values were 8.3 deg, 4.1 deg, and 51% for , , 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).

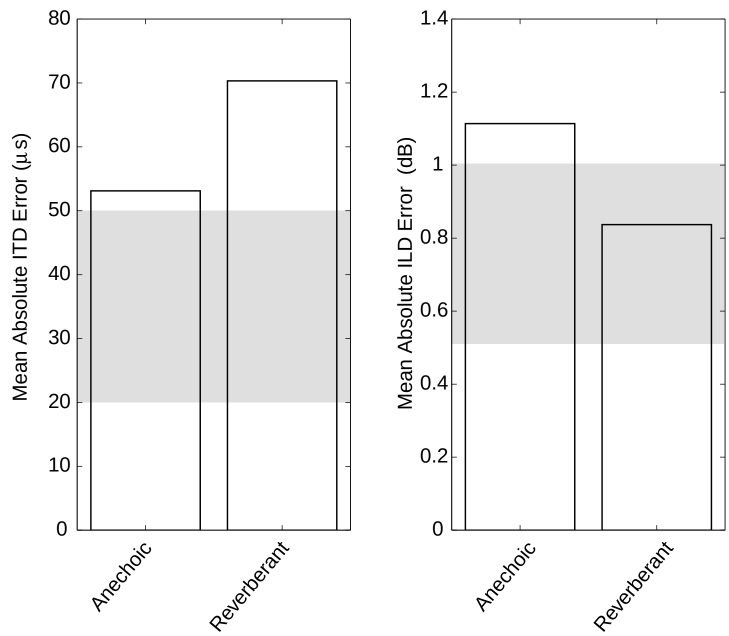

Figure 4: Average errors in ITD (left) and ILD (right) for localization responses in anechoic and reverberant conditions. Gray areas show range of just noticeable differences. Average ILD errors in the reverberant condition are as small as a JND.

In the anechoic condition [13], the magnitude of the average error in ITD (53 s) is only slightly larger than the just-noticeable difference (JND) in headphone experiments (which ranges between 20–50 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 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 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 (, , r) = (45 deg, 0 deg, 1 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

domat”j(ALl
(9by hz(]:”j(([T“9cowz(T(”]:”j(bT(”]:”j(exTj9laT(”]:”j(s hT(”]:”j(e9gyT(”]:”j(sfter%iT:”j(9domat”j(“anshoFz(I]T”mat”j(z((:.mat”j(Fz(Tr r

spectral content of the near ear is only 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 reveberant 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 /4 radians. For the ILD results (left panel of Figure 7), the only frequency for which the *m* 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 - *m* 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 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, 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 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

17. Durlach, N.I., A. Rigapulos, X.D. Pang, W.S. Woods, A. Kulkarni, H.S. Colburn, and E.M. Wenzel,