

NEAR-FIELD VIRTUAL AUDIO DISPLAYS

Douglas S. Brungart
Human Effectiveness Directorate
Air Force Research Laboratory
Wright-Patterson AFB, Ohio

Abstract

Although virtual audio displays are capable of realistically simulating distant sound sources, they are not yet able to reproduce the auditory cues associated with sound sources near the head. Researchers have long recognized that the binaural difference cues that dominate auditory localization are independent of distance beyond 1 m but change systematically with distance when source the approaches within 1 m of the listener's head. More recent research has shown that listeners are able to use these binaural cues to determine the distances of nearby sound sources. However, technical challenges in the collection and processing of near-field Head-Related Transfer Functions (HRTFs) have thus far prevented the construction of a near-field audio display. This paper summarizes the current state of research in the localization of nearby sound sources, and outlines the technical challenges involved in the creation of a near-field virtual audio display. The potential applications of near-field displays in immersive virtual environments and multimodal interfaces are also discussed.

Introduction

Over the past 15 years, virtual audio displays have matured from crude laboratory prototypes into sophisticated head-coupled systems capable of manipulating the perceived azimuth and, to a lesser extent, the perceived elevation of virtual sound sources. However, the current generation of virtual audio displays cannot generate realistic virtual sounds in the "near field" within 1 m of

the head. This paper briefly summarizes recent research in near-field sound localization, analyzes the technical challenges that have prevented the implementation of realistic near-field virtual displays, and discusses the possible applications of near-field virtual audio cues in immersive virtual environments and multimodal interfaces.

Auditory Localization

One of the impressive capabilities of the human auditory system is the ability to localize the direction of a sound source. Lord Rayleigh (1907) first recognized that the lateral position of a sound source can be determined directly from two interaural difference cues: the difference in the arrival time of the sound at the left and right ears, known as the Interaural Time Delay (ITD); and the difference in the loudness of the sound at the left and right ears, known as the Interaural Level Difference (ILD). Since the propagation time of a sound wave is shorter for the ear closer to the source (the ipsilateral ear) than for the ear farther from the source (the contralateral ear), the sound at the more distant ear is delayed relative to sound arriving at the closer ear. This Interaural Time Delay has been shown to dominate directional sound localization at low frequencies (Wightman & Kistler, 1992). At higher frequencies, where the wavelength of the sound is small relative to the dimensions of the head, interference by the head attenuates the sound reaching the farther ear. This head-shadowing effect causes an Interaural Level Difference between the direct sound reaching the ipsilateral ear and the diffracted sound at the contralateral ear.

The ITD and ILD are potent localization cues, allowing the discrimination of changes in azimuth as small as 1° for sounds directly in front of the




Fig. 1. Head-Related Transfer Functions (HRTFs) for a sound source directly to the left of the listener (90°) measured at three distances (1.0 m, 0.25 m, and 0.12 m) from the center of the head. The HRTFs were measured with a Knowles Electronics Manikin for Acoustic Research (KEMAR) and an acoustic point source, and have been normalized by the free-field sound pressure at the center of the head to eliminate the overall effects of distance (Brungart and Rabinowitz, 1999). Note that, as the sound source approaches the head, the magnitude of the HRTF increases rapidly at the ipsilateral ear and decreases slightly at the contralateral ear. This results in a large increase in ILD (the difference between the solid and dotted lines) across all frequencies for nearby sources.

listener (Mills, 1958). However, binaural cues alone cannot fully explain directional auditory localization. In order to determine the elevation of sound sources and to distinguish between sounds in front of or behind the head, listeners rely on the direction-dependent filtering properties of the outer ear or pinna (Musicant & Butler, 1984; Roffler & Butler, 1968). The pattern of folds in the pinna produces a complex pattern of peaks and notches in the high-frequency spectrum of sound that changes systematically with the direction of the source. Since pinna geometry varies across listeners, it appears that listeners are able to adapt to the high-frequency responses of their own ears through experience and learning. For a more complete description of auditory localization, see the review article by Middlebrooks and Green (1991).

The HRTF

The acoustic cues available for sound localization are completely characterized by the linear transformations of the acoustic signal propagating from the sound source to the listener's ears. These transformations are known as the Head-Related Transfer Functions, or HRTFs. In order to measure the HRTF, sound generated by a loudspeaker at a known location in space is recorded by microphones positioned in the ears of a listener. The magnitude and phase of the HRTF are then determined by dividing the complex spectra of the

sounds reaching the ears by the complex spectrum of the sound generated by the loudspeaker. The HRTF includes the effects of diffraction by the head and torso as well as the direction-dependent filtering of the pinnae, and can be used to measure changes in localization cues such as the ILD and ITD as a function of source location.

Virtual Audio Displays

Since all of the available audio localization cues are included in the HRTFs, the HRTFs can be used to add spatial attributes to non-spatial audio signals. This is accomplished by processing the input signal with two digital filters that match the HRTFs of the left and right ears for the desired sound location, and playing the resulting signals through headphones. If care is taken to correct for the frequency response of the headphones, this type of processing can produce a convincing illusion of sound source at the desired location in space (Wightman & Kistler, 1989b). When a head-tracking device is added to the system, an even more compelling simulation can be achieved by continuously updating the position of the synthetic source in response to the head motions of the listener. This type of interactive audio synthesizer is called a virtual audio display (Wenzel, 1991). Although the current generation of audio displays does an excellent job of simulating the azimuth of a sound and an adequate job of simu-

lating elevation, they are unable to account for important changes that occur in the HRTF when the sound source is near the head. These changes, described in the next section, are vital for producing accurate auditory simulations of nearby sounds.

Near-Field Auditory Localization

The HRTF in the Near Field

For nearly a century, researchers have recognized that there are fundamental differences between the auditory localization of distant sounds and the localization of sounds near the head. The earliest studies in this area examined the distance-dependent changes in the near-field HRTF with a mathematical rigid-sphere model of the head (Hartley & Fry, 1921; Stewart, 1911). Recently there has been renewed interest in near-field localization. The rigid-sphere model has been updated and compared to acoustic measurements on a rigid sphere (Brungart & Rabinowitz, 1996; Duda & Martens, 1998), and an extensive set of near-field HRTFs have been measured with a Knowles Electronics Manikin for Acoustic Research (KEMAR) (Brungart & Rabinowitz, 1999).

These studies have shown that, when the source is within 1 m of the listener's head, the HRTF changes substantially with distance. The largest changes occur in the Interaural Level Difference (ILD), which increases dramatically as a lateral source approaches the head (Figs. 1 and 2). Two factors contribute to the increase in ILD for nearby sources. The first factor is an increase in head-shadowing effect for objects near the head. The increased head shadow results in greater attenuation at the contralateral ear and a larger ILD at high frequencies. The second factor follows directly from the inverse relationship between amplitude and distance for a freely propagating sound wave. As distance decreases, the amplitude of the sound increases more rapidly at the ipsilateral ear than at the contralateral ear, resulting in an increased ILD. This proximity effect is especially important because, unlike head shadowing, it increases the ILD across all frequencies. This results in substantially larger low-frequency ILDs in the near field than are ever found in the far field. At 500 Hz, for example, the near-field ILD can exceed 15 dB (Fig. 2, top panel), while




Fig. 2. ILD and ITD as a function of distance at three source directions: 90° (directly to the left of the listener), 45°, and 0° (directly in front of the listener). As in Fig. 1, the HRTFs were measured with the KEMAR manikin (Brungart and Rabinowitz, 1999). Note that the ILD increases much more rapidly with decreasing distance than the ITD.

the far-field ILD never exceeds 5-6 dB. The presence of large ILDs at low frequencies is one of the distinguishing characteristics of the near-field HRTF.

Other localization cues change less dramatically in the near field. The ITD is roughly independent of distance even for the closest sources (Fig. 2, bottom panel). Although there is a slight increase in the ITD at the closest distances, most of this increase occurs at lateral locations where listeners are relatively insensitive to changes in ITD (Hershkowitz & Durlach, 1969). Pinna-based localization cues are also similar in the near and far fields. The high-frequency patterns in the




Fig. 3. Directional auditory localization accuracy for a nearby, free-field sound source. The left panel shows the overall, great-circle angular error of listeners placing a response wand at the perceived location of a near-field sound source. The right panel shows the percentage of front-back confusions experienced by the listeners. Both show that directional localization accuracy decreases slightly when the sound source approaches within a few centimeters of the head (Brungart et al., 1999).

near-field HRTF are more sensitive to changes in the elevation than to changes in distance (Brungart & Rabinowitz, 1999). In the horizontal plane, it appears that the high-frequency pinna cues in the HRTF change with the angle of the source relative to the ear rather than the angle relative to the head, resulting in an auditory parallax effect for nearby sources (Brungart, 1999b). Although these changes in ITD and pinna cues certainly have some influence on sound perception in the near field, the broadband increase in ILD is clearly the most perceptually relevant near-field cue.

Localization Performance in the Near Field

Directional auditory localization is generally similar in the near and far fields. In an experiment measuring the ability of listeners to place a position sensor at the location of a nearby sound source (Brungart, Durlach, & Rabinowitz, 1999), the overall great-circle angular error (which includes azimuth and elevation) between the stimulus location and the response location was 16.9°. When differences in methodology are considered, this is comparable to the results reported for far-field sound localization (Wightman & Kistler, 1989a). However, overall angular error was slightly larger for sound sources very near the head (< 25 cm) than for more distant sources (Fig. 3, left panel). This results primarily from an increase in response bias rather than an increase in response variability: when bias was removed,

the RMS errors in azimuth and elevation were roughly independent of distance. There was also a significantly larger number of front-back reversals (where the symmetry of the ITD and ILD caused the listener to respond at the mirror image of the actual source location) when the sound source was closer than 0.5 m than when it was farther than 0.5 m (Fig 3., right panel). Overall, there appears to be a slight degradation in directional localization accuracy when the sound source approaches within a few centimeters of the head.

The more interesting differences between near- and far-field auditory localization relate to the perception of distance. In the far field, listeners are generally much less adept at distance localization than at directional localization. Several researchers (Coleman, 1962; Gardner, 1968) have found that listeners are unable to accurately determine the absolute distance of an unfamiliar sound source in a free field (although they can of course judge relative distance based on changes in loudness). Only in reverberant environments, where the ratio of direct to indirect energy provides an important absolute distance cue, are listeners able to make distance judgments that are correlated with source distance (Mershon & Bowers, 1979).

The inability of listeners to determine the distance of a sound source in the free field is not surprising, since there is little or no distance dependence in the HRTF beyond 1 m. In the near field, however, there are important binaural distance cues that listeners can use to judge the distance of a sound source even when the amplitude of the source is roved in an anechoic listening environment. Since the ILD increases as distance decreases while the ITD remains roughly constant, it should be possible to estimate the distance of a nearby sound source by comparing the magnitudes of the ITD and ILD. The results of psychoacoustic experiments support this hypothesis (Fig. 4). With a random-amplitude broadband noise stimulus, the listeners' distance judgments were highly correlated with source distance for lateral sources ($r > 0.85$). This suggests that the distance-dependent changes in the near-field HRTF allow listeners to make substantially more accurate free-field distance judgments in the near field than in the far field (Brungart et al., 1999). In the median plane, however, the distance judgments were only moderately correlated with source distance




Fig. 4. Auditory distance perception in the near field. These bars represent the correlation coefficients between the log of the stimulus distance and the log of the response distance for trials directly to the right of the listener (90°) and trials directly in front of the listener (0°). The stimulus in the high-pass condition was high-pass filtered above 3 kHz and the stimulus in the low-pass condition was low-pass filtered below 3 kHz. In all conditions, performance was substantially better for lateral sources than for medial sources, suggesting the importance of binaural cues in near-field distance perception. The results also indicate that low frequencies dominate distance perception in the near field (Brungart, 1999).

($r \approx 0.35$). Since the binaural difference cues are greatest for lateral sources and negligible for medial sources, this pattern of auditory distance perception strongly suggests that listeners are relying on distance-dependent changes in the ILD to make auditory distance judgments in the near field.

There is also strong evidence that low frequencies dominate near-field distance perception in the free field (Fig. 4). When the stimulus was low-pass filtered below 3 kHz, performance was essentially identical to the broadband case (Brungart, 1999a). However, when the stimulus was high-pass filtered above 3 kHz, auditory distance perception was severely degraded for both lateral and medial sources. Thus, it appears that listeners are able to use low-frequency binaural cues to make reasonably accurate absolute distance judgments in the near field.

A final note should be made about near-field auditory distance perception in reverberant environments. As the sound source approaches the head in an echoic room, the ratio of direct-to-reverberant energy increases substantially more

rapidly in the near field than in the far field. For lateral sources, it also increases more rapidly at the ipsilateral ear than at the contralateral ear. Recent psychoacoustic experiments have demonstrated that these reverberation cues play an important role in the localization of sounds in rooms (Santarelli, Kopco, Shinn-Cunningham, & Brungart, 1999), and can be as potent as the binaural near-field distance cues when the listener is familiar with the listening environment. However, near-field reverberation cues also have all the drawbacks of far-field reverberation cues from the standpoint of implementation in a virtual audio display. Therefore we have generally limited our discussion of near-field audio displays to the simulation of anechoic listening environments.

Technical Challenges

It is clear that the ideal virtual audio display would be able to accurately simulate sounds near the head as well as more distant sounds. However, there are some important technical challenges that must be overcome in order to produce a truly effective near-field virtual display. One of the most important challenges is the actual collection of the near-field HRTFs. Since near-field displays require HRTFs at multiple distances, there is a substantial increase in the complexity of the measurement process. Far-field virtual audio displays typically use a library of 100-750 directional HRTFs measured at a single distance. A near-field virtual display would require the same number of HRTFs for at least four different distances, quadrupling the number of measurements required. In addition, a three degree-of-freedom placement system is required to move the speaker around the head. The boom and gimbal systems currently used to collect HRTFs in most spatial audio laboratories do not allow variations in the distance of the source. It is also necessary to use a specially designed acoustic point source instead of a traditional loudspeaker in order to make accurate HRTF measurements near the head. Although it is possible to build an apparatus to overcome these problems, most current HRTF measurement systems are unable to measure near-field HRTFs without substantial modifications.

Controlling the location of the source relative to the head is an even greater challenge in the near field. In the far field, only the orientation

of the head must be tightly controlled during the HRTF measurements. Translations of a few centimeters in the position of the head have little impact on the measurements. In the near field, however, a change in location of only a few centimeters can substantially change the angular position of the source. Therefore it is necessary to completely immobilize the head of the listener during the measurement procedure. Since the near-field measurements are likely to take four times as long as traditional far-field measurements, this can be extremely uncomfortable for human listeners. It is also important to use a consistent coordinate system that defines locations in the near field based on the anatomical features of the head. Brungart, Rabinowitz and Durlach (2000) have suggested using the locations of the ears and the tip of the nose to define this coordinate system. Finally, it is necessary to accurately measure the position of the source relative to the head with this coordinate system. This can be complicated because the location of the center of the head changes from listener to listener.

Because of these difficulties, complete sets of near-field HRTFs have so far been measured only on an acoustic manikin (Brungart & Rabinowitz, 1999). Manikins are well-suited to near-field HRTF measurements for three reasons: 1) they are immobile and allow complete control over the position of the source relative to the head; 2) they can easily be rotated in place to collect a full set of azimuth data from a stationary source; 3) there are no time limitations or comfort issues to deal with when a manikin is used for the measurement. Furthermore, since auditory distance perception in the near field appears to rely primarily on low frequencies where individual differences in the HRTF are minimal, HRTF measurements on a manikin should provide adequate near-field distance information across a wide range of listeners. However, manikins are not ideal for the implementation of veridical audio displays. Accurate reproduction of the individual HRTF corresponding to the unique shape of the listener's ears is known to be important for generating realistic, externalized audio cues. New techniques must be developed before individualized near-field HRTFs can be collected efficiently and used in a high-fidelity virtual audio display.

A final technical challenge that has not yet

been adequately addressed is the interpolation of HRTFs in the near field. Numerous researchers have examined the problem of using HRTFs measured at two adjacent locations to generate an HRTF at a location "in between." It is unclear how well these directional interpolation algorithms can be adapted to interpolating distance in the HRTF, or what resolution in distance is required to successfully interpolate HRTFs in the near field. Further research is needed in this area.

Applications of Near-Field Virtual Audio Displays

Once these implementation issues are overcome, there is a broad range of possible uses for near-field virtual displays. This section discusses the advantages of a near-field display in a number of different applications.

Immersive Multimodal Virtual Environments

The most obvious application of a near-field virtual audio display is the simulation of nearby sound sources in immersive virtual environments. Since the technological requirements are quite different for visual, audio, and haptic displays, there has been a tendency for researchers to specialize and produce immersive displays combining state-of-the-art technology in one modality with less advanced displays in the other modalities. But as virtual environments have grown more sophisticated, advanced visual, audio, and haptic displays are being combined to create a compelling immersive experience. One of the key requirements for adding a sense of presence to these systems is allowing the operator to interact with virtual objects by grabbing and moving them. Thus, the region of space within "arm's reach" is increasingly important in the creation of realistic environments. In an environment where listeners are able to grab virtual sound sources and move them around the head, realistic near-field audio cues are extremely important.

Non-Immersive Multimodal Displays

In many cases, it is desirable to add a virtual audio display to a pre-existing work environment, such as an aircraft cockpit or computer station, and produce audio icons that coincide with the locations of actual physical objects. For example, an engine-failure tone might be associated

with an oil-pressure gauge on a control panel, or a tone indicating the receipt of a new email message might emanate from the mailbox icon on a computer screen. In these situations, where the physical objects are located close to the operator, near-field auditory cues should produce a far more compelling fusion between the audio event and the associated physical object than far-field cues.

Auditory Distance Cueing

One of the greatest shortcomings of current virtual audio displays is their inability to reliably provide robust distance information. The two audio cues that are most widely used to manipulate perceived distance in virtual displays, loudness and reverberation, both have significant disadvantages. Loudness is adequate to convey changes in the distance of the source (Coleman, 1963), but provides little or no information about absolute distance unless the listener has *a priori* information about the intensity of the source. Furthermore, a distance cue based on loudness can result in an auditory signal that is either inaudible when it is far away or uncomfortably loud when it is close. Reverberation cues can provide some absolute distance information, but they are dependent on the specific details of the room and may be less effective if there is a mismatch between the listening environment and the simulated room. Reverberation can also degrade localization ability and speech intelligibility. And reverberation cues are difficult to synthesize veridically, requiring the real-time synthesis of tens or hundreds of reflected signals.

Near-field auditory displays provide a third option for simulating distance in a virtual display. Since near-field distance cues in the free field are based primarily on differences between the signals reaching the two ears rather than the absolute characteristics of the source, they can provide information about the absolute distance of the source without any *a priori* knowledge about the source or the environment. They are also intuitive, and able to convey distance information with little or no training. Thus, it appears that near-field cues are well suited to providing distance information in a virtual audio display.

However, the range of distances they can veridically simulate is limited to the immediate vicinity of the listener. Although there are many

applications where this is appropriate, in some applications it is desirable to produce distance cues for objects tens or hundreds of meters away. This limitation can be overcome by mapping relatively distant locations into the near field and appropriately training the listener. It remains to be seen how well listeners will be able to adapt to this remapping.

Stream Segregation

Near-field cues may also be useful when it is necessary to differentiate between two or more virtual sounds that are co-located in space. In many virtual audio applications, the listener hears multiple simultaneous sounds. If the audio display is using these sounds to provide spatial information to the listener, on some occasions more than one source will be located at the same angle, and they will interfere with one another. With near-field audio cues, the sounds can be placed at different distances, allowing the listener to intuitively segregate the two sounds without losing the spatial information associated with each.

Attention Selection

Near-field auditory displays may also be helpful in the design of interfaces that can selectively tune the operator's attention to nearby and far away events. In many applications, there is a clear differentiation between "peripersonal" objects and events in the immediate vicinity of the operator and "extrapersonal" objects and events remote from the location of the operator (Previc, 1998). In an aircraft cockpit, for example, objects and events related to internal systems on the aircraft, such as an instrument on a control panel or a mechanical warning, would be considered peripersonal because they are stationary with respect to the reference frame of the pilot. Objects and events originating outside the aircraft, such as a missile warning or a wingman communication, would be considered extrapersonal. The use of near-field HRTFs for audio icons related to peripersonal events and far-field HRTFs for icons related to extrapersonal events would provide an intuitive way to direct the operator's attention to internal or external events.

Control of Perceived Urgency

In some cases, it is useful to manipulate the perceived urgency of a virtual sound. Although little research has been done in this area, it is likely that sounds very near the head will naturally be perceived as more urgent than far-field sounds. Near-field sounds may also be more likely to induce a startle reflex, such as the one that occurs when someone sneaks up behind a person and surprises them with a sudden whisper in the ear.

Speech Applications

There are also some important applications of near-field virtual audio cues in speech communication systems. It is known that listeners are better able to monitor and respond to multiple simultaneous communication channels when the channels are spatially separated in direction. It is likely that some improvement can also be obtained by spatially separating the channels in distance inside the near field. This would allow listeners to realize an improvement in performance even if both talkers are located in the same direction.

Near-field cues would also increase the realism and intimacy of virtual conversations. There are important cultural differences between conversational speech at a distance of 1 m and speech whispered directly into the ear. In a sophisticated shared telepresence teleconference, for example, it is easy to imagine a situation where a person would want to whisper something confidential into another person's ear. Many romantic encounters also involve whispered speech at close range. Near-field cues are necessary to create veridical simulations of these intimate conversations.

Entertainment

As has been the case with many advancements in virtual environment technology, it is likely that near-field audio cues will find their first widespread commercial applications in the realm of entertainment and games. This is not surprising, because near-field sounds are among the most compelling, realistic, and impressive virtual sounds around. In binaural recordings, for example, sounds near the head are often especially convincing. A binaural recording of scissors near the head in a "virtual hair-cut" can be a particularly eerie experience. Real-world sounds in the

near field also tend to provoke strong responses. Few sounds are as startling as a sudden whisper in the ear, or as irritating as a fly buzzing around the head. The possible uses of near-field cues in entertainment environments are limited only by the creativity of the audio designer.

Conclusion

To this point, technical limitations on the fidelity of synthetic environments have largely overshadowed the inability of virtual audio displays to accurately simulate nearby sounds. However, as multisensory virtual environments become more sophisticated, near-field auditory cues will become increasingly important to the creation of a compelling virtual world. Despite the engineering challenges involved, it is only a matter of time before veridical near-field audio displays become available. When this happens, near-field cues are certain to be a valuable addition to every virtual sound designer's toolkit.

References

- Brungart, D. (1999a). Auditory localization of nearby sources. III: Stimulus Effects. *Journal of the Acoustical Society of America*, 106, 3589–3602.
- Brungart, D. (1999b). Auditory Parallax Effects in the HRTF for Nearby Sources. In *Proceedings of 1999 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pp. 171–174.
- Brungart, D., Durlach, N., & Rabinowitz, W. (1999). Auditory localization of nearby sources. II: Localization of a Broadband Source. *Journal of the Acoustical Society of America*, 106, 1956–1968.
- Brungart, D. & Rabinowitz, W. (1996). Auditory localization in the near-field. In *Proceedings of the Third International Conference on Auditory Display*. Santa Fe Institute.
- Brungart, D. & Rabinowitz, W. (1999). Auditory localization of nearby sources. I: Head-related transfer functions. *Journal of the Acoustical Society of America*, 106, 1465–1479.
- Brungart, D., Rabinowitz, W., & Durlach, N. (2000). Evaluation of Response Methods for the Localization of Nearby Objects. *Perception and Psychophysics*, 62, 48–65.
- Coleman, P. (1962). Failure to localize the source distance of an unfamiliar sound. *Journal of the Acoustical Society of America*, 34, 345–346.

Coleman, P. (1963). An analysis of cues to auditory depth perception in free space. *Psychological Bulletin*, 60, 302–315.

Duda, R. & Martens, W. (1998). Range dependence of the response of a spherical head model. *Journal of the Acoustical Society of America*, 104, 3048–3058.

Gardner, M. B. (1968). Proximity image effect in sound localization. *Journal of the Acoustical Society of America*, 43, 163.

Hartley, R. & Fry, T. (1921). The binaural location of pure tones. *Physical Review*, 18, 431–442.

Hershkowitz, R. & Durlach, N. (1969). Interaural time and amplitude JNDs for a 500-Hz tone. *Journal of the Acoustical Society of America*, 46, 1464–1467.

Mershon, D. & Bowers, J. (1979). Absolute and relative cues for the auditory perception of egocentric distance. *Perception*, 8, 311–322.

Middlebrooks, J. & Green, D. (1991). Sound localization by human listeners. *Annual Review of Psychology*, 42, 135–139.

Mills, A. (1958). On the minimum audible angle. *Journal of the Acoustical Society of America*, 30, 237–246.

Musicant, A. D. & Butler, R. A. (1984). The psychophysical basis of monaural localization. *Hearing Research*, 14, 185–190.

Previc, F. (1998). The neurophysiology of 3-D space. *Psychological Bulletin*, 124, 123–164.

Rayleigh, L. (1907). On our perception of sound direction. *Philosophy Magazine*, 13, 214–232.

Roffler, S. & Butler, R. (1968). Factors that influence the localization of sound in the vertical plane. *Journal of the Acoustical Society of America*, 43, 1255–1259.

Santarelli, S., Kopco, N., Shinn-Cunningham, B., & Brungart, D. (1999). Near-field localization in echoic rooms. *Journal of the Acoustical Society of America*, 105, 1024 (A).

Stewart, G. (1911). The acoustic shadow of a rigid sphere with certain applications in architectural acoustics and audition. *Physical Review*, 33, 467–479.

Wenzel, E. (1991). Localization in virtual acoustic displays. *Presence*, 1, 80–107.

Wightman, F. & Kistler, D. (1989a). Headphone simulation of free-field listening. I: Stimulus synthesis. *Journal of the Acoustical Society of America*, 85, 858–877.

Wightman, F. & Kistler, D. (1989b). Headphone simulation of free-field listening. II: Psychological validation. *Journal of the Acoustical Society of America*, 85, 868–878.

Wightman, F. & Kistler, D. (1992). The dominant role of low-frequency interaural time differences in sound localization. *Journal of the Acoustical Society of America*, 91, 1648–1660.

Author's Biography

Douglas S. Brungart is currently a research electrical engineer in the Air Force Research Laboratory's Human Effectiveness Directorate. He received his B.S. degree in Computer Engineering from Wright State University in 1993, his S.M. in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1994, and his Ph.D. in Electrical Engineering and Computer Science from MIT in 1998. His current research focuses on auditory distance cueing for virtual audio displays.