Adapting to supernormal auditory localization cues. II. Constraints on adaptation of mean response

Barbara G. Shinn-Cunningham,[a] Nathaniel I. Durlach, and Richard M. Held
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

(Received 23 September 1996; revised 27 January 1998; accepted 2 March 1998)

A series of experiments was performed in which subjects were trained to interpret auditory localization cues arising from locations different from their normal spatial positions. The exact pattern of mean response to these alterations (as a function of time) was examined in order to begin to develop a quantitative model of adaptation. Mean responses were roughly proportional to the normal position associated with the localization cues presented. As subjects adapted, the best-fit slope (relating mean response and normal position) changed roughly exponentially with time. The exponential rate and adaptation asymptote were found for each subject in each experiment, as well as the rate and asymptote of readaptation to normal cues. The rate of adaptation does not show any statistical dependence on experimental conditions; however, the asymptote of the best-fit slope varied with the strength of the transformation used in each experiment. This result is consistent with the hypothesis that subjects cannot adapt to a nonlinear transformation of auditory localization cues, but instead adapt to a linear approximation of the transformation. Over time, performance changes exponentially towards the best-fit linear approximation for the transformation used in a particular experiment, and the rate of this adaptation does not depend upon the transformation employed. © *1998 Acoustical Society of America.* [S0001-4966(98)03406-7]

PACS numbers: 43.71.Hw, 43.72.Ew, 43.66.Qp [RHD]

INTRODUCTION

A series of experiments described in detail in Shinn-Cunningham *et al.* (1998) examined how bias and resolution are affected when subjects are trained with ''supernormal'' localization cues. While the rationale behind the experiments, the experimental protocol, and the experimental conditions are described in detail in other papers (Durlach *et al.*, 1993; Shinn-Cunningham *et al.*, 1998), earlier results are summarized here for convenience. In the experiments, subjects were presented with auditory localization cues simulated over headphones using head-related transfer functions [HRTFs; for a review of these techniques, see Wenzel (1992)]. The goal of the work was to determine if subjects could adapt to learn a new correspondence between the localization cues of the physical HRTFs and the reported azimuthal position. To this end, subjects trained to identify the azimuthal location of an auditory source whose physical cues normally correspond to a different source position.

Normally, a source at azimuth θ and elevation ϕ is simulated by using the HRTF for that position. In these experiments, the HRTF used to simulate a source at position $[\theta,\phi]$, was equal to the HRTF normally corresponding to position $[f_n(\theta),\phi]$, where $f_n(\theta)$ is given by

$$f_n(\theta)=\frac{1}{2}\tan^{-1}\left[\frac{2n\sin(2\theta)}{1-n^2+(1+n^2)\cos(2\theta)}\right]. \qquad (1)$$

The parameter n corresponds to the slope of the transformation at $\theta=0$. ''Normal'' localization cues are presented when $n=1$ [i.e., the function $f_n(\theta)$ is a straight line of slope one through the origin]. This mapping is shown in Fig. 1 for the values of n used in the experiments.

In order to determine whether subjects could adapt to the remapping of HRTF cues, they were repeatedly tested over the course of experimental sessions lasting roughly 2 h, first using the ''normal'' mapping ($n=1$) and then an altered mapping ($n>1$). At the end of the experimental session, testing with the ''normal'' mapping was repeated to look for aftereffects of the learned remapping.

Subjects were seated inside a 5-ft-radius arc of 13 light bulbs, spaced every 10 degrees in azimuth from -60 to $+60$ degrees, which were labeled (left to right) with the numbers 1–13. In each test run, a 500-ms-long wideband click train was simulated from each of the possible locations exactly twice, in random order. Subjects were asked to identify the source azimuth corresponding to the simulated source position while facing straight ahead.

In ''training'' experiments (experiments T_1 and T_3), subjects were not provided with any feedback during these test runs. Instead, they were expected to learn about the transformation of localization cues during training runs (interspersed with the test runs) in which synthetic auditory and real visual light sources were simultaneously turned on from one of the 13 possible locations, chosen at random. During these 10-min-long runs, subjects were instructed to turn their heads to face each audiovisual target. Once they faced the target, the light/sound source was turned off and a new random location turned on. In the training experiments, subjects performed two ''normal-cue'' test runs, five ''altered-cue'' test runs, followed by three ''normal-cue'' test runs.

Training was achieved in the ''feedback'' experiments (experiments F_3, F_{3mid}, F_2, F_{4a}, and F_{4b}) by turning on the

[a]Current address: Boston University, Department of Cognitive and Neural Systems, Room 311, 677 Beacon St., Boston, MA 02215, Electronic mail: shinn@cns.bu.edu

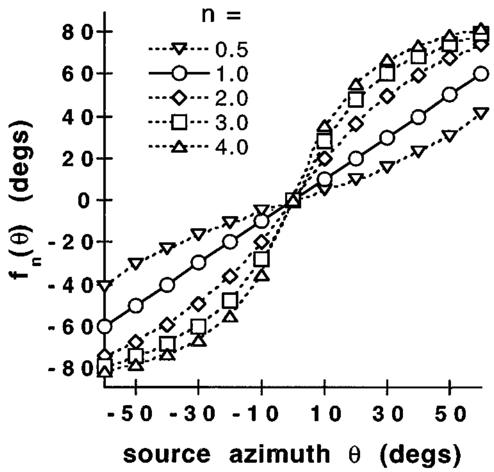

FIG. 1. Plot of the family of functions used to transform auditory localization cues. With this transformation, a source from azimuth θ was synthesized using the HRTF that normally corresponded to the position $f_n(\theta)$.

light at the ''correct'' location for 500 ms after a response in a test run. In these experiments, each subject performed 2 ''normal-cue'' runs, 30 ''altered-cue'' runs, followed by 8 posttraining runs. In all but one feedback experiment, the posttraining runs all used normal cues. In experiment F_{4b}, the eight posttraining runs consisted of four runs in which $n = 0.5$ followed by four ''normal-cue'' runs.

The experimental conditions tested are summarized in Table I. In the previous analysis, bias and resolution were estimated at various points during the course of the experiment using a maximum likelihood estimation technique. In general, bias (a measure of the average error in response, normalized by the standard deviation in responses) grew smaller as subjects were exposed to the remapped cues, consistent with subjects learning the new mapping of physical cue to source location (Shinn-Cunningham *et al.*, 1998). However, errors remained, even after performance had stabilized. In addition, localization errors were not uniform, but varied with stimulus azimuth. The ability to resolve adjacent response locations showed an abrupt change when the remapping was introduced, as expected: resolution improved for stimuli which were physically more distinct than in the ''normal'' mapping (sources in the front region, as seen in Fig. 1) and decreased for stimuli which were more similar than normal (for $n>1$, this occurs for sources at the edges of the range). As subjects adapted to the remapped cues, however, their ability to resolve the *same* physical stimuli showed an overall decrease, indicating that subjects confused adjacent stimuli more often following training than prior to training. It was shown that the amount by which subject performance changed depended primarily on the strength of the transformation [i.e., the value of n in Eq. (1)] and the range/number of source positions presented (Shinn-Cunningham *et al.*, 1998). Conversely, results were largely independent of the complexity of the simulated auditory field. In addition, changes in bias and resolution after exposure to the altered cues were similar for both training and feedback experiments.

TABLE I. Summary of experiments performed. The altered-cue transformation ''strength'' [defined in Eq. (1)] is given in the second column. The ''Exp type'' describes whether subjects were exposed to training runs or given correct-answer feedback. The number of source positions used in the experiment is given in column 5, and the number of acoustic sources simulated in the experiment (target plus additional background sources) is given in column 6.

Experiment	n	Exp type	Position	Sources
T_1	3	training	13	1
T_3	3	training	13	3
F_3	3	feedback	13	3
F_{3mid}	3	feedback	7	3
F_2	2	feedback	13	1
F_{4a}	4	feedback	13	1
F_{4b}	4, 0.5	feedback	13	1

The current paper examines in greater detail how subjects' mean responses change over time when trained with supernormal localization cues. One goal of this analysis is to determine why subjects do not adapt completely to the supernormal cue remapping (i.e., why systematic biases remain in subject responses, even after subject performance has stabilized). In addition, the exact time course of changes in mean response is examined in order to help develop a quantitative model of adaptation (Shinn-Cunningham, 1998).

I. RESULTS

A. Mean response

Subject responses were averaged for each position within each run by combining results from the eight identical experimental sessions performed by each subject. Since each position was presented twice in each run of each session, the mean response for each position was the average of 16 responses. The resulting mean responses were then analyzed to see how mean response related to the physical cues presented to the subjects. An example of mean response for a typical subject (taken from experiment T_1) is shown in Fig. 2. In this figure, the mean response is shown as a function of the normal-cue location of the HRTFs [$f_n(\theta)$ in Eq. (1)]. Starting in Fig. 2(a) and examining the panels in counterclockwise order, the figure shows how mean response changes over the course of the experimental session for the same physical stimuli. In each panel, the correct responses when normal cues ($n=1$) are used are shown by the diagonal solid line. When transformed cues are used, correct responses are shown by the dashed curve [given by the inverse of the transformation function shown in Fig. 2(a)]. Note that in this figure, the HRTF locations presented to the subjects range from -60 to $+60$ degrees in panels (a)–(c) and from approximately -80 to $+80$ degrees in panels (d)–(f). In the lower panels, the range of HRTF positions is greater than the range of response positions (-60 to $+60$ degrees) because the HRTF locations are transformed by the supernormal-cue transformation; however, the feedback presented to the subjects trains them to interpret these physical cues as arising from positions ranging from -60 to $+60$ degrees (examine the dashed line in all panels).

Two important trends are seen in Fig. 2 that are evident for all subjects in all experiments. First, the training with the transformed cues does affect mean response as expected. For

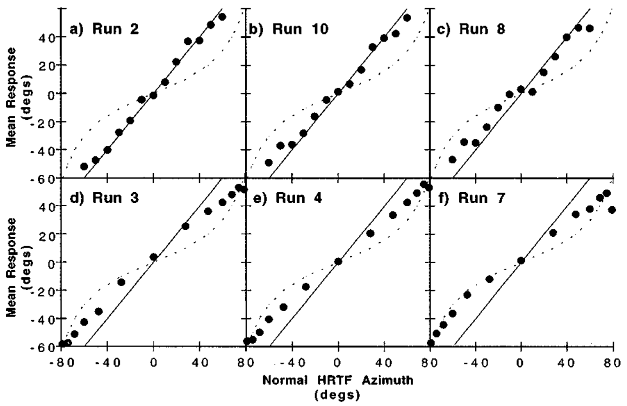

FIG. 2. Plot of mean response for one subject in one experiment as a function of the normal spatial position of the auditory localization stimuli used. In panels (a)–(c), localization cues were not transformed and normal spatial positions ranged from -60 to $+60$ degrees. In panels (d)–(f), transformed localization cues were presented, so that the normal spatial position of the cues ranged from -80 to $+80$ degrees. The order of the runs within the experiment is given by starting in panel (a) and moving counterclockwise to panel (b). The solid line shows the correct response for normal-cue runs. The dashed curve shows the correct response for altered-cue runs. In general, mean response was roughly a linear function of the normal spatial position of the stimuli, even in the altered-cue runs (runs 3, 4, and 7). Subject adaptation is exhibited by changes in the slope of the line relating mean response to normal spatial position over time.

the first two runs, which took place prior to any transformed-cue training, mean response is roughly equal to the correct response [compare the mean response, filled circles, to the correct response, solid line, for run 2 in of Fig. 2(a)]. Thus, for these runs, mean response is proportional to the normal-cue locations with a slope of one. As subjects are exposed to the transformed cues [runs shown in panels (d)–(f)], the mean responses change such that the error in mean response (difference between mean response, filled circles, and the correct response, dashed curve) decreases, even though it does not go to zero. When tested and trained with normal cues following exposure to the transformation [run 8 in panel (c) and 10 in panel (b)], the mean responses begin to change back towards their original pattern.

The second trend seen in the data is that mean response is roughly proportional to the location normally associated with the HRTFs used, even after subjects have been trained using the nonlinear transformation. This result is most obvious when examining Fig. 2(d)–(f). In these runs, correct responses (as defined by the training given subjects) are nonlinear as a function of the normal-cue source location (dashed lines); however, the mean response is roughly linear (with a y intercept of zero) for the transformed-cue test runs, even after 30 min of training with the transformed-cue transformation [cf. results from run 7 in Fig. 2(f)].

B. Test of best-fit hypotheses

The two trends discussed above are seen in the data from all subjects in all experiments. The second observation, that mean response is always proportional to the normal location of the physical HRTFs used, was explicitly tested for each subject in each experiment. Two hypotheses were used to fit mean response as a function of normal-cue source position. Under both hypotheses, a single parameter summarizes how perceived source position is related to the physical stimuli in a given run.

The first hypothesis, suggested by the data, was that in run r, the mean perceived position p of a source whose ''correct'' location is θ is given by

$$p(\theta,r)=k(r)f_n(\theta) \qquad (2)$$

for stimulus $f_n(\theta)$, where the value of $f_n(\theta)$ is determined by the altered-cue transformation used in the experiment [see Eq. (1)]. Under this hypothesis, changes in subject performance are characterized by changes in the best-fit slope $k(r)$ with run r. This hypothesis assumes that subjects cannot adapt to the transformation completely; instead, the average perceived position of the source is constrained to be proportional to the normal-cue source position $f_n(\theta)$. In other words, rather than adapting fully to the given transformation, the subject adapts to a linear approximation to the transformation. For a naive subject, $k(r)$ should start near 1, indicating that the subject perceives localization cues normally. As the subject adapts, $k(r)$ should decrease, consistent with the subject altering his responses to reduce his mean localization error. However, some error between mean perceived position $p(\theta,r)$ and correct response θ will remain, since the imposed transformation $f_n(\theta)$ is nonlinear but subjects can only adapt to a linear transformation.

An alternative hypothesis is suggested by Eq. (3). In this hypothesis, it is assumed that the mean perceived position is given by

$$p(\theta,r)=f^{-1}_{m(r)}(\theta), \qquad (3)$$

where $f_n^{-1}(\theta)$ is the inverse of the function $f_n(\theta)$ defined in Eq. (1) and $m(r)$ varies between 1 and n. The value $m(r)=1$ indicates that subjects perceive source locations in the normal manner. Conversely, the value $m(r)=n$ is consistent with subjects adapting completely to the supernormal transformation $f_n(\theta)$. Thus, this hypothesis assumes that any nonlinearity in mean response will take the form of the nonlinear transformation used to rearrange the acoustic cues. With this hypothesis, subjects achieve zero mean error for all source positions when $m=n$. Partial adaptation occurs when $1<m<n$.

For each subject, run, and experiment, the values of $k(r)$ and $m(r)$ that minimized the mean squared error between measured mean response and predicted mean response was calculated. For most stimuli, the predicted mean response was simply the value of $p(\theta,r)$ predicted from Eq. (2) or (3). However, since responses were constrained to fall between -60 to $+60$ degrees predicted responses were also constrained to fall within the closed interval $[-60,60]$ degrees. Any predictions falling outside this range of possible responses were set to equal the nearest extreme value [i.e., if the mean response calculated from Eq. (2) or (3) was equal to 65 degrees, the prediction used to calculate the mean-square error was 60 degrees].[1] The minimum mean square error (in degrees) for the two curve fits were then compared for each run and subject in each experiment.

Table II summarizes the results of these tests. In general, both approaches fit the data reasonably well. The average

TABLE II. Comparison of linear and nonlinear curve fitting of mean response. Column 1 gives the experiment number, columns 2 and 3 show the mean square error (in degrees) for linear and nonlinear curve fits, respectively. Column 4 shows the number of cases for which the linear curve fit yielded a better fit than did the nonlinear curve fit. Column 5 shows the total number of cases in each experiment. Column 6 shows the percentage of the cases in which the linear curve fit yielded better results. The final row shows average results across all experiments.

Experiment	mse linear	mse *x* form	Linear better	Total cases	Percent
T_1	1.3	1.5	33	40	82.5
T_3	1.7	1.7	41	80	51.3
F_3	1.0	1.4	182	200	91.0
F_{3mid}	0.7	0.8	104	160	65.0
F_2	1.5	1.6	109	160	68.1
F_{4a}	2.0	3.0	113	120	94.2
F_{4b}	1.7	2.6	113	120	94.2
overall	1.3	1.7	695	880	79.0

mean square error in the predictions was 1.3 degrees using the linear curve fit and 1.7 degrees using the nonlinear curve fit. More importantly, however, in almost all of the conditions, the linear curve fit yielded a smaller mean square error than the nonlinear curve fit. Table II shows that the linear curve fit yielded a smaller mean square error than did the nonlinear curve fit in 695 of 880, or 79% of the cases. Both hypotheses will fit normal, unadapted responses equally well [i.e., $k(r)=m(r)=1$ yield identical predictions]. Thus, for runs in which the mean response is near the ''normal'' response (i.e., prior to exposure to the supernormal cues, and as subjects readapt to normal cues at the end of the experiment), there should be no clear advantage to using the linear

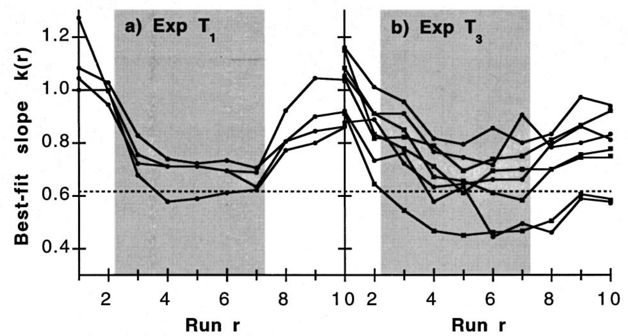

FIG. 3. Plot of best-fit slope $k(r)$ for each subject in experiments with active sensorimotor training. Best-fit slope for the transformation used in each experiment is shown by the horizontal dashed line in each panel. Normal-cue runs are plotted against a white background; altered-cue runs are plotted against a gray background.

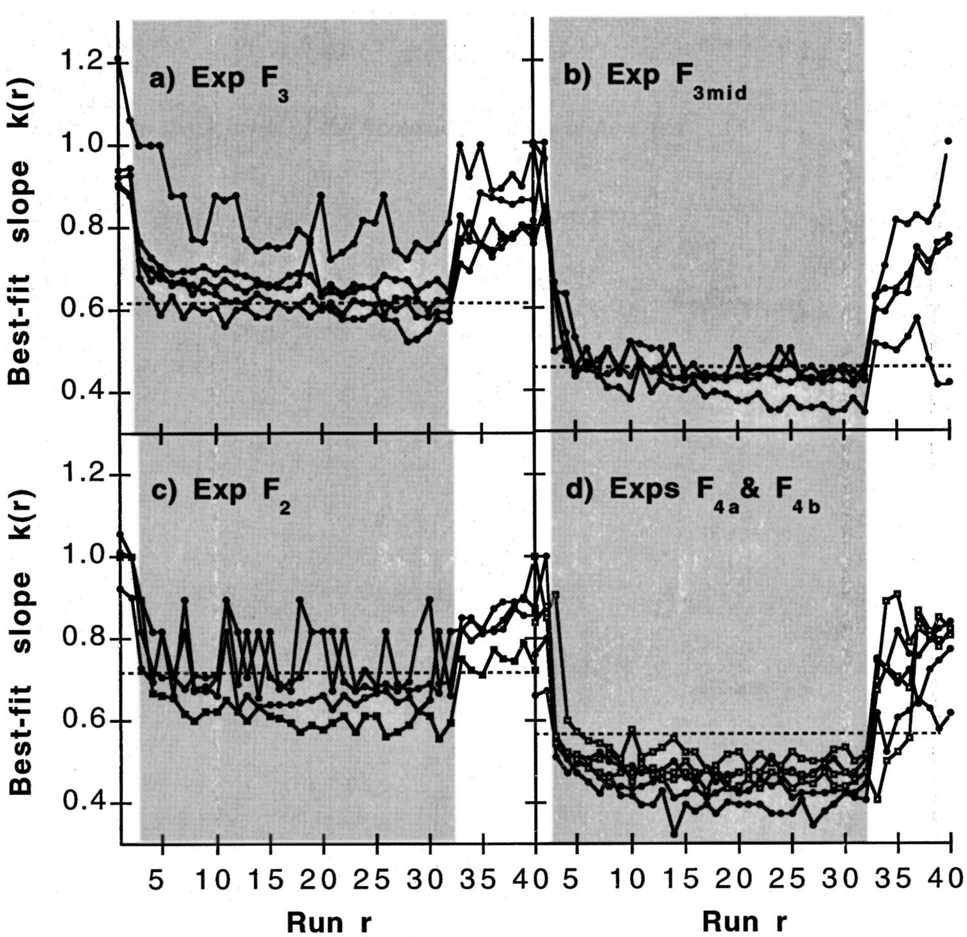

FIG. 4. Plot of best-fit slope $k(r)$ for each subject in experiments with correct-answer feedback. Best-fit slope for the transformation used in each experiment is shown by the horizontal dashed line in each panel. Normal-cue runs are plotted against a white background; altered-cue runs are plotted against a gray background. In panel (d), results are shown for both experiment F_{4a} (solid symbols) and experiment F_{4b} (open symbols), which were identical up through run 32. In experiment F_{4b}, an inverse transformation was used in runs 33–36.

curve fit compared to the nonlinear curve fit; thus, the superiority of the linear curve fit over the nonlinear curve fit is even more compelling.

Although it is possible that other nonlinear curve fits might yield even better fits than those found with the linear curve fit described in Eq. (2), it is unlikely that any other single-parameter curve fit would provide more insight into the pattern of subject responses. If there were any nonlinear trends in the mean response, they should arise from the nonlinear cue transformation employed in the experiment, a hypothesis tested explicitly in the comparison summarized in Table II.

C. Best-fit slope *k*(*r*)

These results demonstrate that the mean response for each subject at each point in time can be summarized by the best-fit slope $k(r)$ defined in Eq. (2). The fact that mean response is roughly proportional to the normal-cue position of the acoustic stimulus implies that subjects cannot completely adapt to a nonlinear cue transformation. Instead, subjects adapt to a linear approximation of the nonlinear transformation employed. Thus, for the type and length of training used in these experiments, some error will remain in their mean responses. Let k_{opt} denote the slope which will minimize the mean squared error in mean response for a given experiment. Given that mean perceived position is constrained by Eq. (2), this slope describes the most complete adaptation achievable for a given experiment. Thus, if subjects adapt as completely as possible, $k(r)$ should approach k_{opt} asymptotically over the course of the exposure period in the experiment. The value of k_{opt} will depend not only on the strength of the transformation employed, but also on the source positions used in the experiment.

Figure 3 plots $k(r)$ as a function of run r for each of the experiments in which active sensorimotor training runs were used in between test runs, while Fig. 4 plots $k(r)$ for the experiments in which correct-answer feedback was employed during each test run. In both figures, each panel shows the best-fit values of $k(r)$ for each subject with a solid line. The dashed horizontal line in each panel shows the best-fit slope k_{opt} for that experiment. Runs in which the altered cues were used are shown with a gray background, while normal-cue runs are shown against a white background.

The plots in Figs. 3 and 4 show that the parameter $k(r)$ changes with run number, consistent with subjects adapting to the supernormal cues over time. In all of the experiments, independent of the experimental conditions, the value of $k(r)$ approaches a value which is roughly equal to k_{opt}, the optimal slope for the given experiment. The change in $k(r)$ is quite rapid; most subjects have slopes near their final asymptotic value by the fifth run with altered cues. When subjects are tested with normal cues following the supernormal-cue exposure period, the value of $k(r)$ increases towards the ''normal-cue'' optimal value of 1.

While there is substantial intersubject variability in the absolute values of $k(r)$, all subjects showed a rapid change in $k(r)$ over the course of the experiment. In general, the

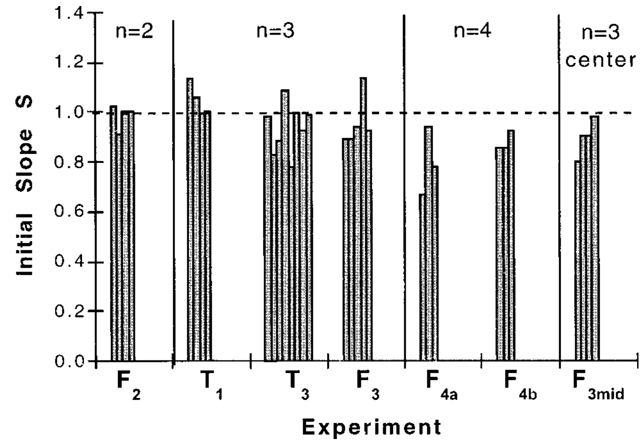

FIG. 5. Estimates of initial values of k prior to exposure to the supernormal (transformed) cues. Each bar represents the estimate of the initial slope for one subject. The horizontal dashed line shows the expected value of initial slope of 1.0.

asymptotic value of $k(r)$ is close to the best-fit slope of k_{opt}. In experiments T_1, T_3, and F_3, the same transformation ($n=3$) and range of positions (from -60 to $+60$ degrees) were presented, so that k_{opt} is equal to 0.6161 in all of these experiments. Looking at the results across all subjects in these three experiments, the final value of $k(r)$ tends to be equal to or greater than the optimal value of k_{opt}. However, there are subjects for whom the best-fit slope $k(r)$ was actually less than k_{opt} as well. Given the intersubject variability, the value k_{opt} is a good predictor for the asymptotic values of slope $k(r)$. In experiment $F_{3\mathrm{mid}}$, k_{opt} is equal to 0.4565 ($n=3$, but only the middle seven response positions, from -30 to $+30$ degrees were presented), while in experiment F_2, k_{opt} is equal to 0.7142 ($n=2$, all 13 source positions used). In these experiments, the majority of the subjects adapt to values of $k(r)$ extremely close to k_{opt}. Finally, in experiments F_{4a} and F_{4b}, in which the transformation was strongest ($n=4$) and all 13 source locations were presented, on average, the asymptotic values of $k(r)$ tend to be smaller than the calculated value of $k_{\mathrm{opt}}=0.5646$.

In experiment F_{4b}, subjects were presented with a transformation of strength 0.5 for four runs (runs 33–36), then normal cues for the final four runs (runs 37–40). In order to examine the effect of this exposure period, slope estimates from runs 33–40 from experiment F_{4a} (which was identical to experiment F_{4b} except that normal cues were presented for runs 33–40) were compared with slope estimates for experiment F_{4b}. As can be seen in Fig. 4(d), results from the three subjects in experiment F_{4b} (small diamond symbols) were indistinguishable from results from the three subjects in experiment F_{4a} (small circle symbols), despite the training with the inverse transformation.

II. FITTING *k*(*r*)

The graphs shown in Figs. 3 and 4 imply that changes in mean localization performance can be summarized by $k(r)$, which appears to change exponentially towards an asymptotic value. In order to better quantify these results, the

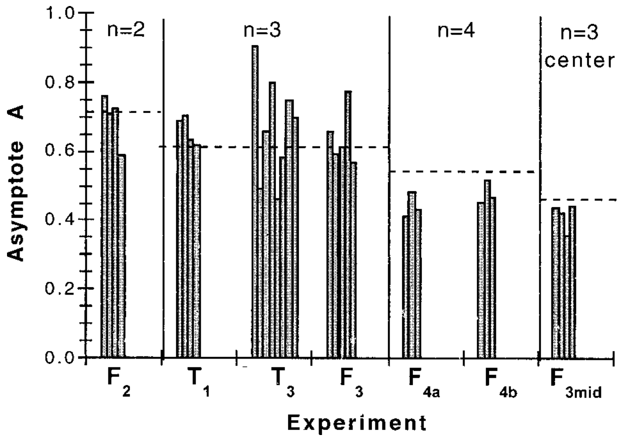

FIG. 6. Estimates of asymptotes of $k(r)$ during adaptation. The best-fit slopes (k_{opt}) for the transformations used in each experiment are shown by the horizontal dashed lines. Each bar represents the estimate of the asymptote for one subject.

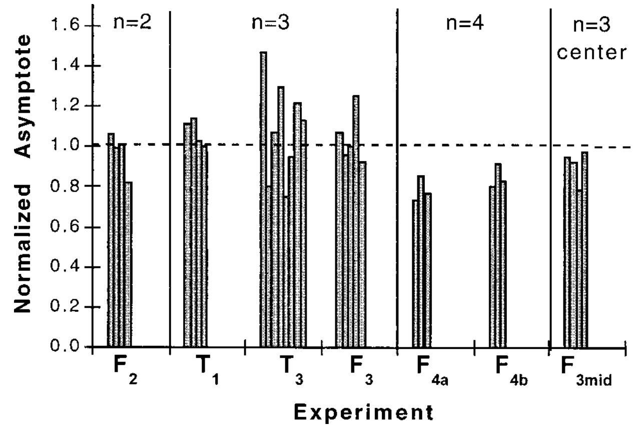

FIG. 7. Normalized asymptotes of $k(r)$. Each bar represents the estimate of the asymptote for one subject, normalized by the best-fit slope (k_{opt}) for that experiment. The dashed line is at 1.0, the normalized value of k_{opt} for all experiments.

values of $k(r)$ found for each subject were analyzed to determine more precisely how performance changed with time.

A. Initial, unadapted slope

The first quantity to be examined was the starting value of $k(r)$. Although some subjects may show a change in $k(r)$ between the first two runs in the experiment (both of which use ''normal'' localization cues), these changes are relatively small compared to the changes which occur as subjects adapt to the intentionally altered cues. Also, while many subjects in experiments T_1 and T_3 show a tendency for $k(2)$ to decrease from the initial value of $k(1)$, subjects in the subsequent experiments were equally likely to show an increase or a decrease in slope between runs 1 and 2. For this reason, estimates of the initial slope prior to training with the supernormal cues were found by averaging the values of $k(1)$ and $k(2)$ for each subject.

Figure 5 plots estimates of S, the initial slope prior to exposure to altered cues, for each subject in each experiment. In order to be consistent with later plots, the data are plotted in a bar graph in which subjects from each experiment are grouped together, and the experiments are ordered on the basis of k_{opt}, the best-fit slope for the transformation of cues used in that experiment.

The plot in Fig. 5 shows little systematic dependence of the initial slope on either experiment or on the transformation strength used in the experiment. This was explicitly tested with a one-way ANOVA, in which S showed no statistical dependence on experiment $[F(6,1)=2.86,\ p>0.005]$. Although there is no statistical dependence of S on experiment, there is obvious subject variability in the estimates of initial slope. Subjects in experiments F_{3mid}, F_{4a}, and F_{4b} tended to have a smaller initial slope than did subjects in other experiments; however, the initial slope values shown by the subjects in these two experiments is within the range of values shown by subjects in other experiments (e.g., note values in experiment T_3 or F_2). On average, across all experiments, the initial slope is 0.94, close to the expected slope of unity for subjects hearing normal cues, prior to exposure to altered cues.

B. Adaptation asymptote

For each subject in each experiment, the asymptote of $k(r)$ during the altered cue exposure was estimated from the values of $k(r)$ plotted in Figs. 3 and 4. For the training experiments (shown in Fig. 3), the asymptote was estimated as the value of $k(7)$, the value from the final, altered-cue run in the experiments. For the feedback experiments (shown in Fig. 4), the asymptote was estimated by averaging the values of $k(30)$, $k(31)$, and $k(32)$, the slopes from the last three runs using the altered cues.

Figure 6 shows estimates of the asymptotes (denoted by A) for each subject in each experiment. In the graph, the data are grouped according to the value of k_{opt}, the best-fit slope fitting the transformation used in that experiment. In each group, k_{opt} is shown by a dashed horizontal line. A one-way ANOVA showed that A has a clear statistical dependence on the experiment $[F(6,1)=6.46,\ p<0.005]$.

In order to examine the extent to which the dependence of the asymptote A on experiment can be accounted for by differences in the best-fit slope k_{opt}, the estimates of asymptote were normalized by the best-fit slope and replotted in Fig. 7. Plotted in this way, the best-fit slope has a value of 1.0 for all experiments (shown by the dashed horizontal line). When normalized by the best-fit slope in the experiment, there is little clear systematic variation in the values of asymptote with experiments. Although the normalized asymptotes in experiments F_{3mid}, F_{4a}, and F_{4b} all have values less than 1.0, there are many subjects in the other four experiments for whom the normalized asymptotes is less than 1.0 as well. It is also interesting to note that S, the *initial* slope value, tends to be less in these three experiments than in the other four experiments (see Fig. 5). This difference may occur because the subjects in these experiments all tend to interpret the ''normal'' HRTFs [taken from subject SDO, Wightman and Kistler (1989)] as closer to straight ahead than the nominal position of the HRTF. In other words, normal intersubject variability in HRTFs may be responsible for the tendency of subjects in some experiments to have smaller-than-expected slopes. A second possibility is that this tendency is the result of training with supernormal cues.

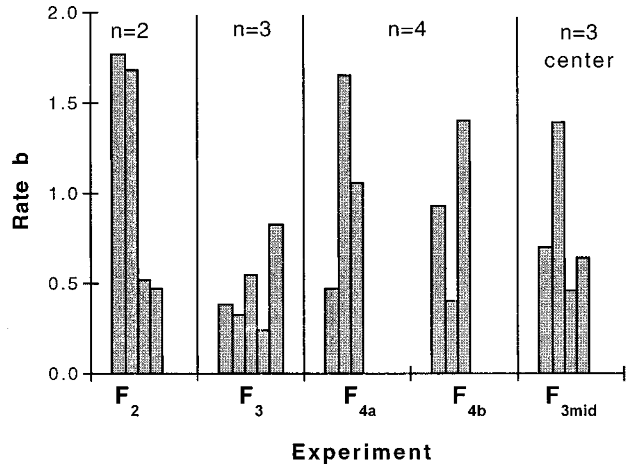

FIG. 8. Rate of adaptation b. Each bar represents the estimate of the exponential rate of adaptation (in units of run^{-1}) for one subject.

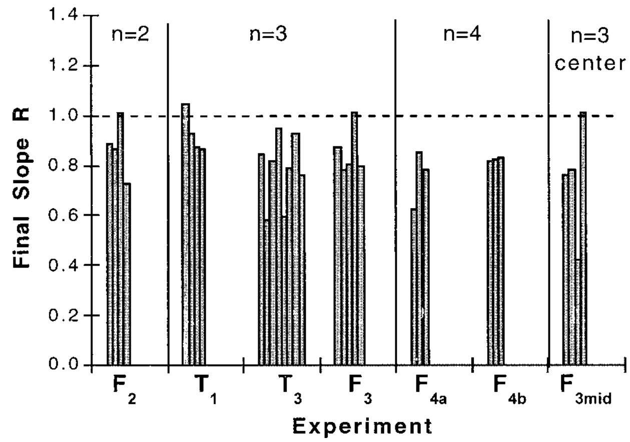

FIG. 9. Estimate of slope at end of experiment after retraining with normal cues. Each bar represents the final estimate of k at the end of the experiment for one subject. The solid error bar shows the initial estimate of slope for that subject (as shown in Fig. 4). The dashed line is at 1.0, the optimal slope for normal-cue runs.

The three experiments in question (in experiments F_{3mid}, F_{4a}, and F_{4b}) used supernormal transformations with the smallest values of k_{opt} and the greatest change in perception from normal. Thus, any supernormal training effects would be greatest in these experiments. Training effects can only explain the tendency for the initial slope S to be less than one if it is assumed that training carries over from session to session, a hypothesis that was not evident when examining the results across sessions. In any case, a one-way ANOVA failed to find any statistically significant differences between the values of normalized asymptote in the different experiments $[F(6,1)=2.28,\ p>0.005]$, further implying that intersubject variability is the likely explanation for the smaller-than-average slope values in experiments F_{3mid}, F_{4a}, and F_{4b}. Thus, intersubject variability is the most parsimonious explanation for the tendency for both initial slope and normalized asymptote to be smaller in experiments F_{3mid}, F_{4a}, and F_{4b} than in other experiments. Since there is no statistically significant dependence of normalized asymptote on experiment, the statistical variability in absolute asymptote A can be accounted for (at least in part) by differences in the best-fit slope k_{opt}.

From examining Fig. 7, it is also clear that the best-fit slope in each experiment is a relatively good predictor for the asymptote. The value of the normalized asymptote is close to 1.0 for most experiments. The normalized asymptote averaged across all subjects in all seven experiments equals 0.99, further showing that the best-fit slope in each experiment is, on average, equal to the asymptotic value of $k(r)$.

C. Rate of adaptation

The rate at which adaptation occurs is quite rapid, as evidenced by the data shown in Figs. 3 and 4. This observation was further refined by assuming that the curves plotted in Fig. 4 changed exponentially (as a function of run) from their initial values S (found in Sec. II A) towards their asymptotic values A (found in Sec. II B) and estimating the decay value of the exponential. For each experiment, $k(r)$ was thus assumed to vary as

$$k(r)=A+(S-A)e^{-b(r-2)}, \qquad (4)$$

where A is the asymptote, S is the initial slope, r is the run number, and b is the rate of adaptation (in units of run^{-1}). With this equation, the slope in the second run equals the initial value of S. As r increases, $k(r)$ approaches A asymptotically.

Estimates of the adaptation rate b were found for each subject in each experiment by finding the value of b which minimized the mean-square error between the estimate [given by Eq. (4)] and the actual values of $k(r)$ which are plotted in Fig. 4.[2] These estimates are shown in Fig. 8 for the feedback experiments.

It is clear from Fig. 8 that estimates of rate show great intersubject variability. Values range from roughly 0.3 to 1.8. No systematic dependence of rate on experiment or on transformation is evident in the data. A one-way ANOVA on b confirmed this observation, showing no statistical dependence of b on experiment $[F(4,1)=1.19,\ p>0.005]$. The average value of b (across the five experiments) equals $0.84\ \mathrm{run}^{-1}$. In other words, on average, by the sixth altered-cue test run $(r=8)$, the slope has changed by 99% of the total change expected after infinite training [from Eq. (4)].

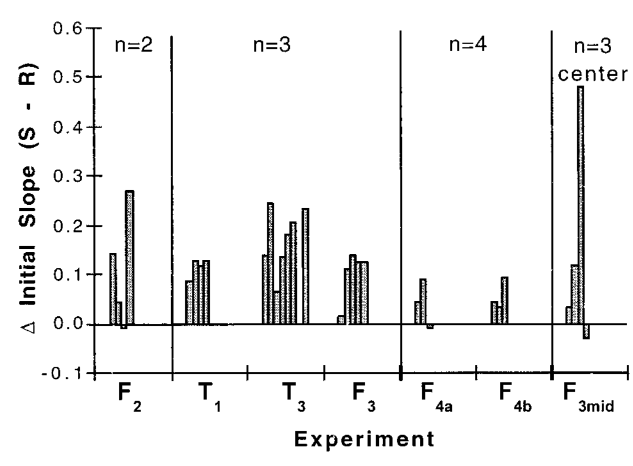

FIG. 10. Difference between initial and final slope (S-R). Each bar plots the difference between the initial slope estimate and the final slope estimate.

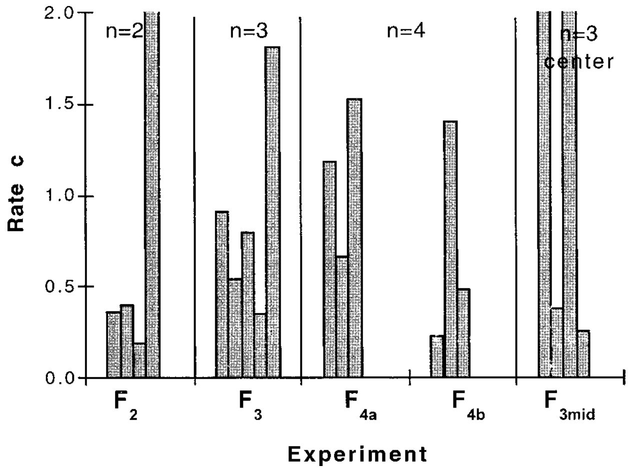

FIG. 11. Estimates of c, rate of adaptation back to normal cues. Each bar represents the estimate of the exponential rate of adaptation (in units of run^{-1}) for one subject.

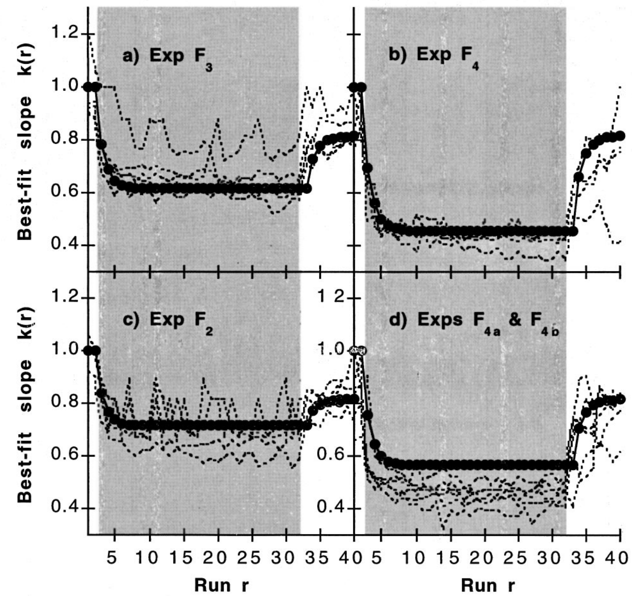

FIG. 12. Predictions of $k(r)$ for each experiment. Dashed lines replot individual subject data, while predictions are given by filled circles and solid lines.

D. Final slope

The final values of $k(r)$ were compared across all experiments to determine whether subject performance returned to normal by the end of the experiment (after retraining with normal cues). The final slope R was estimated as the value of $k(10)$ for the training experiments and $k(40)$ for the feedback experiments. These values are plotted in Fig. 9.

In general, values of the final slope are slightly less than one. In fact, the average value of R across all experiments equals 0.82, which is less than 0.94, the average value of starting slope S. This is shown explicitly in Fig. 10, which plots the difference between initial and final slope estimates for each subject. This difference is generally positive, indicating that subjects have not returned to normal performance after retraining with normal cues. Of most interest, there is no difference between the results of experiments F_{4a} and F_{4b}, despite the fact that subjects in experiment F_{4b} were trained with an inverse transformation for four runs prior to returning to normal cues. One-way ANOVAs showed that neither final slope R [$F(6,1)=0.4211$, $p>0.005$] nor the difference between initial and final slopes (S-R) [$F(6,1)=0.711$, $p>0.005$] were statistically dependent on experiment. However, initial slope was statistically greater than final slope according to a paired t test [$t(56)=3.756$, $p<0.005$].

It is likely that with sufficient time, subjects would readapt back to their original state, and the values of R would approach the starting slope values of S. However, the rate of such a change must be very slow relative to the rate of adaptation to supernormal cues. After eight runs of altered cues, most subjects had neared their asymptotic levels of performance; when returning to normal cues, performance was significantly different from performance prior to exposure to the altered cues. While this difference is statistically significant, the practical impact of this difference is relatively small. For instance, by the end of the readaptation period, the average difference in the mean response for a source at 10 degrees azimuth is only 1.2 degrees, a change within the average standard deviation in response for a source at that location.

E. Rate of readaptation

Readaptation to normal cues takes place during runs 33–40 for the feedback experiments. Since the exposure to the inverse transformation in experiment F_{4b} had no discernible effect on the slope estimates shown in Fig. 4, all five of these experiments were treated identically in trying to quantify how performance changed back towards normal at the end of the experiment.

As with the adaptation period (runs 3–32), the slope k during the readaptation period was fit by an exponential curve given by

$$k(r)=R+(A-R)e^{-e(r-32)}, \tag{5}$$

where r is the run number. The final slope R was determined as described in Sec. II D, while the initial value (in run 32) was determined by the processing described in Sec. II B. Estimates of the rate c were found using the same algorithm described in Sec. II C.

Estimates of c are shown in Fig. 11. There is even more intersubject variability evident in these estimates than was seen in the estimates of b (see Fig. 8). In part, this is due to the fact that the estimates of c depend directly on the estimates of R. Estimates of R depend on only one value of $k(r)$ and are therefore noisier than are estimates of A, which were derived by averaging three values of $k(r)$. In addition, performance had clearly stabilized for all subjects by run 32 during the adaptation period, so that estimates of A will be relatively good predictors of the actual final asymptote. It is less evident that performance during the readaptation period had stabilized by the 40th run. If the actual rate of readaptation is relatively slow (and performance had not stabilized by the end of the experiment), then estimates of R will be less than the final asymptote, and estimates of rate will be larger than their actual values. In a number of cases, estimates of R are less than many other values of k in runs 33–39. When this occurs, estimates of c will be extremely large. However, on the whole the estimates of c shown in Fig. 11 are roughly equivalent to the estimates of b shown in Fig. 8. The average value of b, the rate of adaptation, equaled $0.84\ \text{run}^{-1}$, while the average value of c, the rate of readaptation, equaled $0.94\ \text{run}^{-1}$. A paired t test found no statistical difference between b and c [$t(28)=0.44$, $p>0.005$]. In addition, ANOVA analysis showed that there was no statistical dependence of c on experiment [$F(4,1)=0.08$, $p>0.005$].

F. Predictions of $k(r)$

While there is large intersubject variability in all of the parameters determined in the previous sections, the average values of adaptation rate, starting and ending slope, and normalized asymptote appear to summarize results of all experiments.

Equations (4) and (5) were used to predict $k(r)$ for the feedback experiments by setting $S=1.0$, the optimal initial slope; $A=k_{\text{opt}}$, the optimal slope during the adaptation period; $R=0.82$, the average final slope across the five experiments; and b and c set equal to 0.84 (the average value of b across the five experiments). As such, the predictions should show how a typical, idealized subject adapts over time.

Figure 12 plots the predictions of $k(r)$ for each experiment (solid lines and filled circles). Individual subject data from each experiment (repeated from Fig. 4) is shown in the same figure (dashed lines) for direct comparison.

On average, the predicted curves are quite close to the average results for each experiment. While the predictions cannot capture the large intersubject variability in the data, the fit is quite reasonable. The predictions lie above the actual subject data for experiments F_{4a} and F_{4b}; however, for all other experiments, the predictions fall well within the range of results seen across subjects in the experiments. With only two free parameters (R, the final slope asymptote, and b, the rate of adaptation and readaptation), results from 31 subjects are well summarized.

III. DISCUSSION

Subjects adapt to changes in auditory localization cues; however, there are limits to the adaptation they exhibit. In the current experiments, auditory cues were transformed using a nonlinear transformation; however, subjects adapted to a linear approximation of the transformation. The final slope relating mean response to normal-cue position is roughly equal to the slope of the line which best approximates the transformation employed. This result implies that there may be limits on the plasticity of human subjects in interpreting auditory localization cues. In particular, subjects may be able to accommodate only linear transformations of cues, rather than being able to adapt to arbitrarily complex remappings.

An alternative explanation is that the nonlinearities in the remappings used in the current experiments are not sufficient to cause subjects to adapt to the exact shape of the transformation. Under this hypothesis, subjects may be capable of adapting to nonlinear cue transformations, but only if the nonlinearities in the cue remapping are extreme. If adapting to a simple approximation of the remapping removes most of the error in subject responses, then there may not be sufficient impetus for subjects to adapt to the more complex remapping.

If subjects cannot adapt to nonlinear transformations at all, there are fundamental implications about the way in which subjects interpret auditory localization cues. In addition, similar fundamental limitations may affect sensorimotor rearrangement in other sensory modalities [e.g., possible constraints on visual-motor rearrangement are discussed in Bedford (1989, 1993); recent experiments by Schloerb (1997) have shown that adaptation to interpupillary distance are consistent with this type of constraint as well]. It is interesting to note that subjects, on average, adapted to the best-fit linear approximation of the nonlinear cue transformations used in the current experiments (although some systematic errors in localization judgments, due to the nonlinearity of the employed transformations, remained). In most previous studies of auditory adaptation, subjects appeared to adapt only partially to sensory rearrangements (e.g., see Freedman and Wilson, 1967; Freedman and Gardos, 1965; Freedman and Stampfer, 1964a, 1964b; Freedman *et al.*, 1967; Freedman and Zacks, 1964; Held, 1955; Kalil and Freedman, 1967; Lackner, 1974; Mikaelian, 1974; Mikaelian and Russotti, 1972). This apparent failure to adapt completely may be the result of an inability to adapt perfectly to the type of transformation employed.

In any case, for the current experiments, the single value of the slope $k(r)$ (relating physical cue to mean response) summarizes the adaptive state of the subject during run r. Subject adaptation is exhibited by exponential changes in slope $k(r)$ from a ''normal'' value of 1.0 at the beginning of the experiment towards an asymptotic value roughly equal to the best-fit slope (assuming that mean response was constrained to be proportional to the normal location of the physical stimuli). At the end of the experiment, the slope changes back towards the nominal value of 1.0, but appears to asymptote to a level less than the normal value (or perhaps to change extremely slowly back toward the value of 1.0).

While there is large intersubject variability, a simple exponential model of the changes in $k(r)$ with run r was able to fit the data across all the experiments relatively well. In this model, subjects begin with $k=1$. As training proceeds, k changes exponentially towards the best-fit slope for the transformation used in the experiment with an estimated rate of $0.84\ \text{run}^{-1}$. At the end of the experiment, k changes exponentially toward an asymptotic value near 0.82 at roughly the same rate of $0.84\ \text{run}^{-1}$.[3] Two free parameters, the rate of adaptation (assumed equal to the rate of readaptation) and the final asymptote were used to fit data for five separate experiments.

ACKNOWLEDGMENTS

This work was mainly supported by AFOSR Grant Nos. 90-200, 93NL387, and 96NL145. Supplementary support was received from Navy Grant No. N61339-93-C-0047. Portions of this work were presented at the 127th (June, 1994) meeting of the Acoustical Society of America.

[1]This approach to dealing with the edges of the response range is not perfect. In particular, subjects will only have a mean response equal to the most extreme response if they *always* respond with that extreme value. In practice, subjects always show some variability in their responses, so that the mean response for all sources (even those whose actual mean perceived position is outside the range of responses) will never equal either of the most extreme response values. By setting the predictions exactly equal to the extreme response values, there will be a tendency for the least-square error fits to underestimate slope $k(r)$ and overestimate transformation strength $m(r)$. However, this method of processing the data is the most tractable, consistent method for estimating all response values.

[2]Rates were not found for subjects in experiments T_1 and T_3. In these experiments, rate will not be directly comparable to the rates for the remaining experiments due to differences in the experimental methods. In experiments T_1 and T_3, active sensorimotor training occurs between runs, while in experiments F_3, F_{3mid}, F_2, F_{4a}, and F_{4b} training (via correct answer feedback) occurs during each run. Thus, both the absolute time between runs and the amount of training occurring from run to run are not equivalent in the training experiments and the feedback experiments.

[3]An alternative fit to the data in the readaptation period could be made by assuming that performance asymptotes back towards one. With this assumption, the rate of readaptation would be an order of magnitude slower than the rate of adaptation.

Bedford, F. L. (**1989**). ''Constraints on learning new mappings between perceptual dimensions,'' J. Exp. Psychol. **15**, 232–248.

Bedford, F. L. (**1993**). ''Perceptual and cognitive spatial learning,'' J. Exp. Psychol. **19**, 7–530.

Durlach, N. I., Shinn-Cunningham, B. G., and Held, R. M. (**1993**). ''Supernormal auditory localization. I. General background,'' Presence **2**, 89–103.

Freedman, S. J., and Gardos, G. (**1965**). ''Compensation for auditory rearrangement and transfer to hand-eye coordination,'' unpublished manuscript.

Freedman, S. J., and Stampfer, K. (**1964a**). ''Changes in auditory localization with displaced ears,'' Paper presented at the Psychonomic Society, Niagara Falls, Ontario.

Freedman, S. J., and Stampfer, K. (**1964b**). ''The effects of displaced ears on auditory localization,'' AFOSR Technical Report 64-0938.

Freedman, S. F., and Wilson, L. (**1967**). ''Compensation for auditory rearrangement following exposure to auditory-tactile discordance,'' Perceptual and Motor Skills **25**, 961–966.

Freedman, S. J., and Zacks, J. L. (**1964**). ''Effects of active and passive movement upon auditory function during prolonged atypical stimulation,'' Perceptual and Motor Skills **18**, 361–366.

Freedman, S. J., Wilson, L., and Rekosh, J. H. (**1967**). ''Compensation for auditory rearrangement in hand-ear coordination,'' Perceptual and Motor Skills **24**, 1207–1210.

Held, R. M. (**1955**). ''Shifts in binaural localization after prolonged exposure to atypical combinations of stimuli,'' Am. J. Psychol. **68**, 526–548.

Kalil, R., and Freedman, S. J. (**1967**). ''Compensation for auditory rearrangement in the absence of observer movement,'' Perceptual and Motor Skills **24**, 475–478.

Lackner, J. R. (**1974**). ''Influence of visual rearrangement and visual motion on sound localization,'' Neuropsychologica **12**, 291–293.

Mikaelian, H. H. (**1974**). ''Adaptation to displaced hearing: a nonproprioceptive change,'' J. Exp. Psychol. **103**, 326–330.

Mikaelian, H. H., and Russotti, J. S. (**1972**). ''Lack of bilateral generalization of adaptation to auditory rearrangement,'' Technical Report 715, Naval Submarine Medical Research Laboratory.

Schloerb, D. (**1997**). ''Adaptation to perceived depth related to changes of the effect interpupillary distance in computer-graphic stereoscopic displays,'' Doctoral dissertation, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA.

Shinn-Cunningham, B. G. (**1998**) ''Adapting to remapped auditory localization cues: A decision-theory model,'' Perception and Psychophysics (under review).

Shinn-Cunningham, B. G., Durlach, N. I., and Held, R. M. (**1998**). ''Adapting to supernormal auditory localization cues. I. Bias and resolution,'' J. Acoust. Soc. Am. **103**, 3656–3666.

Wenzel, E. M. (**1992**). ''Issues in the development of virtual acoustic displays,'' Presence **1**, 80–107.

Wightman, F. L., and Kistler, D. J. (**1989**). ''Headphone simulation of free-field listening. I. Stimulus synthesis,'' J. Acoust. Soc. Am. **85**, 858–867.