

Learning Reverberation: Considerations for Spatial Auditory Displays

Barbara Shinn-Cunningham

Depts. of Cognitive and Neural Systems and Biomedical Engineering
Boston University
677 Beacon St.
Boston, MA, 02215 USA
+1 617 353 8693
shinn@cns.bu.edu

ABSTRACT

Reverberation has both beneficial and detrimental effects on auditory localization. This paper reviews evidence that listeners adapt to the reverberation in a room. Results show that reverberation degrades perception of source direction, but with experience in a given room, performance improves. Reverberation enhances distance perception but distance accuracy still improves with experience. (The importance of reverberation for distance perception can be heard in the accompanying demonstration.) These results are considered along with results from previous studies investigating how past experience affects spatial perception. Implications for learning mechanisms and for the design of auditory displays are discussed.

Keywords

reverberation, learning, distance, spatial perception, interaural differences

BACKGROUND

Most past studies of the effects of reverberation on spatial hearing fall into one of three categories: studies examining 1) the effect of reverberation on perception of room characteristics, 2) the effect of a single echo (or very small number of echoes) on perception of source direction (precedence effect studies), or 3) the effect of reverberation on distance judgements.

Reverberation has a large impact on the perceived characteristics of a listening environment (e.g., see [17, 5, 27, 6]), in part because reverberation decorrelates the signals reaching the two ears. The perceived “spaciousness” of a room increases with reverberation time, reverberation level, and/or the amount of decorrelation between left and right ear signals. Headphone signals presented without reverberation (or sources presented in an anechoic environment) generally sound very strange, at least to naïve listeners. In contrast, reverberation dramatically improves the subjective “realism” and the amount of “externalization” achieved in a headphone simulation.

When a signal from one direction (the “lead”) is followed closely in time by a signal from another direction (the “lag”), directional information from the second signal has little perceptual impact (a phenomenon known as the “precedence effect;” for a recent review, see [23]). Studies of directional localization in rooms generally show that the effect of reverberation on localization accuracy is not very large, at least when onset information is available to the listener [19]. Extrapolating from these results, one might expect little “learning” of directional cues in a reverberant room, since the echoes have little perceptual impact in the first place. Nonetheless, reverberation physically distorts steady-state “directional” cues like interaural differences and spectral shape, and there is some evidence that more “realistic” reverberation does interfere with directional perception (e.g., see [6, 19, 28]).

Previous work shows that reverberation actually helps listeners judge source distance, at least for relatively distant sources [24]. The underlying cue enabling these distance judgements is thought to depend on the ratio of direct to reverberant sound energy (or some closely-related statistic) [12]. At least to a first order approximation, the level of the direct sound is inversely proportional to the square of the source distance while the level of reverberation is roughly independent of source location. The direct-to-reverberant energy ratio thus provides a distance cue that is independent of overall source level.

Since the level of reverberation varies from room to room, the way in which listeners interpret the total signals reaching the ears must change with listening environment for spatial perception to be accurate. However, few studies have examined how listeners may calibrate spatial perceptions and “learn” a room’s characteristics. In studies of the precedence effect, repetitions of a particular lead-lag stimulus pair increase the strength of the effect and the suppression of the lag (e.g., see [14, 18]). For instance, after 2-5 repetitions, what may have been an audible, distinct lag burst often disappears as the lag is “fused” with the lead into a single event. These results have been interpreted as evidence that subjects learn a particular room reverberation

pattern in order to suppress disruptive spatial cues from echoes. In studies of distance perception in reverberant rooms, distance perception can improve after presentation of as few as 5 stimuli in a reverberant space [24]. These studies provide evidence for rapid learning of echo or reverberation characteristics, even without any explicit training or feedback. Below we present evidence for long-term learning that results in improved localization accuracy in all spatial dimensions.

LOCALIZATION CUES IN A REVERBERANT ROOM

In anechoic space, the most robust and unambiguous sound localization cues are interaural differences in time (interaural time differences or ITDs) and level (interaural level differences, or ILDs). ITD cues are roughly constant for sources on a cone centered on the interaural axis (the well-known “cone of confusion”). ILD cues vary with both cone of confusion (i.e., the “head shadow” effect often discussed in directional hearing) and the relative distance from the source to the left and right ears [38]. At low frequencies, only the relative distance from source to the ears contributes to the overall ILD. For all frequencies, the spatial information in the ILD that is independent of ITD depends only on the distances from source to two ears and is constant on a sphere symmetrical about the interaural axis [38]. Figure 1 shows how ITD and the unique portion of the ILD vary with source location for sources within reach of a listener. Each contour shows all spatial locations for which sources yield essentially the same ITD (left side of the figure) or low-frequency ILD (right side of the figure). The figure shows that while the ITD conveys information about the cone of confusion, ILD cues convey additional information about both source distance and direction when sources are within a meter of the listener.

Spectral cues are most important for disambiguating locations that cannot be distinguished on the basis of binaural cues, and thus are most important for determining the angle around the interaural axis (i.e., the angle around the “torus of confusion;” see [38]). In reverberant space, reverberation is used to determine the distance of the source [24].

Thus, there are four candidate cues that may underlie sound localization behavior in a reverberant room: ITD, ILD, spectral, and reverberation cues. Only three of these cue dimensions are independent, however, since ILD covaries with both the ITD cone-of confusion cue and the reverberant distance cue. Arguably, the three most important independent spatial dimensions in reverberant space are therefore the angle of the cone of confusion, torus angle, and distance. Analysis based on these spatial dimensions should give insight into localization behavior (e.g., see [15]). Figure 2 shows these spatial dimensions: ITD depends on θ , the cone of confusion angle; spectral cues convey ϕ , the “torus angle;” and reverberant cue varies with r , the source distance.

LOCALIZATION LEARNING IN A REVERBERANT ROOM

Objective

An experiment was designed to determine whether reverberation had any significant effect on localization performance for conditions in which reverberation is relatively weak. In particular, we replicated the experimental procedure of Brungart [13], who investigated three-

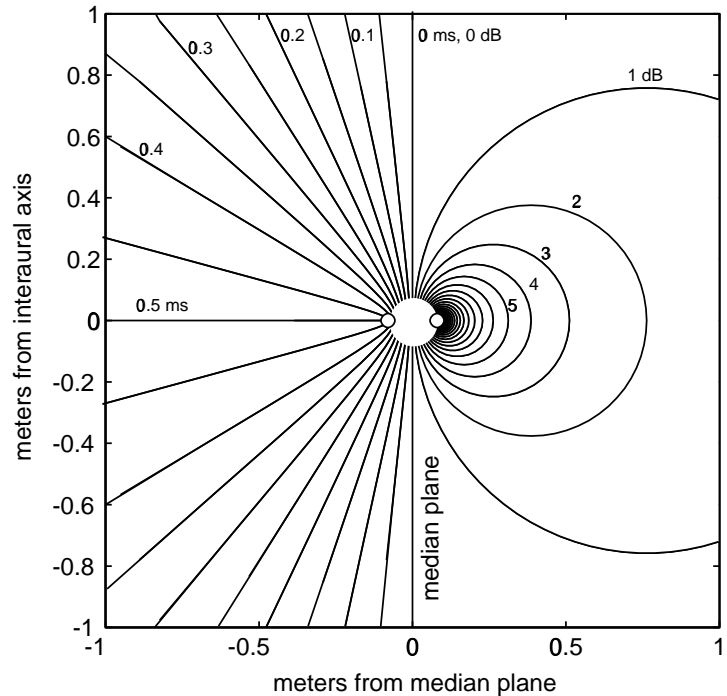


Figure 1: Iso-ITD (left) and iso-ILD contours (right) for sources in a plane containing the interaural axis. Small open circles denote the ears. The ITD or ILD value resulting from source at each spatial location is shown on each contour.

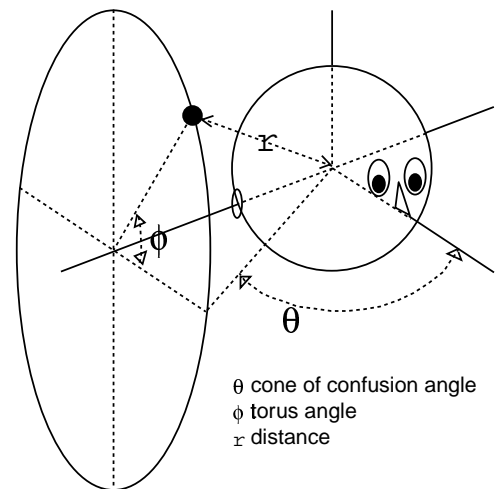


Figure 2: Coordinates for analyzing sound localization.

dimensional localization for sources within a meter of the listener. In these conditions, the level of reverberation is low compared to the level of the direct sound since sources are relatively close to the ears. In addition, since ILDs provide some distance information for nearby sources, we expected that the relatively low levels of reverberation would have little impact on distance perception. Overall, we expected reverberation to have little effect on either directional or distance localization for sources within reach of the listener.

Methodology

Subjects were seated in the center of a reverberant room (reverberation time R_{60} ^a 550 ms). An experimenter positioned a sound source at a random location uniformly distributed within the right hemifield and within one meter of the subject (whose eyes were closed). A broadband noise stimulus with randomly set level was presented from the source (whose position was measured using a Polhemus electromagnetic transmitter/receiver). The subject opened their eyes and positioned a wand in 3-dimensional space at the heard source location (which was measured using a second Polhemus receiver). No feedback was provided to the subjects. Subjects performed multiple sessions over the course of 3-5 days, with each session lasting roughly one hour (approximately 200 trials). (For more detailed descriptions of the experimental procedure, see [13].)

Results

The results shown here focus on how the magnitude of response errors in various spatial dimensions evolved over time. For each source and response, we calculated the source and response location using the coordinate system shown in Figure 2. For each trial of each subject, the source and response coordinates in each of the dimensions were subtracted and the absolute error value was computed. For the distance dimension (r), the absolute error (in meters) was roughly proportional to the source distance (implying that subjects were sensitive to percentage changes in source distance, consistent with a logarithmic perceptual scale for distance). Because of this empirical finding, absolute distance error was normalized by source distance to yield percentage distance error (a metric that is more meaningful than absolute error since humans are logarithmically sensitive to source distance). Since the magnitudes of the errors in all dimensions varied dramatically from trial to trial, a moving average of the error magnitude in each dimension was calculated for each subject across 250 trials.

The resulting error (as a function of time) was averaged across subjects.

Results are shown in Figure 3 (averaged across the 7 subjects in the experiment). In all dimensions, performance improved with practice (the absolute error magnitude generally decreases with trial). In fact, learning continues throughout the experiment in all three spatial dimensions. This is a very interesting result, particularly given that the experiment was performed across multiple (typically five) days. In other words, subjects learned to localize more accurately in the room over time, and this learning lasted from one experimental session to the next. Statistical analysis of the mean error in the initial and final 100 trials showed a statistically significant decrease in error magnitude for all spatial dimensions (paired t-tests; $p < 0.005$). While localization accuracy improved significantly for all spatial dimensions, the magnitude of the improvements was relatively small. Average improvement was roughly 1.5 deg for the cone (θ) and circle (ϕ) angles; percentage distance error decreased by approximately 10% (i.e., % distance error decreased by roughly 3% from its original value of 30%). While improvements were not large, the fact that there was any consistent improvement over hours of practice was surprising.

For comparison, the running averages for Brungart’s anechoic study [13] were computed. There was no evidence for any long-term learning effects in the anechoic data. In addition, there was no statistically significant change in error magnitude with practice. The means and standard deviations of the absolute errors in θ , ϕ , and r (in %) were computed. The average

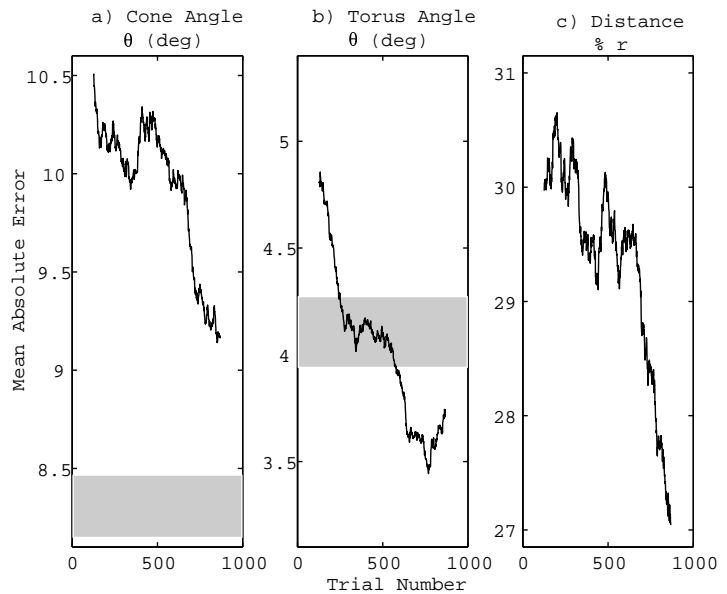


Figure 3: Running average of absolute localization errors (across 7 subjects). Gray stippled areas show range of performance expected from anechoic results [Brungart, 1999 #1328]. Subjects improve with time in all spatial dimensions.

values were 8.3 deg, 4.1 deg, and 51% for θ , ϕ , and r , respectively. From these statistics, the expected range of absolute errors averaged across subjects was calculated for the anechoic data. These ranges are shown in panels a) and b) of Figure 3 for the cone and circle angles (the expected error in percent distance was too large to plot in Figure 3, panel c).

In general, there are large differences in the localization error magnitudes for reverberant and anechoic listening conditions. Even at the end of 1000 trials (approximately 5 hours of listening), average absolute cone angle errors are larger in the reverberant room than in the anechoic condition. Initially, errors in the spectrally-mediated dimension of circle angle are also larger in the reverberant condition. However, within approximately 300 trials, the circle angle error is within the range of errors expected in the anechoic conditions. By 600 trials, performance in the reverberant room actually surpasses performance in the anechoic condition. In contrast with the other dimensions, the average distance error in the reverberant condition is roughly half that in the anechoic condition. Since ILD provides a weak or nonexistent distance cue for sources near the median plane, there are many anechoic trials for which large distance errors are to be expected. What is surprising is that relatively low levels of reverberation provide robust distance information for nearby sources. In other words, even low levels of reverberation provide reliable distance information beyond the binaural ILD cue available in the anechoic condition.

In order to gain further insight into localization errors, we computed the size of the localization error in units of ITD and ILD using the spherical head model (Figure 1). For this analysis, we estimated the ITD (or ILD) that would arise for a sound at the location of the actual source and at location of the subject responses. The absolute magnitude of the response error was then computed across all trials. Figure 4 plots the average absolute response errors in units of ITD (left) and ILD (right).

In the anechoic condition [13], the magnitude of the average error in ITD (53 μ s) is only slightly larger than the just-noticeable difference (JND) in headphone experiments (which ranges between 20–50 μ s, depending on the listener). The average ILD error in the anechoic condition (1.1 dB) is also only slightly larger than reported ILD JNDs (which ranges from 0.5–1.0 dB). Thus, in the anechoic condition, the mean absolute errors in the “binaural” dimensions are close to the absolute limits of the perceptual system. The ITD error is slightly larger in the reverberant condition (70 μ s). The mean ILD error in the reverberant condition is actually as small as JND values reported in the literature. Resolution of stimuli generally degrades with the amount of uncertainty in psychophysical experiments. For instance, in the current experiment, sources could be positioned in any direction in the right hemifield (at any distance up to a meter); there is a large degree of uncertainty. As a result, resolution of the underlying cues in a localization experiment should be worse than in a JND- or discrimination-type experiment (e.g., see [11, 22]). For the anechoic condition, this is the case; both mean ILD and mean ITD error are slightly larger than the range of JNDs reported in the literature (e.g., see [16, 26, 25]). This is also the case for the ITD analysis in the reverberant condition. However, in the reverberant condition, the mean ILD error is as small as the smallest change in ILD that can be detected in a discrimination (“same/different”) task. This ILD result is further proof that in the reverberant condition, ILD cues do not determine performance; instead, the reverberation provides additional information (probably about source distance) that is partially correlated with the ILD “dimension.”

MEASUREMENTS OF REVERBERATION

The results reported above demonstrate that reverberation has a measureable effect on localization performance in both direction and distance. In order to better understand how reverberation might be affecting performance, measurements of the




Figure 4: Average errors in ITD (left) and ILD (right) for localization responses in anechoic and reverberant conditions. Gray areas show range of just noticeable differences. Average ILD errors in the reverberant condition are as small as a JND.

head and room transfer functions were made in a number of subjects in the room used for the experiment. Examples of these measures are shown here for one subject and source location to illustrate how reverberation can distort localization cues.

A maximum-length-sequence technique [31] was used to measure the head-related transfer functions (HRTFs) for sources in the reverberant room. Small microphones were placed at the entrances of the subject's ear canals and the meatus was blocked (i.e., the direction-independent transfer characteristics of the ear canal were not part of the measurements). These raw measures give good estimates of the time-domain impulse response for a source at the location of the acoustic source in the actual room. This raw signal thus describes the total signal reaching the ear of the listener when an impulse is presented in the particular room from the measured location.

Figure 5 shows the impulse response for the right ear for a sound source at $(\theta, \phi, r) = (45 \text{ deg}, 0 \text{ deg}, 1 \text{ m})$. Note that for this spatial configuration, the right ear is the nearer ear. The total signal includes both a large direct sound and many reflections, which die off with time (disappearing by 600 ms). Looking only at the first 20 ms of this impulse response (insert in Figure 5), one can make out the first discrete reflection off the floor in addition to the stereotypical ringing in the direct impulse response (due to the filtering effects of the head and ears of the listener). In order to estimate the impulse response that would occur in an anechoic space, the raw impulse responses from the reverberant room were time windowed to exclude all energy after the "anechoic" HRTF. From these windowed functions, "anechoic" HRTFs were estimated.

Figure 6 compares the magnitude spectrum of the HRTFs for the left and right ears for the sound at $(45 \text{ deg}, 0 \text{ deg}, 1 \text{ m})$ in the reverberant room. Both reverberant and anechoic magnitude spectra were computed. The thick solid line shows the magnitude spectrum of the direct sound (initial 15 ms); the thin dashed line shows that of the total signal.

For this source, which is one meter from the head and off the median plane, the direct sound reaching the near (right) ear is significantly more intense than that reaching the far ear (compare overall levels of top and bottom panels in Figure 6). The reverberation has only small effects on the long-term spectral content of the total signal reaching the near (right) ear (top panel). However, the far ear receives much less direct sound energy and the effect of the reverberation on the magnitude spectrum is more pronounced (bottom panel). While the mean spectral shape is similar for most frequencies, the reverberation adds variability to the spectrum (deviations of 10-20 dB around the spectrum level for the anechoic HRTF). In addition, where the anechoic spectrum shows a prominent notch in the frequency spectrum at 8 kHz, this notch is absent in the reverberant HRTF (the reverberation completely "fills in" this spectral notch). In other words, while the spectral content of the near ear is only




Figure 5: Right ear impulse response for a source in the reverberant room at $(\theta, \phi, r) = (45 \text{ deg}, 0 \text{ deg}, 1 \text{ m})$. Insert shows initial 20 ms in greater detail (note first reflection from floor). Reverberation dies off by 550 ms.




Figure 6: Long-term magnitude spectrum of the right (top) and left (bottom) ear HRTFs; source at $(\theta, \phi, r) = (45 \text{ deg}, 0 \text{ deg}, 1 \text{ m})$. Reverberation degrades the spectrum of the far (left) ear signal.

marginally affected by the reverberation, the spectral cues in the far ear signal are noticeably degraded by the reverberation, even for this relatively near source. In addition, for the far ear, the magnitude spectrum shows random frequency-dependent fluctuations around the true “anechoic” magnitude spectrum.

Of course, binaural cues are more influential on localization performance than spectral cues. Figure 7 presents the interaural phase and interaural level differences in the HRTFs for the source at (45 deg, 0 deg, 1 m).

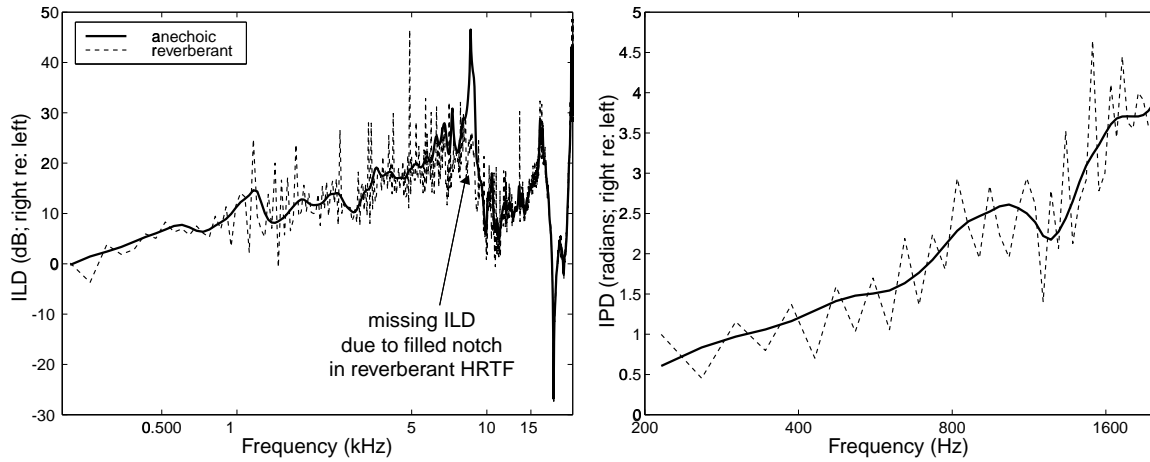


Figure 7: ILD and IPD as a function of frequency for the HRTFs measured in the room. Binaural cues in the reverberant HRTF are essentially noisy versions of those in the anechoic HRTF.

The binaural cues reaching the listener are much less robust to the reverberation than the magnitude spectrum of the signal at the near ear. Given the effect of the signal on the far ear magnitude spectrum, this result is not too surprising; however, the fact that binaural cues are noticeably affected for sources this close to the listener is informative. The reverberation distorts both the ILDs and the interaural phase differences (IPDs) by adding random fluctuations to the interaural differences. That is, on average the ILD and IPD values for the reverberant HRTF are equal to those of the anechoic HRTF (as a function of frequency); however, the reverberation introduces deviations in the interaural cues that are significant. The distortions in the ILD are on the order of 10 dB. Distortions of the IPD are on the order of $\pi/4$ radians. For the ILD results (left panel of Figure 7), the only frequency for which the *mean* ILD is substantially different for the anechoic and reverberant HRTFs is around 8 kHz, where the ILD is large in the anechoic HRTF, but relatively small in the reverberant HRTF. This is the result of the reverberation filling in the notch in the left (far) ear magnitude spectrum at about 8 kHz. For the IPD results, there are no systematic deviations between the mean IPD as a function of frequency for the anechoic and reverberant HRTFs; the reverberant HRTF simply shows much greater variations with frequency. In other words, while the gross features of the ILD and IPD are preserved, the reverberation causes these cues to be less reliable. In perceptual terms, the effect of the reverberation is to add significant noise to the underlying binaural cues.

A CONCEPTUAL MODEL OF THE LEARNING PROCESS

While previous results imply that listeners learn to interpret reverberant distance cues differently in different rooms, the behavioral results presented here are the first to demonstrate *long-term* learning effects that continue across hours (rather than minutes). In addition, the current results demonstrate that directional accuracy improves with long-term experience in a room.

Results are consistent with a localization system that interprets available cues (like ITD, ILD, spectrum, and reverberation level) using some internal template based on past experience. Assume that when presented with a new listening environment, the *a priori* model for mapping spatial cues to location is based on “average” experience. For directional cues like ITD, ILD, and spectrum, the “average” mapping is equal to the ideal mapping in anechoic space, since the reverberant distortions of these cues vary randomly from frequency to frequency and from one environment to another (and thus cancel out, on average). As a result, when listening in anechoic space, *a priori* expectations fit the optimal mapping from localization cue to source position and no long-term learning occurs. This is consistent with the anechoic results, in which no evidence for learning could be identified. In contrast, in a reverberant room, the internal model for direction will not match the ideal mapping because of systematic distortions of spatial cues by reverberation. Because of these distortions, the “default” mapping (based on average expectations) will be modified by experience. Even a small number of trials will help to

recalibrate the perceptual system; however, additional experience may continue to refine localization behavior over hours of experience.

There are a number of processes that could be driving the perceptual calibration that enables listeners to learn to improve localization performance in a room. For instance, it may be that the system simply builds up estimates of the distributions of the spatial cues s/he hears in a room. By comparing the distributions heard with *a priori* expectations for these distributions, the listener can identify any systematic distortions in the cues and adjust their responses to compensate for the systematic distortions. For instance, when in close proximity to a wall, ITD values may tend to be biased toward the wall location. Subjects could adapt to this skewing by recalibrating how ITD maps to azimuthal position. Alternatively, if there is a prominent room resonance at a particular frequency, reverberation will be relatively intense and will cause an increase in variability in spectral and interaural cues around that frequency, making the localization information in these frequencies unreliable. Subjects could learn to weight cues in this frequency range less heavily in order to improve localization accuracy. Long-term adaptation effects similar to this have been observed in a number of previous experiments (e.g., see [35, 36, 34, 21, 20, 30]). A more sophisticated process might build up an explicit internal geometrical model of the reflective surfaces in a room and try to “deconvolve” room effects from the total signals received at the two ears. However, recent evidence suggests that the spatial auditory system cannot even deconvolve directional spectral cues from effects of the source spectrum in an anechoic space [29], so it is unlikely that the calibration of localization performance in a room depends on a detailed geometrical analysis of room characteristics (despite recent models that assume deconvolution is possible in a reverberant room [12]). Further experiments are needed to determine how complex and sophisticated the learning process may be.

IMPLICATIONS FOR SPATIAL DISPLAYS

When designing a spatial auditory display, there are many tradeoffs to consider: whether it is necessary to employ individualized HRTFs, the sampling density of the HRTFs to be stored and used, the sampling rate and length of the HRTFs to be used, whether to include realistic, distance-dependent HRTFs in the simulation, etc. One must take into account the goal of the display device and the resources available in order to determine the optimal approach for a given application (e.g., see [33]).

In the same vein, the decision about whether to include reverberation (and in what form) in an auditory display must be made after considering the specific application. Adding reverberation to a spatial auditory display can be beneficial: distance accuracy is better and the simulation sounds more realistic when reverberation is present. However, there are also costs associated with adding reverberation. Determining whether these costs are outweighed by the benefits can be a complicated problem.

Directional accuracy is worse in the presence of reverberation (e.g., see not only the current results, but [19, 8, 6]). Although the magnitude of the reverberation-induced error in direction can be relatively small, the effect is significant and measurable. Of course, if the listener can be given hours of experience with the simulated reverberant environment, the current results imply they can at least partially overcome the directional-localization distortions caused by reverberation. However, the time required to train subjects is large, and even after 5 hours of practice, performance remains slightly worse than in anechoic conditions (at least in terms of judging cone angle). There are some applications for which the learning observed in the current experiments may prove important; most notably, where highly-trained operators work with three-dimensional data and must localize accurately in both distance and direction (e.g., for air-traffic controllers, fighter pilots, operators of remote vehicles, etc.).

Another factor to be considered is the practical costs associated with including reverberation in a spatial auditory display. The computational demands of simulating realistic reverberation are significant. Even with the power of current computers, it is not possible to simulate the “true” reverberation patterns that would occur in a real room in real time (e.g., see [9, 37]). Thus, simplifications must be made if an interactive spatial auditory display is needed. Of course, the optimal method for generating simplified reverberant distance cues is unknown, since there are only a few studies examining which aspects of the reverberant energy provide distance information to the listener and which aspects are perceptually unimportant (see, for example [7, 9, 10, 1-4]). For instance, while many people have argued that the direct-to-reverberant energy ratio provides the main distance cue (e.g., see [12]), recent experimental results call into question whether this is actually the case ([32]).

Finally, although the focus of this paper is spatial perception, reverberation may decrease performance in other, non-spatial tasks. For instance, speech perception may be adversely affected, at least for highly reverberant simulations. Such effects will be especially pronounced when multiple sources are present simultaneously (and the listener is in a “cocktail party” situation, e.g., see [12]). In such conditions, the addition of reverberation will reduce the correlation of the signals reaching the left and

right ears and interfere with binaural unmasking effects (see Figure 7). In addition, because level differences between the left and right ears will tend to be reduced with the addition of reverberation (see Figure 6), the so-called “better-ear” advantage (which is the most significant factor in spatial unmasking) will also be reduced. In other words, if the goal of the display is to convey source content in addition to source position, the costs of adding reverberation are likely to outweigh potential benefits.

Reverberation can improve the subjective realism of an auditory display as well as distance perception, but depending on the goal of the display, reverberation may cause more harm than good. A careful evaluation of both benefits and drawbacks must be undertaken in order to determine whether adding reverberation to an auditory spatial display is worthwhile.

ACKNOWLEDGMENTS

This work was supported in part by AFOSR grant F49620-98-1-0108 and a research grant from the Whitaker Foundation for Biomedical Research. Scott Santarelli, Norbert Kop o, and Tara Brown gathered the data reported here. The author wishes to thank Douglas Brungart for providing anechoic behavioral data [13].

REFERENCES

1. Bech, S., *Audibility of individual reflections in a complete sound field II*. Journal of the Acoustical Society of America, 1995. **97**: p. 3320.
2. Bech, S., *Spatial aspects of reproduced sound in small rooms*. Journal of the Acoustical Society of America, 1998. **103**(1): p. 434-445.
3. Bech, S., *Timbral aspects of reproduced sound in small rooms. I*. Journal of the Acoustical Society of America, 1998. **97**(3): p. 1717-1726.
4. Bech, S., *Timbral aspects of reproduced sound in small rooms. II*. Journal of the Acoustical Society of America, 1998. **99**(6): p. 3539-3549.
5. Begault, D.R., *Perceptual effects of synthetic reverberation on three-dimensional audio systems*. Journal of the Audio Engineering Society, 1991. **40**: p. 895-904.
6. Begault, D.R., *Perceptual effects of synthetic reverberation on three-dimensional audio systems*. Journal of the Audio Engineering Society, 1992. **40**(11): p. 895-904.
7. Begault, D.R. *Virtual acoustics: Evaluation of psychoacoustic parameters*. in *Third International Workshop on Human Interface Technology IWHIT '97*. 1997. Aizu, Japan.
8. Begault, D.R., E.M. Wenzel, A.S. Lee, and M.R. Anderson. *Direct comparison of the impact of head-tracking, reverberation, and individualized head-related transfer functions on the spatial perception of a virtual speech source*. in *Audio Engineering Society 108th Conference*. 2000. Paris, France.
9. Begault, D.R., E.M. Wenzel, L.L. Tran, and M.R. Anderson. *Octave-band thresholds for Modeled Reverberant Fields*. in *Audio Engineering Society 104th Convention*. 1998. Amsterdam.
10. Bradley, J.S. and G.A. Soulodre, *The influence of late arriving energy on spatial impression*. Journal of the Acoustical Society of America, 1995. **97**(4): p. 2263-2271.
11. Braida, L.D. and N.I. Durlach, *Peripheral and central factors in intensity perception*, in *Auditory Function: Neurobiological Bases of Hearing*, G.M. Edelman, W.E. Gall, and W.M. Cowan, Editors. 1988, John Wiley & Sons: New York. p. 559-583.
12. Bronkhorst, A.W. and T. Houtgast, *Auditory distance perception in rooms*. Nature, 1999. **397**(11 February): p. 517-520.
13. Brungart, D.S. and N.I. Durlach, *Auditory localization of nearby sources II: Localization of a broadband source in the near field*. Journal of the Acoustical Society of America, 1999. **106**(4): p. 1956-1968.
14. Clifton, R.K. and R.L. Freyman, *The precedence effect: Beyond echo suppression*, in *Binaural and Spatial Hearing in Real and Virtual Environments*, R. Gilkey and T. Anderson, Editors. 1997, Erlbaum: New York. p. 233-256.
15. Duda, R.O., *Elevation dependence of the interaural transfer function*, in *Binaural and Spatial Hearing in Real and Virtual Environments*, R. Gilkey and T. Anderson, Editors. 1997, Erlbaum: New York. p. 49-76.

16. Durlach, N.I. and H.S. Colburn, *Binaural phenomena*, in *Handbook of Perception*, E.C. Carterette and M.P. Friedman, Editors. 1978, Academic Press: New York. p. 365-466.
17. Durlach, N.I., A. Rigapulos, X.D. Pang, W.S. Woods, A. Kulkarni, H.S. Colburn, and E.M. Wenzel, *On the externalization of auditory images*. Presence, 1992. **1**: p. 251-257.
18. Freyman, R.L., R.K. Clifton, and R.Y. Litovsky, *Dynamic processes in the precedence effect*. Journal of the Acoustical Society of America, 1991. **90**(2): p. 874-884.
19. Hartmann, W.M., *Localization of sound in rooms*. Journal of the Acoustical Society of America, 1983. **74**: p. 1380-1391.
20. Held, R.M., *Shifts in binaural localization after prolonged exposure to atypical combinations of stimuli*. American Journal of Psychology, 1955. **68**: p. 526-548.
21. Hofman, P.M., J.G.A. Van Riswick, and A.J. Van Opstal, *Relearning sound localization with new ears*. Nature Neuroscience, 1998. **1**(5): p. 417-421.
22. Koehnke, J. and N.I. Durlach, *Range effects in the identification of lateral position*. Journal of the Acoustical Society of America, 1989. **86**: p. 1176-1178.
23. Litovsky, R.Y., H.S. Colburn, W.A. Yost, and S.J. Guzman, *The precedence effect*. Journal of the Acoustical Society of America, 1999. **106**(4): p. 1633-1654.
24. Mershon, D.H., W.L. Ballenger, A.D. Little, P.L. McMurtry, and J.L. Buchanan, *Effects of room reflectance and background noise on perceived auditory distance*. Perception, 1989. **18**: p. 403-416.
25. Mills, A.W., *On the minimum audible angle*. Journal of the Acoustical Society of America, 1958. **30**: p. 237-246.
26. Mills, A.W., *Lateralization of high-frequency tones*. Journal of the Acoustical Society of America, 1960. **32**: p. 132-134.
27. Plenge, G., *On the differences between localization and lateralization*. Journal of the Acoustical Society of America, 1974. **56**: p. 944-951.
28. Rakerd, B. and W.M. Hartmann, *Localization of sound in rooms. III. Onset and duration effects*. Journal of the Acoustical Society of America, 1986. **80**: p. 1695-1706.
29. Rakerd, B., W.M. Hartmann, and T.L. McCaskey, *Identification and localization of sound sources in the median sagittal plane*. Journal of the Acoustical Society of America, 1999. **106**(5): p. 2812-2820.
30. Recanzone, G.H., *Rapidly induced auditory plasticity: The ventriloquism aftereffect*. Proceedings of the National Academy of Sciences, 1998. **95**: p. 869-875.
31. Rife, D.D. and J. Vanderkooy, *Transfer-function measurement with maximum-length sequences*. Journal of the Audio Engineering Society, 1989. **6**: p. 419-444.
32. Santarelli, S., N. Kop o, and B.G. Shinn-Cunningham, *Distance judgements of nearby sources in a reverberant room: Effects of stimulus envelope*. Journal of the Acoustical Society of America, 2000. **107**(5).
33. Shinn-Cunningham, B.G., *Applications of virtual auditory displays*. Proceedings of the 20th International Conference of the IEEE Engineering in Medicine and Biology Society, 1998. **20**(3/6): p. 1105-1108.
34. Shinn-Cunningham, B.G., *Adapting to remapped auditory localization cues: A decision-theory model*. Perception and Psychophysics, 2000. **62**(1): p. 33-47.
35. Shinn-Cunningham, B.G., N.I. Durlach, and R.M. Held, *Adapting to supernormal auditory localization cues I: Bias and resolution*. Journal of the Acoustical Society of America, 1998. **103**(6): p. 3656-3666.
36. Shinn-Cunningham, B.G., N.I. Durlach, and R.M. Held, *Adapting to supernormal auditory localization cues II: Changes in mean response*. Journal of the Acoustical Society of America, 1998. **103**(6): p. 3667-3676.
37. Shinn-Cunningham, B.G., H. Lehnert, G. Kramer, E.M. Wenzel, and N.I. Durlach, *Auditory Displays*, in *Binaural and Spatial Hearing in Real and Virtual Environments*, R. Gilkey and T. Anderson, Editors. 1997, Erlbaum: NY. p. 611-663.
38. Shinn-Cunningham, B.G., S. Santarelli, and N. Kop o, *Tori of confusion: Binaural localization cues for sources within reach of a listener*. Journal of the Acoustical Society of America, in press.