

Neural correlates of instrumental learning in primary auditory cortex

David T. Blake*, Fabrizio Strata, Anne K. Churchland, and Michael M. Merzenich

Coleman Laboratory and Keck Center for Integrative Neuroscience, Department of Otolaryngology, University of California, San Francisco, CA 94143-0732

Contributed by Michael M. Merzenich, May 9, 2002

In instrumental learning, Thorndike's law of effect states that stimulus–response relations are strengthened if they occur prior to positive reinforcement and weakened if they occur prior to negative reinforcement. In this study, we demonstrate that neural correlates of Thorndike's law may be observed in the primary auditory cortex, A1. Adult owl monkeys learned to discriminate tones higher than a standard frequency. Responses recorded from implanted microelectrodes initially exhibited broad spectral selectivity over a four-to-five octave range. With training, frequency discrimination thresholds changed from close to one octave to about $\frac{1}{12}$ octave. Physiological recordings during the week in which the monkey came under behavioral control signaled by a drop in measured threshold had stronger responses to all frequencies. During the same week, A1 neural responses to target stimuli increased relative to standard and nontarget stimuli. This emergent difference in responsiveness persisted throughout the subsequent weeks of behavioral training. These data suggest that behavioral responses to stimuli modulate responsiveness in primary cortical areas.

In instrumental learning, an animal progressively associates its actions with future outcomes (1). One principle in instrumental learning is Thorndike's law of effect, which states that stimulus–response relationships are strengthened if their pairing leads to future positive reinforcement and weakened if their pairing leads to future negative reinforcement. This learning should be reflected by changes in associated neural response. In the current experiment, a chronic multisite recording implant technology (2) was used to document possible neural substrates of instrumental learning in the primary auditory cortex of the awake primate.

Learning-induced representational change in mature primary sensory cortex occurs in a task-dependent manner (3–8) and can include changes in receptive field size (3, 6, 7, 9), observed cortical column size (4, 8), and cortical representational area associated with behaviorally important sensory inputs (4, 7). One study (6), especially relevant to the experiments reported here, used operant training to drive cortical representational changes in the primary auditory cortex, A1, of adult owl monkeys. Animals oriented in a behavioral apparatus and attended to a series of tone pairs. If the elements of a pair differed in frequency, and the animal removed itself from the apparatus within time limits, it was rewarded with a food pellet. Misses resulted in a brief time-out. In trained animals, the cortical area that responded to the frequency range of the target and standard tones expanded. This change in representational area was inversely correlated with the progressively improved frequency discrimination thresholds.

In the present study, owl monkeys were again engaged in auditory frequency discrimination, but in a different behavioral task. Animals oriented in a listening posture and received a series of single tones, which were initially at a constant standard frequency. After the tonal stimulus stepped up in frequency to a behavioral target, the animal moved its head from the listening posture to receive a reward. Removal at other times led to a time-out. Implanted microelectrodes were used to record unit responses across a densely sampled sector of A1 throughout task

performance. This technological advance allowed sampling from a restricted set of neurons at each site for several months, more than long enough to train a primate to master and progressively improve at a behavioral task.

Neural correlates of learning were studied in parallel with psychophysical performance, and physiological changes were related to the behavioral cues for reward. The results of this study provide a more complete picture of the representational plasticity that can contribute to learning-driven changes in frequency discrimination performance.

Methods

Animal Behavior. Animals were engaged in a limited-hold reaction time behavior (6, 9). The animal initiated a trial by making an orienting response. The orienting movement was leaning the head forward to break an invisible infrared beam in front of the animal's nose and maintaining the head for variable time periods in that slightly forward position. After the orienting response criteria were met, a series of standard stimuli, identical tone pips, were delivered in the free sound field. Each tone pip was 50 ms long with 5-ms raised sinusoidal onset and offset ramps. Two to six standard stimuli were followed by stimuli delivered at the same interstimulus interval that were higher in frequency, the target tone pips. The target frequency tone pips continued to be presented to the animal until it ended a trial by removing its head from the head beam, although the trial was correct only if the animal ended the trial before the presentation of the third target frequency tone pip. The onset-to-onset tone pip interval was 250 ms in animal one and 400 ms in animal two. This task is shown schematically in Fig. 1. There were always at least two standards before the frequency changed to the target, and there could be as many as six. A trial can be conceptualized as the series A,A . . . A,B,B,B,B . . . There were two to six standards (A) and the target (B) repeated until the animal signaled the end of the hold. In each animal, two standards were chosen within the sampling range of each animal's A1 implant. In animal one, standard frequencies of 880 and 1,396 Hz were used. In animal two, standards were 1,480 and 7,040 Hz. Each trial used one of the two standards, and the choice between standards was randomized. Two standards were used to prevent the animal from adopting fixed comparison criteria. Targets were always higher than standards and were chosen so that the more difficult frequencies were close to the animal's threshold. Correct responses occurred when the animal removed itself more than 150 ms after the first target and less than 150 ms after the third target. Earlier responses were false positives; later responses were misses. False positives and misses resulted in time-outs between 2 and 10 s long. Rewards were a few drops of a vitamin C-enriched fruit drink.

Threshold was determined with standard physiological signal detection theory criteria (10). On a weekly basis, all psychophysical data for an animal were pooled. The pooling was used to increase the statistical power of the sample. Of the 200 to 300 trials performed daily, only 25% were within the range of the

*To whom reprint requests should be addressed. E-mail: dblake@phy.ucsf.edu.




Fig. 1. Schematic drawing of a trial. A trial began with an orienting response, the animal breaking an infrared beam in front of its nose. Two to six standard tones were then repeated. After the tones changed to the target frequency, which was higher than the standard frequency, the animal could remove its head from the beam to receive a fruit-juice reward.

animal's threshold to ensure the animal's behavior was largely correct. Only trials that ended in the first two target windows could be used to assess the threshold, because there could not be false positives in the third target window (a target was always presented in one of three windows). After pooling the data, the times of hits, misses, and false positives were determined for all frequency changes assessed for the week. Asymptotic performance was determined from the trials with the largest changes in frequency. Chance performance was the ratio of (hits/hits + misses) expected if the animal randomly chose a response window, as determined by the ratio (false positives/all trials). Threshold was the first frequency at which the psychometric function rose above midway between chance and asymptotic behavior. Some weeks, the animals missed at chance levels; no behavioral threshold was determined for those weeks. For each day, a determination of above or below threshold could be made for each frequency used. However, four Weber fractions were tested in 1 day, so pooling of data over weeks was done to increase the reliability of this measure. Still, some weeks the animal performed above threshold for all tested frequencies—week 5 in animal one and weeks 2 and 3 in animal two. The target frequency ranges were 1,480–4,699 Hz in animal one, and 1,480–2,793 Hz and 7,040–13,290 Hz in animal two. Control nontarget ranges were matched in octave range and were adjacent to the target ranges: 440–1,397 Hz in animal one and 2,960–6,645 Hz in animal two. These frequency ranges were never presented to the animals during frequency discrimination training. They are controls because the responses recorded from the implant before training began were approximately equally strong in the target and in these control nontarget frequency ranges.

Physiological Recordings. Data were obtained from two chronically implanted owl monkeys, *Aotus nancymae*. Microelectrodes were implanted into the physiologically defined primary auditory cortex (A1). The A1 target was 2–3 mm anterior interaural on the superior aspect of the superior temporal sulcus (11, 12). Transdural recording through cranial burr holes prior to implantation confirmed A1 response characteristics and expected tonotopy (11, 12). Array implantation surgery was performed under areflexic barbiturate anesthesia. Techniques for implantation are described in a methods paper (2). Recordings were made with parylene-insulated iridium microelectrodes (Micro Probe, Potomac, MD) with tip exposures between 5 and 7 μ m long, to maximize the probability of sampling single units (13, 14). The implant microelectrode recordings were transdural. Implant best frequencies spanned the range of frequencies found on the exposed surface and ranged from 110 Hz to 20 kHz.

Histological confirmation of A1 electrode positioning was performed in animal one. After 56 mo of implantation, the brain was perfused with electrodes in place before cresyl violet staining with standard histological procedures. Electrode positions were

confirmed in A1 by the presence of relatively dense cell body staining in middle cortical lamina on coronal sections.

After implantation, a recovery period of several weeks ensued before recording was initiated. During training and recording sessions, the primate sat in a primate chair with its head positioned 24 in in front of a free field speaker. Single units were isolated online by using the Magnet system (Biographics, Winston-Salem, NC), and 1.5 ms of spike waveform was stored for each unit discharge event beginning 0.5 ms before a voltage-threshold crossing. Single unit quality was confirmed offline by waveform analysis that used three criteria: signal-to-noise ratio, coefficient of variation (CV) of maximal positive slope on the principal waveform deflection, and CV of maximal negative slope on the principal deflection. The signal-to-noise ratio, or the mean peak-to-peak magnitude divided by the noise standard deviation, had to exceed five, and CVs for each unit had to be below 0.25. Multiunit recordings consisted of recordings that were manually selected as single units but did not meet our single-unit statistical criteria.

In animal one, implantation preceded this behavioral study by 18 mo. In animal two, implantation occurred 3 mo before this behavioral training was initiated.

Sound Presentation. All experiments were conducted in a double-walled sound attenuation chamber. Sound levels were calibrated with a Brüel and Kjaer (Brüel and Kjaer Instruments, Marlborough, MA) sound level meter by using the “A” filter. Sounds for characterization stimulus sets were created digitally, recorded on an audio CD, and played through a McIntosh (Binghamton, NY) audio amplifier. Sounds in the behavior were created digitally by using LABVIEW software (National Instruments, Austin, TX) with 100-kHz sampling rates and played through the same audio amplifier. Sounds were played from a free field speaker positioned approximately 24 in in front of the animal. Each sound in the behavior was a 50-ms tone pip with 5-ms raised sinusoidal onset and offset ramps. The onset ramps can be described by the equation $(1 - \cos(2\pi t/10 \text{ ms}))/2.0$ for $0 < t < 5$. Offset ramps are the time reverse of onset ramps. For animal one, tones in the behavior were 70 dB sound pressure level (SPL). For animal two, they were 50 dB SPL. Receptive field characterization stimuli were 100-ms duration tones. For animal one, these were isolated tones played every 1,000 ms at each of eight intensities and 84 frequencies, spaced each $\frac{1}{12}$ octave. The order of frequency and intensity was randomized. The use of tuning curve stimuli in this animal necessitated pooling the data weekly to achieve reasonable standard errors. To increase the reliability of sampling tonal responses, for animal two, tones were played at one tone per octave per 700 ms. The stimulus frequencies were again spaced each $\frac{1}{12}$ octave. Five minutes of such randomly delivered stimuli were presented at 30, 50, and 70 dB SPL to characterize frequency-selective responses at different sound amplitudes. Responses to each frequency at each intensity were sampled approximately 30 times. The response measure used in this study was the stimulus-evoked unit discharge rate recorded 5–35 ms after tone onset. Characterization stimuli were presented to the animals before each training session while they sat passively in the primate chair. Recordings in which the animal's head did not stay reasonably constant were discarded. There was no attentional control of the animal while recording characterization stimuli responses and no behavioral contingency was associated with the presentation of the characterization stimuli.

Results

Changes in frequency discrimination behavior across training had significant physiological correlates in both trained animals. With the implant recordings, each electrode sampled from the same small group of neurons over time. This sampling method allowed each animal to act as its own control for later points in

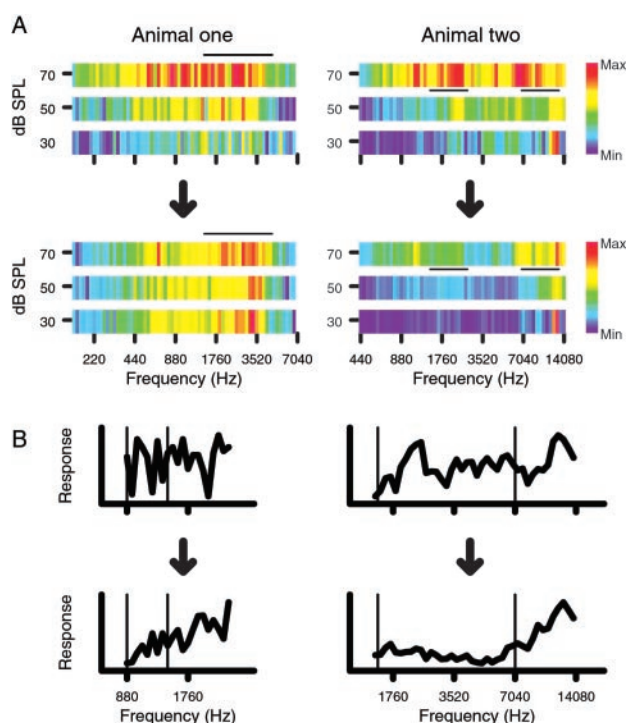


Fig. 2. (A) (Upper) The sum of all responses sampled from the array as a function of frequency and intensity in the first week of behavioral training. (Lower) Responses in the last week. The black bars indicate the range of frequencies and intensities used as targets in the behavior. (Left) Animal one; (Right) animal two. (B) The sum of all responses to different frequency ranges at the trained intensity. At the end of training, every standard frequency response was lower than the response to at least 10 of the next 12 higher frequencies. Each of the eight plots is individually normalized to its maximum and minimum response. Standard frequencies are indicated by thin vertical lines.

time. Fig. 2A shows the responsiveness of the sampled auditory cortex to pure tones that differed as a function of frequency and intensity as recorded from the implant in each animal, in the first and last week of training. The color of each pixel indicates the total number of spikes elicited from all implanted electrodes in the first 30 ms after tone onset at one frequency and intensity. The black bars mark the frequency and intensity range that contained the rewarded stimulus targets over the 6- to 7-wk training period. Animal one had two standards that were close together, so its overlapping target ranges are displayed together under the black bar in Fig. 2. At the trained intensity, responsiveness was maintained in the target frequency range, whereas responsiveness was lost in the nontarget ranges. This effect was particularly clear if the standard frequencies were compared to the nearest targets that were just higher in frequency. In animal one, standards were 880 and 1,396 Hz, and in animal two, they were 1,490 and 7,040 Hz. The standard frequencies are marked by thin vertical lines.

Animals reliably performed the task within the first weeks. Thresholds were assessed by the signal detection theory (10). Threshold was the first frequency at which performance was closer to asymptotic performance than to chance. Fig. 3 illustrates behavioral thresholds for each animal over successive training weeks. Thresholds for each animal in the first week were 12 semitones for one monkey, and no threshold could be measured for the second. By the end of training, thresholds were 1 and 0.5 semitones.

The population data for the changes in responsiveness of the implant recordings to the target and nontarget ranges are shown in Fig. 3 Lower, with a statistical summary shown in Table 1. The

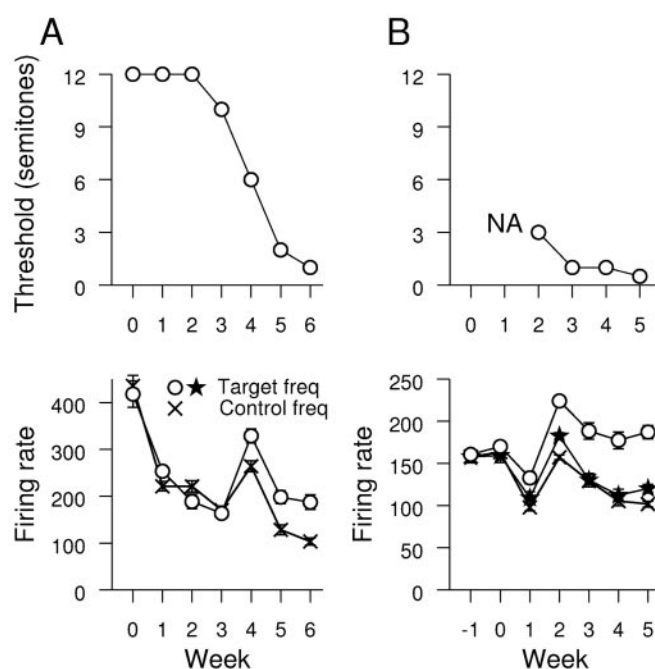


Fig. 3. Changes in threshold and selectivity with time. (A) Thresholds vs. time are shown on the top for animal one. Selectivity of target and nontarget ranges are shown on the bottom. (B) Animal two. Targets were presented in frequency ranges symbolized with a circle and star. Nontarget ranges are symbolized by an X. In all cases, significant differences emerged between the target and nontarget frequency ranges.

relevant measure for this plot is all action potentials recorded from all electrodes in the first 30 ms of response to a pure tone delivered at the trained intensity. For each day's recordings, responses to target and nontarget frequency ranges are derived from the same neurons in the same recording sessions during a period without behavioral contingencies associated with the sounds. Both animals show similar and robust responses to target and control frequency ranges at the beginning of training. A breakpoint occurred in animal one at week four and in animal two at week two, as shown in Fig. 3. At the breakpoint, responses to all frequencies were elevated, with a more prominent effect for target frequencies. Later, significantly greater responses were observed to the target frequency ranges compared to nontarget ranges. The two target frequency ranges in animal one overlapped and were considered together for this statistical test.

To test the stability of these recording measures across weeks, in animal two, physiological data were collected for five sessions immediately before behavioral training. The plot in Fig. 3 shows

Table 1. t test results for Fig. 3 t statistics and associated probabilities for significant changes are shown

Week	Animal one		Animal two			
	t-stat	P	t-stat-1	P-1	t-stat-2	P-2
0	—	—	—	—	—	—
1	—	—	4.53	0.00002	—	—
2	2.756	0.005	10.8	0	5.45	3e-7
3	—	—	5.01	2e-6	—	—
4	2.666	0.008	6.77	8e-10	—	—
5	5.03	5e-7	10.8	0	3.14	0.002
6	6.44	1e-10	—	—	—	—

t-stat-1 is the statistic for the upper target frequency range for animal two, and t-stat-2 is the statistic for the lower target frequency range.




Fig. 4. Daily changes across the breakpoint in animal two for the upper frequency range. Circles show average response to the target frequency range. Xs show the daily response to the standards. Stars show the responses to frequencies not used in the behavior. All physiological responses are derived immediately before the corresponding behavioral sessions. The animal first achieved performance above chance before day 0 recordings.

that all three frequency ranges were similar in responsiveness to the first week of the behavior.

In addition to comparing target and control frequency ranges, responses to the standards were compared to the target range responses. The target range was always just higher in frequency than the standard tones. A sign test was used to determine whether the responses to the standards were significantly lower than the 12 adjacent higher semitones for all data pooled weekly. In animal one, significant differences for both standards compared to their next higher 12 semitones (each $P < 0.01$, or at least 10 of 12 higher in response) occurred only in the last 2 weeks of the behavior. In animal two, responses to standards were each lower than the next 12 higher semitones on weeks 1, 4, and 5 (each $P < 0.01$). The plots of relative responsiveness showing the standard and target frequencies are shown in Fig. 2B. The thin vertical lines indicate the standard frequencies.

To investigate the behavioral correlate of the breakpoint seen in Fig. 3, changes in measured thresholds were compared to the changes in responsiveness of A1 to the target and nontarget frequencies. In Fig. 3 *Upper*, the threshold changes as a function of time are shown. For both animals, decreases in threshold are seen across the breakpoint week. In animal one, the threshold changes from 10 semitones during the week before the breakpoint to 6 semitones during the breakpoint week to 1 semitone the following week. In animal two, the breakpoint week is the first week the animal had measurable thresholds; for earlier weeks, all behavior was at chance for all targets.

The breakpoint changes are illustrated on a day-by-day basis in Fig. 4 for animal two. Firing rates begin to rise during the first recording session after the animal performed the behavior at better than chance discrimination. The firing rates observed in the target frequency ranges more than doubled after the breakpoint compared to before the breakpoint. Responses to tones with target range frequencies were consistently larger than responses to either standards or nontarget frequencies after the breakpoint. A similar phenomenology, elevated responses to all frequencies, and particularly to target sounds, was observed at the lower frequency range and in animal one shortly after the animals came under behavioral control.

Discussion

Two animals were engaged in frequency discrimination training. Both improved their frequency discrimination thresholds. Both animals had a psychophysical change in which thresholds approached asymptotic limits for their species. Physiological correlates of this change were elevated responses to all frequencies

and differentially intensified responses to target stimuli relative to other frequencies. In particular, responses to target tones were elevated relative to a standard frequency, even though identification of both frequencies was necessary for task performance. The emergence of strengthened responses to stimuli preceding a rewarded response—correct target identification—relative to those preceding a time-out is consistent with A1 providing a neural substrate for Thorndike's law of effect (1) in this behavior.

This emergent differential responsiveness establishes a role for the sensory cortex, at a level as low as the primary sensory cortex, in changing stimulus–response relationships to enhance the future probability of rewards. Both standard frequencies and targets were used in the task, and discriminating between them required an internal representation of both frequencies. However, the animal held its position until the target occurred. Once the target occurred, the signal had to be communicated to the frontal cortex to initiate a motor response (15).[†] One interpretation of these changes would be that sensory contexts that are triggers for rewarded motor responses are enhanced. In a broader view of the sensory cortex, other findings can be reinterpreted. Somatosensory maps of area 3b in primates have restricted hairy skin representations; animals rarely initiate motor acts on the basis of stimulation of the hairy skin (16). If monkeys are trained to retrieve small pellets from long narrow cylinders, the digit tips used in the task become represented with smaller receptive fields and large representations (7). These distal fingertips receive sensory input that guides the fingers in retrieving food, a basic reinforcer. In another study, stimuli that were delivered to adjacent digits simultaneously before a rewarded motor response were represented together in area 3b (8). In the motor neuropathology of focal dystonia, there are parallel abnormalities in motor capabilities and in the hand representation in area 3b in nonhuman primates and humans (17–20). These findings form reasons to rethink the functional roles of the sensory cortex to include remapping sensory contexts to motor acts that enhance the future probability of satisfying events. Further, plasticity in the primary sensory cortex may be specifically sensitive to nearly simultaneous stimulation of sensory epithelia prior to motor acts that lead to reward.

The differences in responsiveness to standard and target frequencies emerged as an apparent consequence of the animal's responses to the targets predicting a reward and its responses to the standards predicting a time-out. This difference in behavioral context suggests that neuromodulator release played a role in differentiating the target from the standard. In particular, the first presentation of the target tone, which predicted that a motor act would lead to reward, should be followed by activity of ventral tegmental area neurons that release dopamine to the cerebral cortex (21). Activity of cholinergic neurons in the nucleus basalis might also be expected to follow such a cue (22, 23). Either of these neuromodulators, if released after a sound cue, causes increases in the representation of that cue in A1 of the rat (24–26). Although the neuromodulatory release after the target tone would be sufficient to cause reorganization, myriad other cues are present in the actual behavior consisting of the activity of all other neurons that project to auditory pathways up to and including A1. Changes in any of these projections should lead to representational changes in A1. The neuromodulatory cues are the most specific to the behavioral context, which makes them attractive candidates as regulators of the target and standard response enhancement and suppression.

The physiological changes at the breakpoint are a correlate of the animal reacting more specifically to the presentation of targets. At that point, the thresholds of the animals began to drop markedly and approached asymptotic limits for their species (6).

[†]Woolsey, C. N. & Bard, P. (1936) *Am. J. Physiol.* 116, 165 (abstr.).

At the same time, the population of neurons in A1 responded more strongly to target than to nontarget stimuli or standards. Daily analysis of firing rates in both animals demonstrated that physiological changes were observed only *after* changes in the behavior. The elevated responsiveness and differentiation of frequency ranges cannot occur without shifts in responsiveness of individual neurons. Although changes in responsiveness or selectivity may be expected if electrodes move or electrode impedances change, there is no reason to suspect that the highly statistically significant effect—an emergent differential responsiveness between targets and to nontargets—would be attributed to such changes. To create a stronger response to the target range, either single neurons must shift in their selectivity, or increases in responsiveness must be differential between neurons representing the target and nontarget ranges, or both. The contributions of these different single neuron changes that would lead to differential selectivity cannot be determined with high confidence, although examples of each of these effects were present in the data.

Although physiological changes were observed for all standards relative to the next 12 higher semitones, the changes in responsivity seem different between the higher and lower standards in both animals. Responsiveness at frequencies just higher than 880 Hz in animal one and just higher than 1,490 Hz in animal two is not as strong as responsiveness just above the higher targets of 1,396 and 7,040 Hz in both animals. This difference in effect may occur because the higher-frequency target ranges are always higher than both standards. In animal two, the lower target range was not elevated in responsiveness compared to the nontarget range shown in Fig. 3. This was the only target/nontarget comparison in this study in which the target range was lower in frequency than the nontarget range. The animal could have attended to that range as potential targets in trials by using the lower standard. Unfortunately, the implant did not offer complete enough sampling at even lower frequencies to make a more consistent comparison possible. The lower frequency range in animal two was also the only frequency range on the edge of the implant, which raises the possibility that a substantial A1 sector that was relevant to the behavior was unsampled.

In both animals, changes in responsiveness of A1 were also associated with two time points: the beginning of behavior and the breakpoint. These points were also times at which there were substantial behavioral changes. At the beginning of the task, the animal went through a transition from head-positioning behavior without auditory stimulation to head-positioning behavior with auditory stimulation. The physiological breakpoint occurred when the animal began to behave, i.e., began to improve markedly at the task. The association of behavioral changes with the increased responsiveness of A1 neurons suggests that the neurons of the locus coeruleus, which has responses that relate to changes in behavioral contingencies (27), may be involved in this effect. There is some question whether the responses in the first week were above or below normal. In animal two, physiological data for the 5 days before the behavior began were not significantly different from the first 5 days of the behavior. This

lack of change suggests that the difference in responsiveness between the first 2 weeks of behavior is likely due to suppression of responses in the second week and not to enhanced responses in the first.

The discrimination thresholds of our animals were within the normal range found for training primates (6, 28–30). Our animals had discrimination thresholds, ($\Delta F/F$), of 0.06 and 0.03. In a previous study, owl monkeys were also trained in frequency discrimination and the changes in A1 were evaluated (6). In that study, the standards were also fixed, and the targets were variable and usually higher than the standards [but see their figure 8 (6)]. The physiological changes in that study were evaluated with dense microelectrode mapping experiments in which A1 was sampled with 20–25 recordings per mm². The different physiological techniques used allowed those investigators to find significant increases in the area of cortex that responded to the frequency ranges around the trained frequency. Those experiments also demonstrated increases in the proportion of tuning curves for which the neurons exhibited high selectivity 10 dB above the threshold. This property is thought to be a characteristic of the auditory inputs that is reflected in tuning of neurons at higher levels of auditory processing (31). The implant technique cannot confirm changes in distributional area with statistical significance, because the implant allowed the sampling of 10–15 sites continuously throughout the behavior instead of 150 sites at the behavioral endpoint. Instead, the present study adds to the phenomenology of the behavior by finding the behavioral and physiological breakpoint and by demonstrating a simple neural code that may be the basis of the behavior. Further analysis of this work may focus on responses recorded during the behavior and on more direct testing of this putative neural code.

Conclusion

Structured operant training creates a behavioral context in which behavioral reinforcement cues are timed relative to sensory events. This relationship causes changes in the relative sensitivity of primary auditory cortex. In adult primates, large-scale remodeling of A1 occurred: the representations of target frequencies became stronger than task standards and stronger than the representations of nontask frequency ranges. These changes occurred within 8 weeks, although a closer analysis of the time course suggests that most of the change occurred within 2 weeks, once the animals began to perform the task with progressively lower thresholds. Physiological changes were temporally specific; only responses to sounds preceding a motor act that led to reward were enhanced. Those preceding an act that led to a time-out were comparatively weaker. Such a first-order learning rule, an embodiment of Thorndike's law of effect in A1, can powerfully alter behavior to enhance the probability of future rewards by pairing behavioral contexts with motor acts.

We thank David Moorman and Lucas Zier for participating in data collection for these experiments. Gregg Recanzone made useful commentary in his review of the paper. This work supported by the Coleman Fund, the Hearing Research Institute, the Sooy Fund, and National Institutes of Health Grants 1F32NS10154 and NS10414.

- Domjan, M. (1997) *The Principles of Learning and Behavior* (Wadsworth, Belmont, CA), 4th Ed.
- deCharms, R. C., Blake, D. T. & Merzenich, M. M. (1999) *J. Neurosci. Methods* **93**, 27–35.
- Jenkins, W. M., Merzenich, M. M., Ochs, M. T., Allard, T. & Guic-Robles, E. (1990) *J. Neurophysiol.* **63**, 82–104.
- Recanzone, G. H., Merzenich, M. M., Jenkins, W. M., Grajski, K. A. & Dinse, H. R. (1992) *J. Neurophysiol.* **67**, 1031–1056.
- Recanzone, G. H., Merzenich, M. M. & Schreiner, C. E. (1992) *J. Neurophysiol.* **67**, 1071–1091.
- Recanzone, G. H., Schreiner, C. E. & Merzenich, M. M. (1993) *J. Neurosci.* **13**, 87–103.
- Xerri, C., Coq, J. O., Merzenich, M. M. & Jenkins, W. M. (1996) *J. de Physiologie* **90**, 277–287.
- Wang, X., Merzenich, M. M., Beitel, R. & Schreiner, C. E. (1995) *J. Neurophysiol.* **74**, 2685–2706.
- Recanzone, G. H., Jenkins, W. M., Hradek, G. T. & Merzenich, M. M. (1992) *J. Neurophysiol.* **67**, 1015–1030.
- McDonough, R. N. & Whalen, A. D. (1995) *Detection of Signals in Noise* (Academic, San Diego).
- Imig, T. J., Ruggero, M. A., Kitzes, L. M., Javel, E. & Brugge, J. F. (1977) *J. Comp. Neurol.* **171**, 111–128.
- Recanzone, G. H., Schreiner, C. E., Sutter, M. L., Beitel, R. E. & Merzenich, M. M. (1999) *J. Comp. Neurol.* **415**, 460–481.

13. Galambos, R. & Davis, H. (1943) *J. Neurophysiol.* **6**, 39–58.
14. Hubel, D. H. (1957) *Science* **125**, 549–550.
15. Bard, P. (1937/38) *Harvey Lectures* (Academic, New York), pp. 143–169.
16. Merzenich, M. M., Kaas, J. H., Sur, M. & Lin, C. S. (1978) *J. Comp. Neurol.* **181**, 41–73.
17. Byl, N. N., Merzenich, M. M. & Jenkins, W. M. (1996) *Neurology* **47**, 508–520.
18. Byl, N. N., Merzenich, M. M., Cheung, S., Bedenbaugh, P., Nagarajan, S. S. & Jenkins, W. M. (1997) *Phys. Ther.* **77**, 269–284.
19. Bara-Jimenez, W., Catalan, M. J., Hallett, M. & Gerloff, C. (1998) *Ann. Neurol.* **44**, 828–831.
20. Elbert, T., Candia, V., Altenmuller, E., Rau, H., Sterr, A., Rockstroh, B., Pantev, C. & Taub, E. (1998) *NeuroReport* **9**, 3571–3575.
21. Schultz, W., Dayan, P. & Montague, P. R. (1997) *Science* **275**, 1593–1599.
22. Richardson, R. T. & DeLong, M. R. (1991) *Adv. Exp. Med. Biol.* **295**, 233–252.
23. Richardson, R. T. & DeLong, M. R. (1990) *J. Neurosci.* **10**, 2528–2540.
24. Kilgard, M. P. & Merzenich, M. M. (1998) *Science* **279**, 1714–1718.
25. Bao, S., Chan, V. T. & Merzenich, M. M. (2001) *Nature (London)* **412**, 79–83.
26. Kisley, M. A. & Gerstein, G. L. (2001) *Eur. J. Neurosci.* **13**, 1993–2003.
27. Aston-Jones, G., Rajkowski, J. & Kubiak, P. (1997) *Neuroscience* **80**, 697–715.
28. Prosen, C. A., Moody, D. B., Sommers, M. S. & Stebbins, W. C. (1990) *J. Acoust. Soc. Am.* **88**, 2152–2158.
29. Sinnott, J. M., Petersen, M. R. & Hopp, S. L. (1985) *J. Acoust. Soc. Am.* **78**, 1977–1985.
30. Fay, R. R. (1974) *J. Acoust. Soc. Am.* **56**, 206–209.
31. Ramachandran, R., Davis, K. A. & May, B. J. (1999) *J. Neurophysiol.* **82**, 152–163.