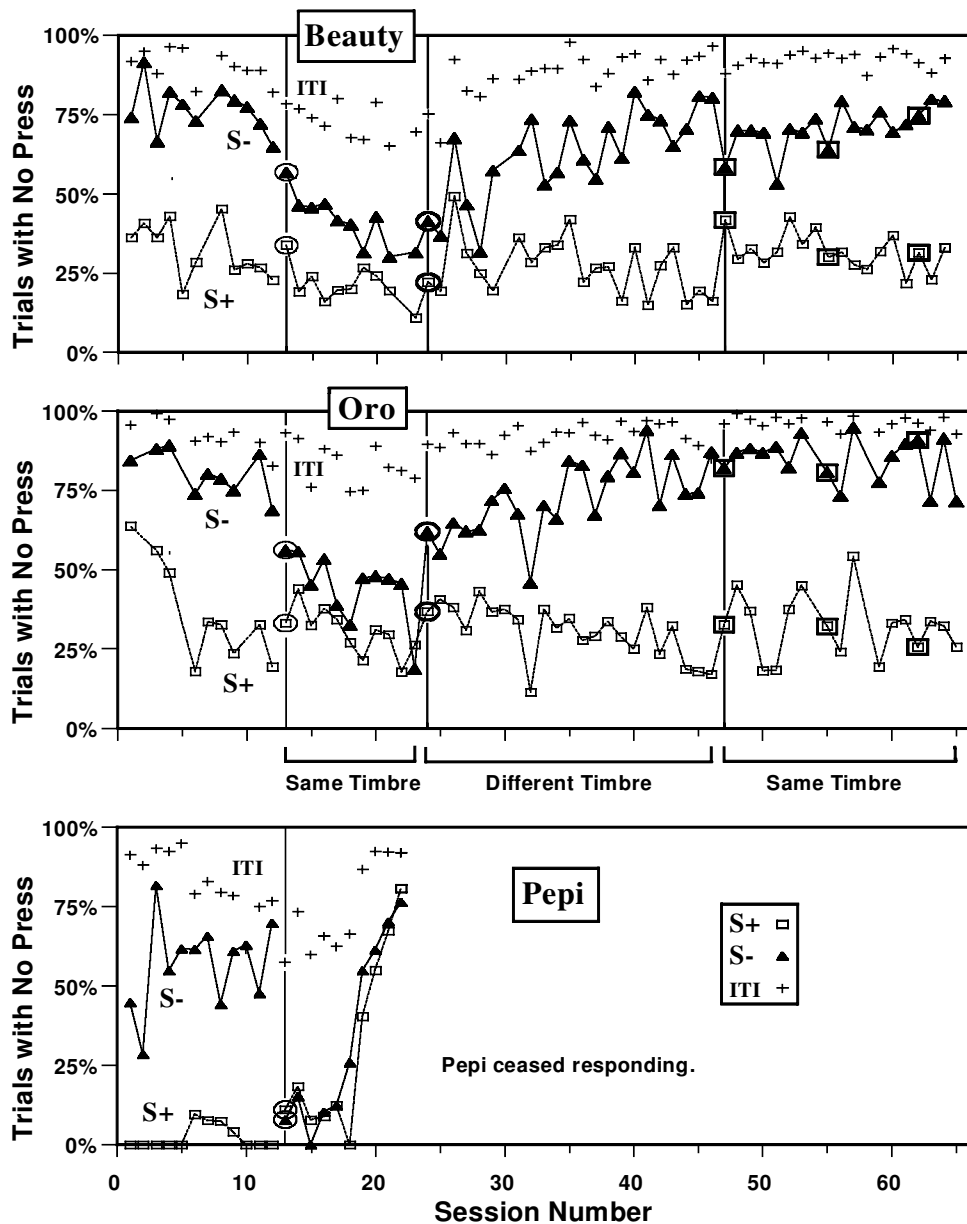
Figure 11. Experiment 3 (transition to MIDI and timbre study): percentages of S+ and S− trials during which the fish did not press the button. During the first 12 sessions, the stimuli consisted of a single track from each of the two commercial recordings the subjects had last heard 10 months earlier at the conclusion of the iterated-reversal experiment; the immediate resumption of accurate discrimination suggests long-term memory of training. The circled data points (Sessions 13 and 24) indicate the introduction of the first- and second-generation synthetic music. From Session 47 until the end of the experiment, S+ and S− had identical timbre; each boxed data point pair indicates the first presentation of a new timbre: Session 47, rock piano; Session 55, chiffloot; and Session 62, guitar. Crosses denote silence.

reversing the pitch track. The result is two stimuli that are "identical with respect to their mean frequency, note duration and rhythm, but noticeably different in their melodic patterns" (Poli & Previde, 1991, p. 12). Such a pitch-track-reversed melody differs from a completely backwards melody only in rhythm, sacrificing the linkage between each pitch and its original duration in order to preserve the original sequence of durations per se. (If the original melody—unlike the

Rachmaninoff—has a palindromic rhythm track, the pitch-track-reversed melody and the completely backwards melody are identical.) The melodies were recorded using the same MIDI apparatus as that described above. To enhance clarity, the tempo was somewhat slower than that of the usual piano performance.

Over the course of 53 sessions, the instrumental timbre alternated six times between strings and acoustic piano, but in any given ses-

Paganini Variations

Paganini Variations, Pitch Reversed

Figure 12. Experiment 4 (melody discrimination): musical scores of the normal (S+) and pitch-reversed (S−) melodies.

sion, S+ and S− always had the same timbre. The discrimination protocol was unchanged from the previous experiments, except that the session termination trigger was raised from 50 to 55 pellets. As a preliminary test for a possible perceptual invariance (Hulse, Takeuchi, & Braaten, 1992), S+ was transposed down a fifth in Sessions 51–52.

After 53 sessions, it was discovered that the MIDI technology had introduced artifactual differences between S+ and S−. A musically attuned visitor to the lab detected minor variations in note durations between the two stimuli, which apparently had resulted from incorrect handling of the attack–sustain–decay envelopes of S+'s recorded notes. There also turned out to be subtle volume differences, which were said to have resulted from the combination of the MIDI score editing process and the volume compression function in the digital audio tape (DAT) recorder used to create stimulus master tapes. These differences were confirmed by careful examination of the waveforms, using a computer oscilloscope. Therefore, a second generation of S− was created by direct numerical substitution of the S+ pitches into the MIDI score in reverse order and by recording a new master tape with all volume compression disabled. The corrections were verified by waveform analysis.

After eight sessions with the second-generation S−, both S+ and S− were shortened by the removal of each one's distinctive initial note, and seven sessions later they were further shortened by the removal of their remaining initial six and final four notes. The purpose of these stimulus contractions was to look for behavioral disruption that might

indicate that the subject had been discriminating the melodies on the basis of such local features as the absolute pitch of the initial note.

Results and Discussion

Beauty's performance inexplicably deteriorated to the point where he had to be dropped as a subject.

Figures 13 and 14 report the performance of Oro, the only fish to complete the melody experiment. The data show that without timbre cues, Oro clearly discriminated between the two melodies [in each stimulus condition (i.e., second-generation S−, first-note-removed, and both-ends-clipped), 5 final sessions, $t(4) > 10$, $p < .001$].

Oro responded more to S+ than to S− even in the very first session with melodic stimuli. Familiarization with the S+ during pretraining could well have contributed to this response bias, perhaps partly because music heard in a sound-versus-silence context had consistently become the S+ in a subsequent sound-versus-sound discrimination. However, the MIDI artifacts in the first-generation stimuli might also have played a role.

Session 54, in which the corrected S− was introduced, marks a significant reduction of the separation between
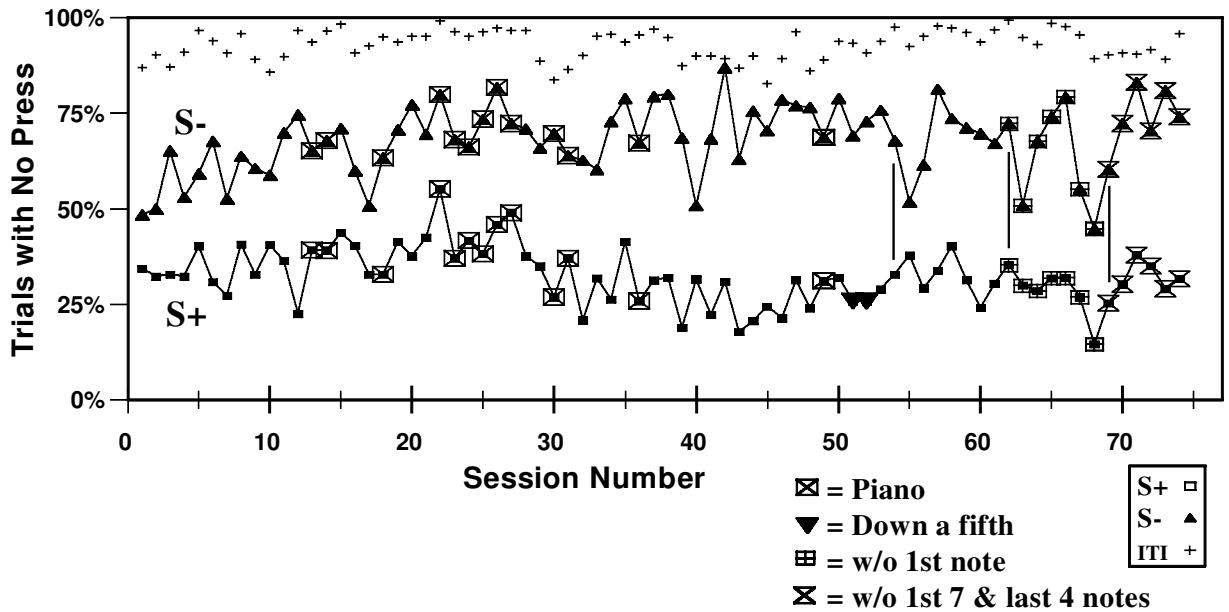
Figure 13. Experiment 4 (melody discrimination): trials with no responses. The MIDI timbre was violin, except for 13 sessions of piano, as marked. The corrected S— was introduced at Session 54. The first notes of S+ and S— were removed at Session 62. The six initial and four final notes were removed at Session 69. Crosses denote silence.

the S+ and the S— nonresponse percentages [two-tailed, two-sample *t* test, assuming unequal variances, contrasting the mean difference between S— and S+ nonresponse percentages before and after the change; 8 observations just before the stimulus change vs. 8 observations just after; *t*(11) = 2.77, *p* < .019]. This implies that the rather subtle artifacts in the original S— had indeed been exploited as discriminants.

One oddity is the fact that although the stimuli changed significantly in Sessions 54 (second-generation S—) and 62 (removal of the first note of S+ and S—), in each case

a pronounced anomaly in Oro's performance occurred in the following session.

Studies of melody discrimination in animals have often found that the behavior was controlled by local features, such as timbre, pitch register, or the absolute pitch of the initial note, rather than by the melodic patterning per se (D'Amato & Salmon, 1984; Hulse & Cynx, 1985; Hulse & Page, 1988; Poli & Previde, 1991). (In this paper, for the term local properties, I follow the convention of referring not only to features like distinctive individual notes, which are clearly local in a temporal sense, but also to features
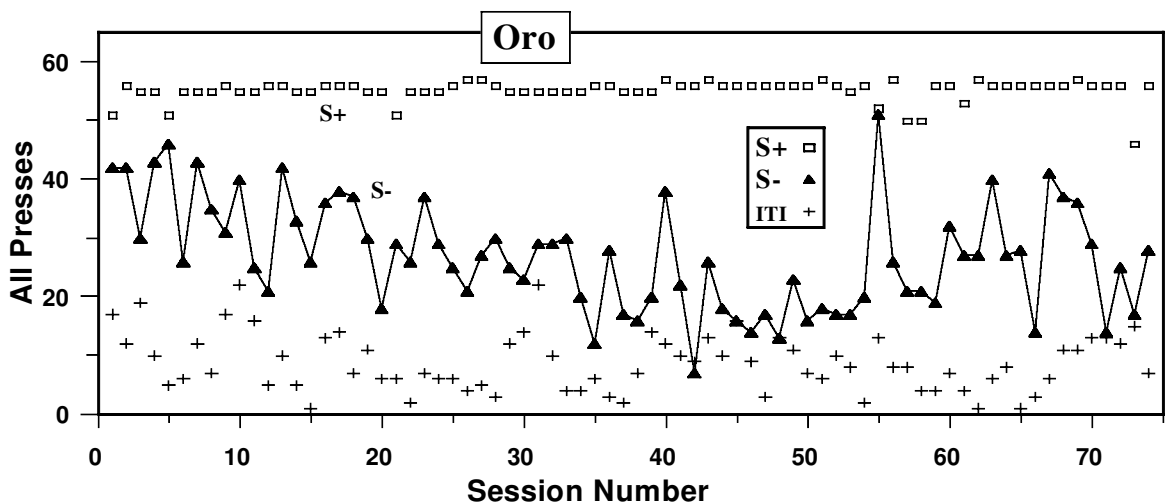

Figure 14. Experiment 4 (melody discrimination): all responses. Sessions terminated with the S+ trial in which the 55th pellet was dispensed. Crosses denote silence.

like timbre and pitch register, which, although recognizable in any local segment, can be constant over the duration of the stimulus and which, therefore, might arguably be called global properties.) It should not be surprising that with sequential stimuli, such as sound, animals often exploit local features in preference to global pattern features. In the acoustic domain, features that are instantly recognizable presumably confer a survival advantage not only on listeners, but also on creatures in whose interest it is to be identifiable by sounds they make. Oro did not have timbre available as a discriminant; but there is no conclusive evidence that he did or did not use other local features, and it is not known which if any global features he might have used.

It warrants mention that in this type of discrimination experiment, inaccurate responding need not imply that the subject is musically insensitive because musical sensitivity can actually confound discriminative accuracy if both stimuli are musically plausible or stylistically similar. Indeed, throughout this experiment, the stimuli were monitored through a speaker in the laboratory office, and S+ and S− became so familiar, and the pitch-track-reversed S− sounded so pleasing, that human observers occasionally became confused as to which reinforcement contingency was in force in the tank.

The finding that Oro did not need timbre cues to discriminate melodies contrasts sharply with the result in rats reported by Poli and Previde (1991). Using as stimuli forward and pitch-track-reversed versions of the melody "Frère Jacques," they tested for discrimination on the basis of melodic pattern. The rats had neither previous experimental experience nor prior familiarity with music. In an attempt to preclude the exploitation of local features, such as the absolute pitch of distinctively positioned notes, each 15-sec stimulus presentation began at a randomly chosen position within the melody. Different groups of subjects were asked to discriminate two melodic stimuli that differed in melody only, in timbre only, or in both attributes. However, the stimuli were never heard as unitary melodies from start to finish and were presented in a continuous stream of randomly sequenced exposures, with responses in the presence of S+ being reinforced only on an FR-5 schedule. This would seem to be such a formidable task as to make it unsurprising that animals failed except when timbre, a consistent property of every note in each stimulus, was available as a discriminant. However, Poli and Previde did suggest that rats might well prove able to discriminate melodies without timbre cues under other conditions.

CONCLUDING REMARKS

The distinction between local features and pattern features, although it arose in connection with melody discrimination, is equally relevant to the categorization experiment. To say that, in classifying music as blues or classical, the koi might have been responding to deep generic attributes is not to say that they were necessarily using anything other than local cues, because the diag-

nostic property could have been a particular combination of timbre, loudness, pitch, brief rhythmic or melodic figures, and other such local features. A human scanning a radio dial requires only a fraction of a second to identify the genre of a given station, hardly enough time for pattern features to manifest themselves, so a brief sample of music could well contain something recognizable to an animal that had a rote store of auditory feature memories. Local features might control classification behavior even when there is extensive exposure to a stimulus. Conceivably, animals could acquire a functional understanding of music entirely on the basis of the distribution and frequency of various combinations of essentially local features, without regard to the artistic pattern features in terms of which musical genres are commonly defined.

Prior to this series of experiments, the prevailing opinion appeared to be skepticism as to whether koi could discriminate one piece of music from another under any circumstances. Now it appears that these animals can discriminate polyphonic music, discriminate melodic patterns, and even classify music by artistic genre. As far as I know, these experiments presented the most complex auditory stimuli to which fish have ever been shown capable of making sophisticated discriminative responses. However, there remains the task of elucidating the stimulus control functions of the constituents of that complexity, such as pattern features, local features, and simple properties of the stimulus as a whole. The present work was intended as exploratory rather than explanatory, and the results clearly extend the known envelope of auditory capabilities in fish.

REFERENCES

BARON, M. R. (1965). The stimulus, stimulus control, and generalization. In D. I. Mostovsky (Ed.), *Stimulus generalization* (pp. 62-71). Stanford, CA: Stanford University Press.

CHASE, A. R., & HILL, W. (1999). Reliable operant apparatus for fish: Audio stimulus generator, response button, and pellet-dispensing nipple. *Behavior Research Methods, Instruments, & Computers*, **31**, 470-478.

COOMBS, S., JANNSEN, J., & MONTGOMERY, J. C. (1992). Functional and evolutionary implications of peripheral diversity in lateral line systems. In D. B. Webster, R. R. Fay, & A. N. Popper (Eds.), *The evolutionary biology of hearing* (pp. 267-294). New York: Springer-Verlag.

D'AMATO, M. R., & SALMON, D. P. (1982). Tune discrimination in monkeys (*Cebus apella*) and rats. *Animal Learning & Behavior*, **10**, 126-134.

D'AMATO, M. R., & SALMON, D. P. (1984). Processing of complex auditory stimuli (tunes) by rats and monkeys (*Cebus apella*). *Animal Learning & Behavior*, **12**, 184-194.

FAY, R. R. (1988). *Hearing in vertebrates: A psychophysics data-book*. Winnetka, IL: Hill-Fay Associates.

FAY, R. R. (1992). Analytic listening in goldfish. *Hearing Research*, **59**, 101-107.

FAY, R. R. (1994). Perception of temporal acoustic patterns by the goldfish (*Carassius auratus*). *Hearing Research*, **76**, 158-172.

FAY, R. R. (1995). Perception of spectrally and temporally complex sounds by the goldfish (*Carassius auratus*). *Hearing Research*, **89**, 146-154.

FAY, R. R. (1998). Auditory stream segregation in goldfish (*Carassius auratus*). *Hearing Research*, **120**, 69-76.

FAY, R. R., & REAM, T. J. (1986). Acoustic response and tuning in saccular nerve fibers of the goldfish (*Carassius auratus*). *Journal of the Acoustical Society of America,* **79**, 1883-1895.

HARTMAN, W. M. (1988). Pitch perception and the segregation and inte-

gration of auditory entities. In G. M. Edelman, W. E. Gall, & W. M. Cowan (Eds.), *Auditory function: Neurological basis of hearing* (pp. 623-645). New York: Wiley.

HERRNSTEIN, R. J. (1992). Levels of stimulus control: A functional approach. In C. R. Gallistel (Ed.), *Animal cognition* (pp. 133-166). Cambridge, MA: MIT Press.

HERRNSTEIN, R. J. (1994). *Animals as classifiers*. Lecture given at the Rowland Institute for Science, Cambridge, MA, 20 April 1994.

HULSE, S. H., & CYNX, J. (1985). Relative pitch perception is constrained by absolute pitch in songbirds (*Mimus, Molothrus, Sturnus*). *Journal of Comparative Psychology*, **99**, 176-196.

HULSE, S. H., & PAGE, S. C. (1988). Toward a comparative psychology of music perception. *Music Perception*, **5**, 427-445.

HULSE, S. H., TAKEUCHI, A. H., & BRAATEN, R. F. (1992). Perceptual invariances in the comparative psychology of music. *Music Perception*, **10**, 151-184.

JACOBS, D. W., & TAVOLGA, W. N. (1968). Acoustic frequency discrimination in the goldfish. *Animal Behavior*, **16**, 67-71.

MALOTT, R. W., & SIDDALL, W. (1972). Acquisition of the people concept in pigeons. *Psychological Reports*, **31**, 3-13.

NELSON, J. S. (1984). *Fishes of the world* (2nd ed.). New York: Wiley.

POLI, M., & PREVIDE, E. P. (1991). Discrimination of musical stimuli by rats (*Rattus norvegicus*). *International Journal of Comparative Psychology*, **5**, 7-18.

POPPER, A. N., & FAY, R. R. (1993). Sound detection and processing by fish: Critical review and major research questions. *Brain, Behavior, & Evolution*. **41**, 14-38.

POPPER, A. N., PLATT, C., & SAIDEL, W. (1982). Acoustic functions in the fish ear. *Trends in Neurosciences*, **5**, 276-280.

PORTER, D., & NEURINGER, A. (1984). Music discriminations by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 138-148.

SEGAL, M., & HARRISON, J. M. (1978). The control of responding by auditory stimuli: Interactions between dimensions of stimuli. *Journal of the Experimental Analysis of Behavior*, **30**, 97-106.

STURDY, C. B., PHILLMORE, L. S., PRICE, J. L., & WEISMAN, R. G. (1999). Song-note discriminations in zebra finches (*Taeniopygia guttata*): Categories and pseudocategories. *Journal of Comparative Psychology*, **113**, 204-212.

VAUGHAN, W., JR., & GREENE, S. L. (1984). Pigeon visual memory capacity. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 256-271.

WATANABE, S., & SATO, K. (1999). Discriminative stimulus properties of music in Java sparrows. *Behavioural Processes*, **47**, 53-57.

APPENDIX
Music Bibliography

Classical

Johan Sebastian Bach, *3 Oboe Concertos*, BWV 1053, 1055, and 1059, Philips 412851-2.

G. F. Handel, *Trio Sonatas*, Denon 38C37-7026.

Wolfgang Amadeus Mozart & Giuseppe Ferlendis, *Oboe Concertos*, Philips 420 179-2.

Ludwig van Beethoven, *The Sonatas for Piano & Cello*, CBS M2K 42446.

Antonio Vivaldi, *3 Cello Concertos & Sonatas*, Decca Record Co. 433 052-2.

Antonio Vivaldi, *Guitar Concertos*, Philips 412 624-2.

Franz Schubert, *Trout*, Grammophon 413-453-2.

Baroque Favourites, Naxos 8.550102.

Johan Sebastian Bach, *Brandenburg Concertos 1–3*, MCA Classics MCAD-5956.

Johan Sebastian Bach & Antonio Vivaldi, *Concerti for 2 Violins*, CBS MK 37278.

Maurice Andre, *Baroque en Famille*, EMI 5-55488-2.

J. S. Bach, C. P. E. Bach, J. C. F. Bach, & W. F. Bach, *Trio Sonatas*, CBS MK-37813.

Blues

John Lee Hooker, *Blues Before Sunrise*, Masters CLCD 61028-2.

John Lee Hooker, *The Real Folks Blues*, Chess CHD-9271, MCA. (Excluded: Track 9).

John Lee Hooker, *Boom Boom*, K-TEL 3676-2

Muddy Waters, *I'm Ready*, CBS ZK 34928.

Koko Taylor, *I Got What It Takes*, ALCD 4706.

The Best of Chess Blues, Volume 2, Chess CHD-31316, MCA.

The Blues—Volume 3, Chess CHD-9276, MCA.

Blending the Blues, MCA, MSD-36021. (Excluded: Tracks 9, 11, & 14)

Legendary Blues, Volume 1, Blues Journey SS1896.

Sonny Boy Williamson, *The Real Folk Blues*, Chess CHD-9272, MCA.

Essential Blues 2, House of Blues, (2 CDs), Platinum Entertainment Inc. 514161183 2.

The Blues Story, Masters, MCAD 61056-2.