

SOUND-DISCRIMINATION LEARNING AND AUDITORY DISPLAYS

Beverly A. Wright and Matthew B. Fitzgerald

Department of Communication Sciences and Disorders
and Institute for Neuroscience
Northwestern University
2240 Campus Drive
Evanston, IL 60208
b-wright@northwestern.edu

ABSTRACT

Human listeners can learn to discriminate between sounds that are initially indistinguishable. To better understand the nature of this learning, we have been using behavioral techniques to examine training-induced improvements on basic auditory discrimination tasks. Here we report how multiple-hour training differentially affects the discrimination of sound frequency, intensity, location, and duration, and how learning on a given discrimination condition generalizes, or fails to generalize, to stimuli not encountered during training. We discuss how these data contribute to our understanding of discrimination learning and of the mechanisms underlying performance on particular trained tasks, and explore the implications of this learning for the design and evaluation of auditory displays.

1. INTRODUCTION

Relatively little is known about how practice influences the performance of human adults on basic auditory discrimination tasks. We are interested in this issue because auditory learning provides a window into the mechanisms underlying performance on the trained task, and into the learning process itself [1]. Among other benefits, a greater understanding of these issues will help guide the search for the physiological substrates of learning, and aid the development of perceptual training schemes.

With these motivations, we examined learning using simple, pure-tone stimuli, on five basic auditory discrimination tasks: frequency, intensity, interaural-time-difference (ITD), interaural-level-difference (ILD), and duration [2]. Our overarching questions were (1) can listeners improve their ability to discriminate stimuli along each of these dimensions with practice, and if so, (2) does this learning generalize to the trained discrimination performed with untrained stimuli? Because we used the same basic format for all five experiments, any differences in the learning patterns across these trained discriminations likely reflect differences in the plasticity of the underlying mechanisms. Here we report the influence of training on discrimination thresholds assessed by comparing the mean proportional improvements on trained and untrained conditions between listeners who were, and those who were not, given multiple-hour practice on a single discrimination condition. Learning on these five tasks followed one of two general patterns. For ITD and intensity discrimination, additional practice did not

lead to greater learning than that seen in untrained listeners. In contrast, for ILD, duration, and frequency discrimination, such practice yielded greater learning, but only on a subset of conditions. The implications of these learning patterns for the design and evaluation of auditory displays are discussed.

2. GENERAL METHOD

The format was the same for each of the five experiments [1, 2]. At the beginning of each experiment, we gave a group of naïve listeners a pretest during which we measured their discrimination thresholds for tonal stimuli on six conditions. We collected five threshold estimates per condition (300 trials) over the course of a single ~2.5 hour session. We then divided the listeners into two groups. One group, referred to as trained listeners, received training on one of the conditions from the pretest. This training consisted of 12-15 threshold measurements per day (720-900 trials, ~1 hour), for 6 to 10 days. The other group, referred to as control listeners, received no training. Finally, at the end of the training phase, we retested all listeners on a posttest that employed the same conditions as the pretest. We randomized the condition order in the pre- and posttests across listeners, but each listener received the conditions in the same order for both of these tests. All stimuli were digitally generated and presented over headphones.

We measured the discrimination thresholds using an adaptive, two-interval, forced-choice procedure. We adjusted the signal level within each 60-trial block using the three-down/one-up rule to estimate the 79% correct point on the psychometric function [3]. All listeners received visual feedback as to whether each of their responses was right or wrong throughout the entire experiment.

Here we assessed learning based only on the mean changes in threshold values, computed as the proportional improvement for each listener: $[(\text{pretest threshold} - \text{posttest threshold})/\text{pretest threshold}]$. The advantage of this calculation is that it normalizes data for which starting values vary across listeners and/or conditions. Note that this measure emphasizes the amount of change relative to the starting value, rather than the absolute magnitude of threshold change.

We analyzed the proportional-improvement scores for each experiment using a 2 group (trained vs. control) \times n condition analysis of variance, in which each included condition ($n=3$ to 5) employed the same discrimination task with a different standard stimulus (described for each experiment, below). We did not use

repeated measures on condition in these analyses because the number of listeners sometimes differed across conditions. When there was a significant group x condition interaction, we used *t*-tests to compare the scores of the trained and control listeners separately for each condition. If the proportional-improvement score of the trained listeners was greater than that of controls on the trained condition, we concluded that the trained listeners learned on the trained condition during the training phase. If that score was greater for the trained than control listeners on an untrained condition, we concluded that the trained listeners generalized their training-phase learning to that condition. We used an alpha value of 0.05 for all analyses.

3. RESULTS

We observed two different general learning patterns across the five discrimination conditions. In the first pattern, listeners who were trained for 6-10 hours showed no more improvement than controls who participated only in the 2.5-hour pre- and post-tests. This pattern occurred for learning on ITD and intensity discrimination, based on our mean data. In the second pattern, listeners who received multiple-hour training improved more than controls on the condition on which they were trained, and generalized their learning to untrained stimuli that differed from the trained condition along some dimensions, but not to others. This pattern occurred for ILD, duration, and frequency discrimination.

3.1. Trained learning equal to control learning

Trained listeners learned no more than controls on both ITD and intensity discrimination. In the ITD-training experiment (Fig. 1; [4]), trained listeners practiced discriminating the lateral position of a standard stimulus of 300-ms, 0.5-kHz tones presented to both ears at 70 dB SPL with an ITD of 0 μ s from a signal stimulus that differed from the standard only in that the ITD favored the right ear. In the untrained conditions, the standard differed from that in the trained condition either only in location (150- μ s ITD vs. 0- μ s ITD), frequency (1 kHz vs. 0.5 kHz), or interaural cue (ILD vs. ITD), or in both the frequency and cue (4 kHz, ILD vs. 0.5 kHz, ITD). The mean proportional-improvement scores did not differ between the trained and control listeners (main effect for group, $p=0.095$; group x condition interaction, $p=0.239$). However, listeners did learn, because, on each condition, the combined scores of both groups were always greater than zero (based on 95%-confidence intervals).




Figure 1. Threshold proportional-improvement scores for ITD discrimination for the trained (left-most) and untrained conditions. Results are shown for the trained (black bars) and control (white bars) listeners. Error bars indicate one standard error of the mean. (Data from [4]).

In the intensity-training experiment (Fig. 2; [5]), the trained condition employed a standard of two 15-ms, 1-kHz, tone pips whose onsets were separated by 100 ms, presented to both ears with an ITD of 0 μ s at 76 dB SPL. The listener's task was to discriminate this standard from a signal of greater intensity. We chose this particular standard because we previously had used a very similar one to train both duration and frequency discrimination (see below). Here, the untrained conditions differed from the trained one only in the standard stimulus level (46 or 91 dB SPL vs. 76 dB SPL), frequency (4 kHz vs. 1 kHz), duration (temporal interval between tone pips, 50 ms vs. 100 ms) or location (200 μ s vs. 0 μ s). As for ITD discrimination, the mean proportional-improvement scores did not differ between the trained and control listeners (main effect for group, $p=0.684$; group x condition interaction, $p=0.694$), but listeners did learn, because, their scores combined across groups were greater than zero on each condition (based on 95%-confidence intervals). Thus, for both ITD and intensity discrimination, 6-10 hours of training yielded no greater reduction in threshold than did exposure to only the 2.5-hour pre- and posttests. For each discrimination task, this learning was relatively uniform for all stimuli.




Figure 2. Same as Fig. 1, but for intensity discrimination [5].

3.2. Trained learning greater than control learning

In contrast to the learning patterns for ITD and intensity discrimination, trained listeners learned more than controls on ILD, duration, and frequency discrimination and generalized their learning only to task-dependent subsets of untrained conditions. In the ILD-training experiment (Fig. 3; [4]) the standard in the trained condition consisted of 300-ms, 4-kHz tones presented to both ears at 70 dB SPL, and the signal differed from the standard only in that the ILD favored the right ear. The untrained conditions differed from the trained one only in the standard location (6-dB ILD vs. 0-dB ILD) or frequency (0.5 kHz or 6 kHz vs. 4 kHz), or in both the frequency and cue (0.5 kHz, ITD vs. 4 kHz, ILD). Here, the control listeners improved on all conditions (based on 95%-confidence intervals). However, the trained listeners learned more than controls (group x condition interaction, $p=0.008$), but only on the trained condition ($p<0.001$) and the untrained location (6-dB ILD; $p=0.036$) and not on the untrained frequencies or the ITD cue.




Figure 3. Same as Fig.1, but for ILD discrimination (Data from [4]).

We examined learning on duration discrimination (Fig. 4; [6, 7]) using a monaural presentation of essentially the same standard as in the intensity-training experiment (described above). The listener's task was to discriminate the standard from a signal in which the two tone pips were separated by a longer temporal interval. The untrained conditions differed from the trained one only in the standard frequency (4 kHz vs. 1 kHz) or duration (200 or 500 ms vs. 100 ms). In this case, the controls did not improve on any condition (based on 95%-confidence intervals). The trained listeners learned significantly more than controls (group x condition interaction, $p=0.026$) on the trained condition ($p=0.002$) and the untrained frequency (4 kHz, $p<0.001$), but not on the untrained durations.

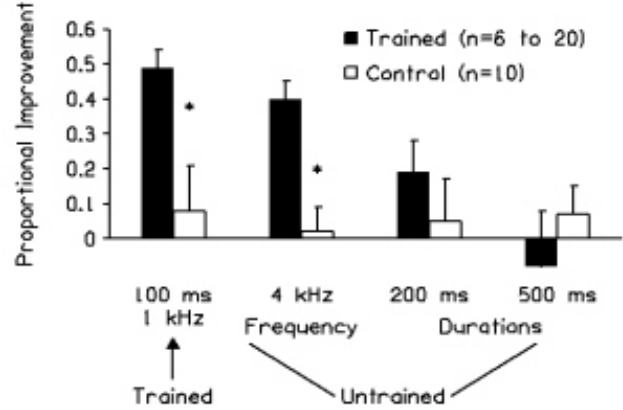


Figure 4. Same as Fig.1, but for duration discrimination (Data from [6,7])

Finally, to examine learning on frequency discrimination (Fig. 5; [7]), we trained listeners with the same standard as in the duration-training experiment but required them to discriminate that standard from a signal of lower frequency. In the two untrained frequency-discrimination conditions, the standard differed from that in the trained condition only in its duration (temporal interval between tone pips, 50 ms vs. 100 ms) or frequency (4 kHz vs. 1 kHz). Here, controls improved for the 4-kHz, but not the two 1-kHz stimuli (based on 95%-confidence intervals). The trained listeners improved more than controls (group x condition interaction, $p=0.024$), but, surprisingly, not on the trained condition ($p=0.129$). However, there appears to have been training-induced learning, because the trained listeners learned more than controls on the untrained duration ($p=0.01$). This learning did not generalize to the untrained frequency.

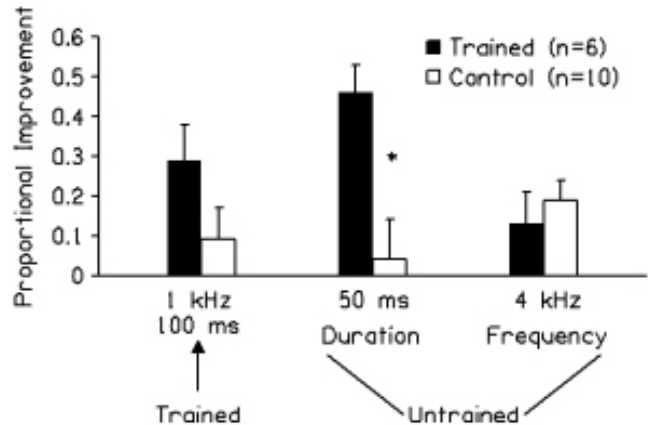


Figure 5. Same as Fig. 1, but for frequency discrimination [7]

Thus, for ILD, duration, and frequency discrimination, 9-10 hours of training yielded greater threshold improvements than did exposure to only the 2.5-hour pre- and posttests. Further, this training-induced learning generalized to some untrained stimuli but not others, and did so in a pattern unique to each trained discrimination task. For ILD discrimination, learning generalized to an untrained location, but was specific to the trained frequency

and cue. For duration discrimination, learning generalized to an untrained frequency, but was specific to the trained duration. Finally, for frequency discrimination, learning generalized to an untrained duration, but was specific to the trained frequency.

4. DISCUSSION

4.1. Two Types of Discrimination Learning

From the present results, it appears that improvements on basic auditory discrimination tasks result from two different types of learning, that this learning affects task-specific processing mechanisms, and that these mechanisms are not equally malleable. The threshold improvements on ITD and intensity discrimination may reflect primarily conceptual or procedural learning. In both instances, control listeners improved on all conditions with exposure only to the pre- and posttests. We, and others, have proposed that such rapid and general learning reflects the acquisition of the general procedures needed to perform the task, and as such does not result from fundamental changes in stimulus processing [4, 8, 9]. By this account, learning on ITD and intensity discrimination is primarily procedural, because additional training did not benefit listeners on these tasks. This lack of training-phase learning may have occurred either because these discriminations are already over-learned, or because the mechanisms that govern them are relatively inflexible.

In contrast, the threshold improvements in ILD, duration, and frequency discrimination may largely result from perceptual or stimulus learning. In these cases, control listeners improved on some tasks, suggesting procedural learning. However, additional training always resulted in greater improvements on a subset of stimuli. Such slow and stimulus-specific learning may reflect fundamental changes in stimulus processing [4, 8, 9]. By this view, the different generalization patterns for the different tasks indicate both that training affected separate mechanisms for the different discriminations, and that these mechanisms are organized in different ways. Training modified a mechanism that, (1) for ILD discrimination, processes multiple locations, but only the ILD cue with the trained frequency, (2) for duration discrimination, processes multiple frequencies, but only the trained duration, and (3) for frequency discrimination, processes multiple durations, but only the trained frequency. Others have reported results for auditory duration [10] and frequency [11, 12, 13] discrimination that are consistent with this interpretation. Further, training-induced learning on duration discrimination has also been shown to be specific to the trained duration in the somatosensory system [14], and to generalize from the somatosensory to auditory [14], and from the auditory to motor [15] systems, suggesting that training on duration discrimination affects a multi-system timing mechanism that processes different durations separately.

4.2. Discrimination Learning and Auditory Displays

For a designer of auditory displays, the primary take-home messages here are that practice can: (1) improve a listener's ability to discriminate small variations along a variety of acoustic dimensions, (2) affect the discriminability of some dimensions

(ILD, frequency, duration) more than others (ITD, intensity), and (3) enhance the perception along the trained dimension for some, though not all, stimuli not encountered during training, even when those stimuli are not auditory (see previous paragraph).

A given display is most likely to induce such discrimination learning along a particular dimension if, like in the present experiments, it requires the listener to actively distinguish stimuli along that dimension, and gives examples of, and feedback for, that discrimination. Active discrimination of stimuli along the dimension of interest appears to be important because listeners trained to discriminate stimuli along one acoustic dimension rarely generalize their learning to discriminations along untrained dimensions. For example, learning on a duration-discrimination task yielded no improvement on a frequency-discrimination task that employed the same standard stimulus [7]. It also appears that improvements at threshold may be aided by an initial exposure to clearly discriminable stimuli, which presumably instructs the listener of the discrimination task to be performed [16]. Finally, improvements in discrimination may also be facilitated by receiving some form of feedback about performance. Supporting this idea, cortical changes in response to passively perceived sounds only occurred in rats when stimuli were paired with electrical stimulation of a brain structure that indicates the behavioral relevance of stimuli [17].

Assuming these learning circumstances are provided, it appears that a display may induce learning regardless of the initial discriminability of the stimuli it employs. If a display includes many stimuli that are initially difficult to discriminate, as in the present experiments, we can infer that use of the display would yield learning and generalization patterns similar to those reported here. This situation may occur in more displays than might be assumed initially, because we have observed that the discrimination thresholds of naïve listeners are frequently 2 to 4 times higher than the thresholds reported in the literature. Further, improvements in the ability to distinguish stimuli that initially are barely discriminable, like here, appear to enhance the saliency of differences between stimuli that can already be distinguished. Supporting this idea, listeners who reduced their discrimination thresholds through training were subsequently more accurate at labeling individual suprathreshold stimuli, than untrained listeners [18]. We do not know whether listeners would show the same patterns if presented with a display that only contains stimuli that are already relatively easy to discriminate. However, providing indirect evidence that this might be the case, listeners trained to attach one of two labels to individual stimuli--the majority of which could be easily labeled--subsequently showed the same learning and generalization for stimuli that were initially indistinguishable as did listeners trained to discriminate those near-threshold stimuli [19].

5. SUMMARY AND CONCLUSIONS

In summary, a systematic examination of learning on five basic auditory discrimination tasks, assessed by threshold reductions, revealed two different learning patterns. Improvements on ITD and intensity discrimination were rapid and general, and therefore appeared to reflect procedural learning. Improvements on ILD, duration, and frequency discrimination were slow and stimulus specific, and, as such, appeared to arise from stimulus learning. Overall, the present results indicate that learning affected a

different mechanism for each trained task, and that these mechanisms differ in their plasticity as well as their organization.

While much remains to be discovered about the forces that drive auditory discrimination learning, it is clear that learning could affect the perception of many auditory displays. In most cases, this learning would simply increase the distinctiveness of the various stimuli in the display. These potential increases in saliency should be kept in mind when initially evaluating a display, because the perception of the display may differ between naïve and experienced listeners. Finally, discrimination learning could also be exploited by specifically designing displays to induce it, for example, to help treat clinical disorders [20].

6. ACKNOWLEDGMENTS

Substantial portions of this paper are to be published in "Learning and generalization on five basic auditory discrimination tasks as assessed by threshold changes," in *Auditory Signal Processing--Physiology, Psychoacoustics, and Modeling (Proceedings of the 13th International Symposium on Hearing)*, edited by D. Pressnitzer, A. de Cheveigne, S. McAdams, and L. Collet, Springer-Verlag, New York, copyright 2004, and are reprinted with permission from the publisher. This work was supported by NIH/NIDCD.

7. REFERENCES

- [1] B.A. Wright, "Why and how we study human learning on basic auditory tasks," *Audiol. Neurootol.* vol. 6, pp. 207-210, 2001.
- [2] B.A. Wright and M.B. Fitzgerald, "Learning and generalization on five basic auditory discrimination tasks as assessed by threshold changes," in *Auditory Signal Processing--Physiology, Psychoacoustics, and Modeling (Proceedings of the 13th International Symposium on Hearing)*, edited by D. Pressnitzer, A. de Cheveigne, S. McAdams, and L. Collet, Springer-Verlag, New York, in press.
- [3] H Levitt, "Transformed up-down methods in psychoacoustics," *J. Acoust. Soc. Am.*, vol. 49, pp. 467-477.
- [4] B.A. Wright, and M.B. Fitzgerald, "Different patterns of human discrimination learning for two interaural cues to sound-source location," *Proceed. Nat. Acad. Sciences*, vol. 98, pp. 12307-12312, 2001.
- [5] D.A. Abrams and B.A. Wright (unpublished data).
- [6] B.A. Wright, D.V. Buonomano, H.W. Mahncke, M.M. Merzenich, "Learning and generalization of auditory temporal-interval discrimination in humans," *J. Neurosci.*, vol. 17, pp. 3956-3963, 1997.
- [7] B.A. Wright "Generalization of auditory-discrimination learning," *Assoc. Res. Otolaryngol. Abs.*, Abs. 413, vol. 21, pp. 104, 1998.
- [8] G.H. Recanzone, C.E. Schreiner, and M.M. Merzenich, "Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys," *J. Neurosci.* vol. 13, pp. 87-103, 1993.
- [9] K. Robinson and A.Q. Summerfield, "Adult auditory learning and training," *Ear. Hear.* vol. 17, pp. 51-65, 1996.
- [10] U.R. Karmarkar and D.V. Buonomano, "Temporal specificity of perceptual learning in an auditory discrimination task," *Learning and Memory*, vol. 10, pp. 141-147, 2003.
- [11] K. Delhommeau, C. Micheyl, R. Jouvent, and L. Collet, "Transfer of learning across duration and ears in auditory frequency discrimination," *Perception and Psychophysics*, vol. 64, pp. 426-436, 2002.
- [12] L. Demany and C. Semal, "Learning to perceive pitch differences," *J. Acoust. Soc. Am.*, vol. 111, pp. 1377-1388, 2002.
- [13] D.R.F. Irvine, R.L. Martin, E. Klimkeit, and R. Smith, "Specificity of perceptual learning in a frequency-discrimination task," *J. Acoust. Soc. Am.*, vol. 108, pp. 2964-2968, 2000.
- [14] S.S. Nagarajan, D.T. Blake, B.A. Wright, N. Byl, and M.M. Merzenich, "Practice-related improvements in somatosensory interval discrimination are temporally specific but generalize across skin location, hemisphere, and modality." *J. Neurosci.*, vol. 18, pp. 1559-1570, 1998.
- [15] D.V. Meegan, R.N. Aslin, and R.A. Jacobs, "Motor timing learned without motor training," *Nature Neuroscience* vol. 3, pp. 860-862, 2000.
- [16] M. Ahissar, and S. Hochstein, "Task difficulty and specificity of perceptual learning," *Nature* vol. 387, pp. 401-406, 1997.
- [17] M.P. Kilgard, and M.M. Merzenich, "Cortical map reorganization enabled by nucleus basalis activity," *Science* vol. 279, pp. 1714-1718, 1998.
- [18] J.A. Mossbridge, B.N. Scissors, and B.A. Wright, "Improved representations of individual stimuli induced by a two-interval discrimination procedure," *Assoc. Res. Otolaryngol. Abs.*, Abs. 880, vol. 26, pp. 222, 2003.
- [19] J.A. Mossbridge, and B.A. Wright, "Similar improvements in auditory discrimination produced by multiple-hour training with a labeling or discrimination procedure," *J. Acoust. Soc. Am.*, vol. 111, pp. 2358, 2002.
- [20] M.M. Merzenich, W.M. Jenkins, P. Johnston, C. Schreiner, S.L. Miller, P. Tallal, "Temporal processing deficits of language-learning impaired children ameliorated by training," *Science*, vol. 271, pp. 77-81, 1996.