Spatial Audio Displays for Improving Safety and Enhancing Situation Awareness in General Aviation Environments

Brian D. Simpson, Douglas S. Brungart, Robert H. Gilkey, and Richard L. McKinley
Air Force Research Laboratory, AFRL/HECB, Bldg 441, 2610 Seventh Street,
Wright-Patterson AFB, OH 45433-7901 Tel: (937) 255-3660 ext 462 Fax: (937) 656-7680
Wright State University, Department of Psychology, 335 Fawcett Hall,
Dayton, OH 45435 Tel : (937) 775-3337 Fax : (937) 775-3347

brian.simpson@wpafb.af.mil

ABSTRACT

Most current displays in general aviation (GA) environments employ, at best, relatively simple audio displays that do not fully exploit the operator's sensory processing capabilities, and thus may not allow pilots to take full advantage of the available information. This paper describes a series of four experiments conducted in an actual flight environment at the USAF Test Pilot school that evaluate the utility of spatial auditory displays as a navigation aid and an attitude indicator. Performance was measured in tasks that required pilots to fly in the direction of a spatial audio "navigation beacon," and use an auditory artificial horizon display to detect changes in attitude and maintain straight and level flight when no visual cues were available. The results from these studies indicate that spatial audio displays can effectively be used by pilots for both navigation and attitude monitoring. In addition, because the stimulus employed in the audio horizon display was self-selected music, pilot acceptance was high and it was possible to provide ongoing information to the pilot without the risk of annoyance. Although performance with standard display configurations remains superior, spatial audio displays may be used to supplement these existing displays, and to provide cues when critical information is either obscured or missing, such as in low-visibility conditions or when the pilot's attention is focused on another task. Thus it appears that spatial audio may be used to support situation awareness and improve overall safety in the general aviation environment.

1.0 INTRODUCTION

Aviation professionals have been considering the use of spatialized auditory displays (3D audio) in military cockpits for target acquisition, threat avoidance, and the separation of multiple talkers [1,2,3,4] and in general aviation (GA) cockpits for taxiway navigation and collision avoidance [5,6,7,8] for many years. However, despite the long-term interest in this technology, nearly all studies have been in conducted in simulator environments, and the few flight tests that have been published have yielded primarily qualitative data [3]. This paper reports results of a flight test of spatial auditory display concepts that have been developed to enhance pilot situation awareness (SA) and potentially improve safety in GA flight environments.

The accident rate in GA environments is extremely high relative to the rates in scheduled commercial air-carrier aircraft, and, in most cases, pilot error is cited as a primary or contributing factor. This suggests the need for new display technologies that increase the pilot's knowledge of the condition of the aircraft and the status of the surrounding airspace (situation awareness, SA). On the other hand, simply providing additional visual instrumentation to a pilot who is only trained for VFR operations may do little to increase SA. In this

Simpson, B.D.; Brungart, D.S.; Gilkey, R.H.; McKinley, R.L. (2005) Spatial Audio Displays for Improving Safety and Enhancing Situation Awareness in General Aviation Environments. In *New Directions for Improving Audio Effectiveness* (pp. 26-1 – 26-16). Meeting Proceedings RTO-MP-HFM-123, Paper 26. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int/abstracts.aps.

context, the use of spatialized auditory displays is particularly appealing. The allure of 3D audio is based on at least three facts: 1) The visual system is already overloaded. Pilots must monitor, in addition to the out-the-window view, a variety of visual flight instruments which report the condition and location of the plane and status of the airspace, and further must consult charts and other printed material in certain situations. Auditory displays have the potential to add new information without further straining visual resources or to provide redundant information to reinforce visual displays. 2) Hearing is a 4π steradian sense. That is, the auditory system can respond to events (e.g., traffic) above, behind, and below the pilot (i.e., in locations that are outside the visual field of view). 3) Hearing is a 24-hr sense. That is, the auditory system is always "on," even in sleep, and thus has the potential to alert a pilot even when distracted by other tasks. Indeed, the auditory system is designed to respond to even subtle changes in the auditory environment and direct visual attention to the source of the new event. The task of the display designer is to exploit this natural potential of the auditory system in a manor that does not distract or annoy the pilot under normal flight conditions.

This paper reports the results of flight tests of two displays to help pilots control an aircraft with limited visual information: an auditory navigation aid and an auditory attitude indicator. Both displays provide usable information about the orientation of the aircraft in a way that allows them to be consciously used when desired and to attract attention when needed, without acting as an ongoing annoyance within the cockpit.

2.0 GENERAL METHODS

2.1 Subjects

Three members of Class 04A at the United States Air Force Test Pilot School served as subject pilots. The pilots had between 8 and 15 years of flight experience in various aircraft.

2.2 Apparatus

A twin-engine turboprop passenger and cargo aircraft (C-12C Huron test aircraft) was used as the flight environment for data collection. A pallet and rack containing the hardware necessary for stimulus generation, presentation, and data collection were securely mounted approximately in the middle of the aircraft, where the experimenter had sufficient access to monitor the status of the equipment and run the experimental software.

The primary component of the auditory display system was a PC-based laptop computer that functioned as the experimental control computer and the audio server. The server used the Sound Lab (SLAB) API developed at NASA-AMES [9]. SLAB enabled the system to generate spatialized auditory stimuli such that sounds presented over headphones appeared as if they came from actual locations in space external to the pilot's head. This spatialization process is accomplished by modifying an arbitrary signal to contain the direction-dependent acoustic transformations (the head-related transfer functions, HRTFs) that a sound undergoes as it travels from a source to the eardrum. An audio mixer (4-channel Allen & Heath Phoenix) mixed the spatial audio cues with the aircraft intercom system, and the sounds were sent to a stereo headset (Bose AH-TS ANR headset modified to transduce a stereo signal) worn by the subject pilot.

The spatialization of auditory stimuli was updated in real time in one of two ways, depending on the experimental condition: 1) with respect to the orientation of the pilot's head or 2) with respect to the position and attitude of the aircraft. Spatialization with respect to the pilot's head was achieved using an inertial headtracking system consisting of a microelectromechanical system (MEMS) inertial measurement unit (IMU), a Global Positioning System (GPS) receiver (Garmin Industries), and a PC-based laptop computer. Spatialization with respect to the aircraft was achieved using a Time, Space, and Position Information (TSPI) GPS-Aided Inertial Navigation Reference (GAINR) system. This system served as the primary source of truth data regarding the aircraft's position and attitude. The C-12C aircraft was also instrumented with a separate data acquisition system.

2.3 Procedure

On each day of data collection, a pre-test briefing was given in which weather conditions, traffic, and any issues regarding the experimental system were discussed. If a decision was made to proceed with data collection on that day, a full test of experimental instrumentation was conducted before the session commenced.

During data collection, the subject pilot was in the left-hand seat of the aircraft, the safety pilot was in the right-hand seat, and the experimenter was behind them, adjacent to the experimental control computer. In each session, four experiments were conducted, one baseline experiment on auditory localization and three experiments examining the utility of 3D audio as a navigation aid or as an attitude indicator. All four experiments were conducted in a single test flight, and two such flights were conducted for each pilot - one for the GAINR-coupled 3D audio condition and one for the headtracker-coupled 3D audio condition. In order to complete all conditions, there were a total of 8 flights lasting approximately 2 hours each. All tests were conducted at airspeeds of approximately 170 Knots Indicated Airspeed (KIAS) at an altitude of approximately 12,000 feet, though these values varied slightly depending on weather and traffic conditions. Only one flight took place on each day, and all flight tests were conducted under day visual meteorological conditions (VMC).

3.0 EXPERIMENT I. AUDITORY LOCALIZATION

In order to assess the relative benefit of spatial audio cues for various flight tasks, it was important to first obtain a baseline measurement of the system performance and the localization ability of the subject pilots. A simple auditory localization task was implemented to examine these issues.

3.1 Specific Methods

3.1.1 Stimuli

The stimulus utilized in the auditory localization task consisted of an ongoing train of broadband noise bursts. The azimuthal direction of the stimulus was generated in reference to the aircraft heading at the time the trial was initiated. For example, if the aircraft magnetic heading was 180° and the stimulus was generated at 3 o'clock, the stimulus had a magnetic azimuth of 270°. The stimulus could originate from any of 12 clock positions (30° spacing) and at any of 3 elevations (-30°, 0° and +30°, referred to as low, mid, and high, respectively). The stimuli were spatialized relative to the orientation of the pilot's head (headtracker-coupled 3D audio) or the position and attitude of the aircraft (GAINR-coupled 3D audio), depending on the condition.

3.1.2 Procedure

The auditory localization test consisted of two tasks. First, a baseline auditory localization test was conducted while the aircraft was on the ground and stationary, with the engines running. Next, the same localization task was conducted in-flight.

In both tasks, the experimenter initiated the start of a trial by entering an input on a graphical user interface (GUI) on the experimental control computer, at which point the auditory cue was presented to the subject pilot. The subject pilot's task was to determine the direction of the cue and respond by stating a clock position and elevation (e.g., "2 o'clock, high"). The experimenter entered this response on the computer, then initiated the start of the next trial. One auditory cue was presented from each of 12 clock positions at each of the 3 elevations, for a total of 36 trials, the order of which was randomized for each subject pilot. For the in-flight portion of this experiment, the subject pilot first had to fly the aircraft to a desired altitude and achieve a desired airspeed at a safe heading. Once accomplished, the subject pilot was required to maintain straight and level unaccelerated flight while localizing the auditory stimuli. Due to hardware constraints, the localization tasks were completed for the GAINR-coupled 3D audio condition in the first session (i.e., the first flight) and for the headtracker-coupled 3D audio condition in the second session (i.e., the second flight) for all subject pilots.

3.2 Results

A 3-pole coordinate system [10] is employed here to examine the data from the auditory localization task. The coordinates of this system, Left/Right (L/R), Front/Back (F/B), and Up/Down (U/D), provide a particularly useful framework for examining auditory localization data because localization in each of these spatial dimensions appears to be mediated by different physical cues. In this system, the azimuthal position of a sound source is described in terms of a L/R coordinate, providing a measure of laterality, a F/B coordinate, providing a measure of how far a sound is in front of or behind a listener. Dividing azimuth in this way ensures that accurate performance in the L/R dimension will not be obscured by front-to-back or back-to-front reversals, which are not uncommon in spatial audio displays. For instance, a sound at 0° azimuth, 80° elevation that was judged to be at 180° azimuth, 80° elevation would suggest an error of 180° in azimuth using a traditional 2-pole coordinate system. However, in the 3-pole system, this error shows up as a 20° error in the F/B dimension, and no error in the L/R dimension. Finally, in this system the U/D coordinate is simply the elevation angle of the sound source.

Results from the localization tasks are shown in Figure 1, where mean angular errors in localization are plotted for each individual pilot, as well as the overall mean, in the L/R dimension (leftmost column), F/B dimension (middle column) and U/D dimension (rightmost column). In each panel, the black bars represent the data from the GAINR-coupled condition and the white bars represent the data from the headtracker-coupled condition. The top row depicts the results from the ground localization task and the bottom row depicts the results from the in-flight localization task. In both the ground and in-flight tasks, localization errors in the L/R dimension were generally the lowest, consistent with previous results in the literature [11]. Moreover, localization performance was best when the audio cues were coupled to the headtracker (i.e., when the cues were spatialized with respect to the orientation of the listener's head), with performance improvements of approximately 3° to 7° over the GAINR-coupled condition. This head-referenced display allowed the pilot to use head motion to turn and face the source directly, thus bringing the source into the "auditory fovea," where spatial acuity is greatest.

The difference across conditions was even more pronounced in the F/B dimension, where mean localization errors were reduced from approximately 37° in the GAINR-coupled display to 13° in the headtracker coupled display. The large F/B errors in the GAINR-coupled condition likely represent front-back confusions, which were greatly reduced when the pilot utilized the dynamic localization cues associated with head motion. Although similar cues could be available to the pilot in the GAINR-coupled condition by changing the heading of the aircraft, the pilots were required to maintain a straight and level course and in general were not exposed to those dynamic cues.

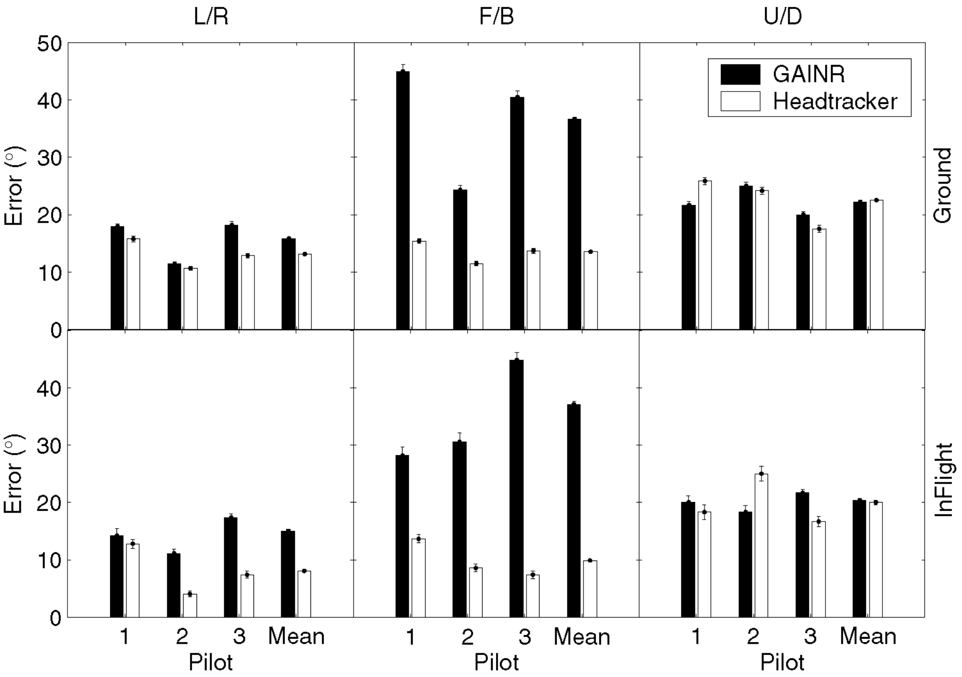

Figure 1: Localization errors for the individual pilots and the mean in the Left/Right, Front/Back, and Up/Down dimensions for the ground localization task (upper panels) and the in-flight localization task (lower panels). Both the GAINR-coupled (black bars) and headtracker-coupled (white bars) 3D audio conditions are shown. Error bars indicate ±1 standard error.

Localization in the U/D dimension was only slightly better than chance and was independent of the 3D audio condition. These results are not altogether surprising because the cues believed to mediate spatial hearing in the U/D dimension – fine-structure spectral cues in the HRTF– are highly individualized. Previous literature on spatial audio suggests that U/D localization improves when the display employs the listener's own HRTFs [12]. However, because the current display employed a generic set of HRTFs, those pilot-specific cues were not employed.

Performance for the in-flight localization task (bottom row) was not substantially different from the baseline ground localization task, although errors were slightly lower overall. This may be attributable to a learning effect, as the in-flight task always followed the ground task.

4.0 EXPERIMENT II. SPATIAL AUDIO AS A NAVIGATION AID

Once it was established that the 3D audio display could adequately generate the cues required to support accurate sound localization, the utility of this display as a flight navigation aid was evaluated. In Experiment II, a task was implemented in which spatialized audio cues were provided as a navigation aid to direct the flight path of pilots along a given heading.

4.1 Specific Methods

4.1.1 Stimuli

The stimulus employed in the navigation task was either a nonspatialized (diotic) verbal cue providing a command heading to the pilot (e.g., "Set Course 3-1-5"), which served as a baseline measure of performance, or a spatialized audio cue that appeared to originate from a particular direction. The spatialized cue consisted of the same noise stimulus used in Experiment I, interrupted at periodic intervals by the spatialized phrase “Set Course.” Note that no numerical heading cue was presented verbally in the 3D audio conditions. As in the localization task, the spatialized stimuli were positioned relative to the orientation of the pilot’s head (headtracker-coupled) or relative to the position and attitude of the aircraft (GAINR-coupled), and only one condition was tested in each block.

4.1.2 Procedure

Before the start of each data collection session, the subject pilot was required to maintain straight and level unaccelerated flight, during which the heading depicted on the Horizontal Situation Indicator (HSI) and the current time were determined to see if any discrepancy between the GAINR system and HSI existed. Once everything was checked, the experimenter initiated the start of a trial on the control computer, at which point the command heading was presented. The task of the subject pilot was to fly the aircraft from the initial heading to the command heading. When the pilot believed this had been accomplished, he indicated this to the experimenter, at which point the trial ended. In each experimental session, 20 navigation trials were collected, 10 in which the nonspatialized verbal cue was employed and 10 in which the spatialized audio cue was employed. Each cue was tested on alternate trials, and the command heading was generated randomly on each trial. In the event that this command heading was determined to be unacceptable due to traffic, planning, or other reasons, a new command heading was generated. As before, spatial audio cues were GAINR-coupled in the first experimental session for each subject and headtracker-coupled in the second session.

4.2 Results

Results from the navigation task are shown in Figure 2. For each pilot, the mean difference between the command heading and the final heading (angular error) is plotted for each condition. The overall mean is also shown. As can be seen, performance in the baseline verbal cue condition was always the best. Pilots were, on average, able to guide the aircraft to within 1°-2° of the command heading. This level of performance is not surprising given the fact that pilots regularly use this method for navigation, and the meaning of such a display is unambiguous. On the other hand, angular errors in the GAINR-coupled and headtracker-coupled conditions were substantially higher (8°-9°). One might expect performance in the headtracker-coupled condition to be slightly better than performance in the GAINR-coupled condition – a head-referenced display would enable the pilot to use exploratory head movements to determine the direction of the command heading with a high degree of accuracy. However, in order to truly minimize errors, the pilot had to be certain that the orientation of his head was perfectly aligned with that of the plane, and that is a nontrivial task. Moreover, the accuracy of the headtracker progressively degraded throughout the trials. Both of these issues likely contributed to these larger than expected errors in this condition.

Despite the larger errors found in the headtracker-coupled condition, the, pilots were able to get to their perceived command heading in a fast and efficient manner. Taking into account the differences in initial heading across trials, the results indicate that the pilots approached the command heading at similar rates in the verbal cue and headtracker-coupled cue conditions (1.25 deg/sec and 1.28 deg/sec, respectively), but much more slowly in the GAINR-coupled display (0.79 deg/sec). It is likely that the pilots had to initiate small aircraft maneuvers in the GAINR-coupled condition in order to resolve front/back confusions, which would result in comparatively longer trial durations. Nevertheless, as shown in Figure 3, the pilots were able to fly a relatively direct course to the command heading in all conditions.

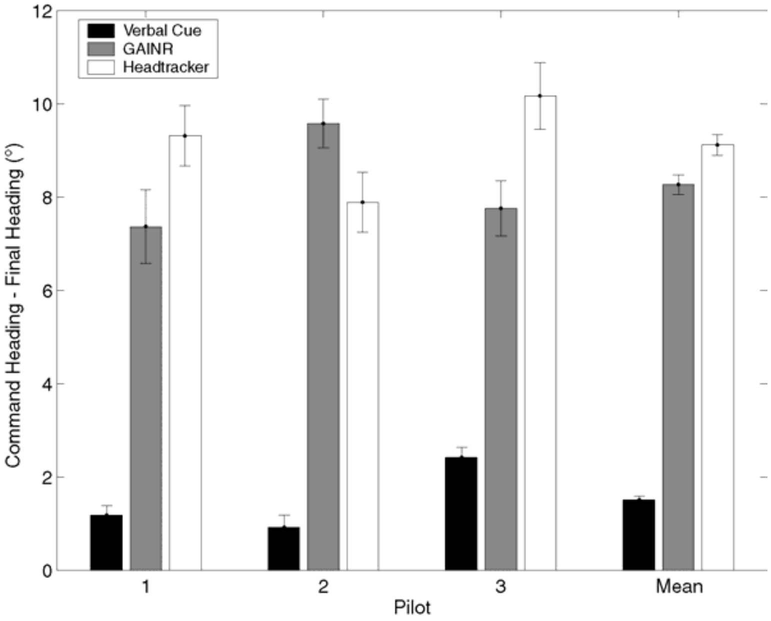

Figure 2: The difference between the initial heading and the command heading (in degrees) in each of the display conditions. Mean data are shown for each individual subject pilot as well as the mean of all three subjects. Error bars indicate ±1 standard error.

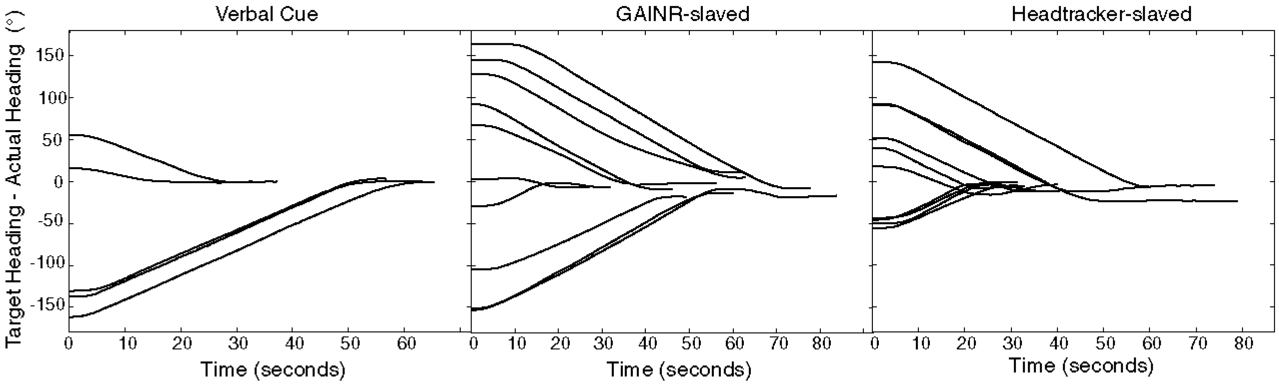

Figure 3: Flight paths as a function of time in each display condition for a single representative pilot.

5.0 EXPERIMENT III. SPATIAL AUDIO AS AN ATTITUDE INDICATOR: DETECTING CHANGES IN ATTITUDE

For Experiment III, an audio artificial horizon was developed. This display was designed to provide continuous attitude information to the pilot in order to enhance situation awareness.

5.1 Specific Methods

5.1.1 Stimuli

For this task, an audio artificial horizon cue was employed that provided the subject pilot with a continuous stream of information about the roll and pitch of the aircraft. Although the concept of an audio artificial horizon is not new, previous attempts have employed only relatively simple auditory stimuli such as tones in noise that vary in frequency, modulation rates, etc. [13]. These displays, while informative, can require a great deal of training to be effective. More importantly, they are not well tolerated by pilots because the ongoing nature of the sounds proves to be an annoyance and potentially a distraction. Because the goal is to provide *continuous* information about the attitude of the aircraft, these displays may be unsatisfactory.

In order to address these issues, an alternative audio display was developed that employs music as a continuous auditory attitude indicator. It was believed that such a display would be widely accepted by the pilot population, as it is not uncommon for pilots to listen to music over long periods on cross-country flights. For this task, the audio signal was obtained by allowing each pilot to select an audio CD with music they would enjoy over the course of the flight. The music on the CD was converted into a monaural '.wav' file by summing together the left and right audio channels, and then this '.wav' file was processed to add pitch and roll information and played back to the pilot over stereo headphones during in-flight testing. Roll was indicated by panning the music from the left ear to the right ear in response aircraft attitude. This was accomplished by decreasing the level of the signal in the ear associated with the higher of the two wings, with a maximum ILD of 48 dB when the magnitude of the roll was 30 degrees or greater. Pitch information was indicated in two ways. First, a spectral cue was provided in which the signal had a low-frequency emphasis when the aircraft was pitched down and a high-frequency emphasis when the aircraft was pitched up. This was implemented by passing the signal through a harmonic filter with peaks spaced every 350 Hz (low-frequency emphasis) or every 2500 Hz (high-frequency emphasis). A second cue for pitch related to the apparent image width of the audio signal. This manipulation was accomplished by varying the interaural cross-correlation of the signal such that level flight resulted in a perfectly correlated signal that was heard as a relatively compact auditory image; any deviation in aircraft pitch from level resulted in a lower correlation value and thus a wider auditory image. This decorrelation was achieved by convolving the left ear signal with a flat-spectrum 128-point random noise and the right ear signal with a time-reversed version of exactly the same 128-point random noise. These two cues (the frequency shaping and the interaural decorrelation) were scaled to produce a gradual transition from an unfiltered, perfectly correlated signal in a straight and level orientation to a maximally shaped, uncorrelated signal for a plane pitched up or down by 10 degrees. Because this display was intended to provide information about the attitude of the aircraft, it was believed that the most appropriate display would be one that was plane-referenced (i.e., coupled to the GAINR).

5.1.2 Procedure

In this experiment, the task of the subject pilot was to identify the direction of a change in the aircraft's attitude from level flight. After controlling the aircraft for a short time in order to become familiar with the dynamics of the audio horizon display, the subject pilot donned a blindfold and the safety pilot took over the controls of the aircraft. Data collection commenced when the subject pilot indicated that he was ready, at which point the experimenter started the trial and the safety pilot began to adjust the attitude of the aircraft (pitch up, pitch down, roll left, or roll right). The dimension and direction of change were chosen randomly from trial to trial. The attitude changes were always made at a rate of 1 degree per second, thus providing a repeatable and reliable stimulus. This slow rate also minimized any vestibular or "seat-of-the-pants" cues indicating motion. The trial ended when the subject pilot identified the direction of change with a verbal response (e.g., "Roll Left") and the experimenter entered the response on the control computer. If the maneuver brought the aircraft to the boundaries set as the limits for safe attitude (±30° of roll or ±10° of pitch), the trial was terminated. The audio condition was alternated from trial to trial, with the artificial horizon presented on the first trial.

5.2 Results

The results from the Change in Attitude task are shown in Figure 4, where the percent correct identification of the direction of attitude change is plotted for each actual attitude change condition, averaged across dimension (i.e., roll [left and right] and pitch [up and down]) and across subjects. The black bars represent performance in the task when no audio cue is given and the white bars indicate performance when the audio horizon cue is provided. As can be seen, the percentage of correct identifications improved substantially in both pitch and roll when the audio horizon was provided, with mean values of nearly 90% with the horizon cue as compared to approximately 43% with no audio cue. Moreover, pilots correctly identified the change in attitude when the aircraft pitched up 100% of the time, suggesting that this was a very robust cue. Overall, it appears that the cues for pitch may have been more salient and provided a slightly larger benefit relative to the baseline than the cues for roll.

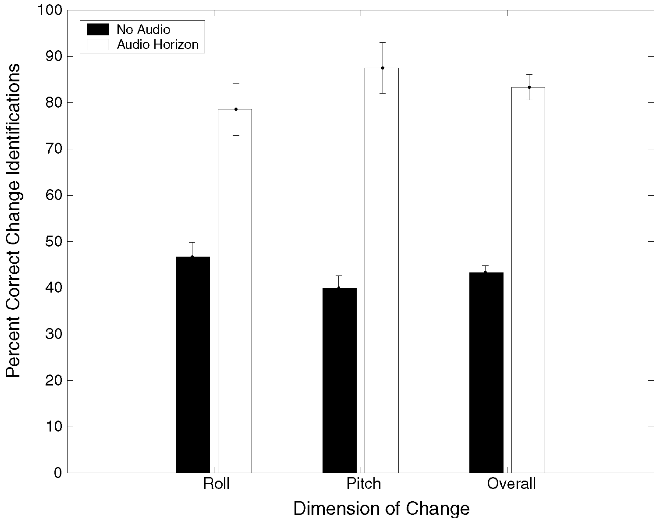

Figure 4: The percentage of correct identifications of the dimension and direction of change in the aircraft attitude. The black bars represent the baseline no-audio condition and the white bars represent performance with the artificial auditory horizon. Mean correct detections, averaged across dimensions, are also shown. Error bars indicate ±1 standard error.

Thresholds for the correct change in detection in each direction are shown in Figure 5. Threshold values for roll were reduced by nearly 7° with the audio horizon cue and threshold in pitch were reduced more than 3°. Recall that the response method in this task was verbal report, which is not ideal for time critical measures. Because of the inevitable delay between the subject pilot's response and the experimenter's input, there is reason to believe that the threshold values reported here underestimate the true performance of the pilots with this display. While these improvements are substantial, it is difficult to assess whether they are

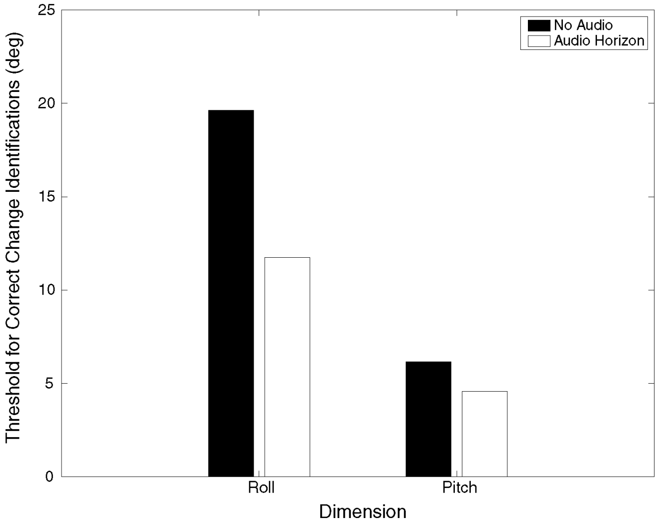

Figure 5: Mean threshold values for correct identification of the direction of change in roll and pitch, averaged across each of these dimensions. The black bars represent the baseline no-audio condition and the white bars represent performance with the artificial auditory horizon.

sufficient in all situations. Nevertheless, in the absence of visual information, the audio horizon provides a much more reliable cue for aircraft attitude than simple "seat-of-the-pants" cues.

6.0 EXPERIMENT IV. SPATIAL AUDIO AS AN ATTITUDE INDICATOR: MAINTAINING STRAIGHT AND LEVEL UNACCELERATED FLIGHT

In Experiment IV, the audio horizon display, in conjunction with a spatialized navigation aid, was evaluated as a means of supporting the pilot's ability to maintain straight and level flight in the absence of any out-the-window views or any visual instrument reference.

6.1 Specific Methods

6.1.1 Stimuli

The audio horizon cue utilized in Experiment III was also utilized in Experiment IV as an aid to the pilot for maintaining level flight. In addition, the audio cue used to indicate a command heading in Experiment II was used to provide a navigation aid to the pilots for maintaining straight flight. The navigation aid was always presented from a direction consistent with the aircraft's heading at the start of the trial.

6.1.2 Procedure

In this task, the subject pilot was required to fly the aircraft while blindfolded and maintain straight and level flight. Before the start of each trial the subject pilot entered a trim input on an acceptable heading. He then donned the blindfold and data collection was initiated. Each trial was 5 minutes in duration, during which the safety pilot monitored the status of the aircraft. If the aircraft flew beyond the attitude limits of this (±30° of roll and ±4° of pitch), the safety pilot took control of the aircraft and the trial was paused until straight and level flight was regained, at which point the trial resumed. Two such trials were conducted on each flight, one in which no audio cue was provided, and one in which the audio horizon display and the navigation aid were provided. The order of conditions varied across subjects.

6.2 Results

The ability of the pilots to maintain straight and level flight is shown in Figure 6. Here, RMS errors in roll (left panel) and pitch (right panel) are plotted for each individual subject and the overall mean in each display condition. Note that the ordinates on the two panels differ.

As can be seen, the audio horizon display (white bars) led to relatively stable performance in both roll and pitch. RMS errors for roll were extremely low for Pilots 2 and 3 (approximately 2°-3°) but higher for Pilot 1 (≈ 11°). Moreover, RMS errors in pitch were much lower overall than those for roll (≈ 1°). Note that the data from the baseline no-audio condition (black bars) are merely shown for comparison purposes. That is, the subject pilots found that they were generally unable to identify the attitude of the aircraft without the audio horizon and thus made few control inputs during this test condition. Therefore, the results shown reflect only the effectiveness of the initial trim input and the stability of the aircraft, and provide little information about pilot performance. The impact of this can be seen more clearly in Figure 7, which depicts the flight paths in the baseline no-audio condition (left panels) and 3D audio condition (right panels). In each panel, the deviation from straight and level in roll (top row) and pitch (middle row) are plotted as a function of flight time, as is the deviation from initial heading (bottom row). The dashed lines in each panel indicate the boundaries of safe travel in that particular dimension (±30 in roll and ±4 in pitch). The discontinuities in the data represent instances where the boundaries of safe travel were breached, the trial, and then restarted after straight and level flight was regained.

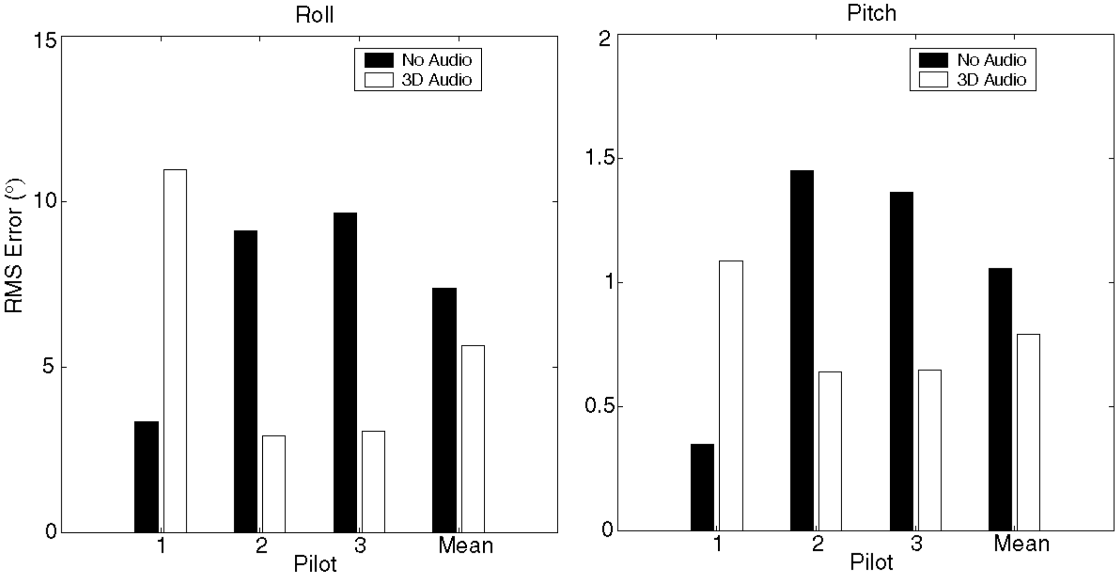

Figure 6: RMS errors between the ideal path and the actual path in the task requiring the subjects to maintain straight and level unaccelerated flight. The black bars represent the data in the no-audio condition and the white bars represent the data when the auditory artificial horizon display was provided.

As can be seen in Figure 7, the flight paths in the no-audio condition appear to be smooth relative to the flight paths in the 3D audio condition in roll, pitch, and heading – evidence that the pilots initiated very few control inputs in the no-audio condition. The surprisingly good performance for Pilot 1 in this no-audio condition, achieved with few inputs, further supports the spurious nature of the data from this condition. It is important to note that such results artificially reduce the difference between performance in the two display conditions and serves to underestimate the true benefit of the audio horizon display.

The right-hand panels show the deviation from straight and level flight when the artificial auditory horizon cue was provided. In both pitch and roll, Pilots 2 and 3 deviated little from the ideal path and never reached the boundary of safe travel. Pilot 1 reached this boundary only once in each dimension. Performance in the heading was also relatively good for Pilots 2 and 3. These results suggest that this display could be used effectively by pilots as an attitude indicator.

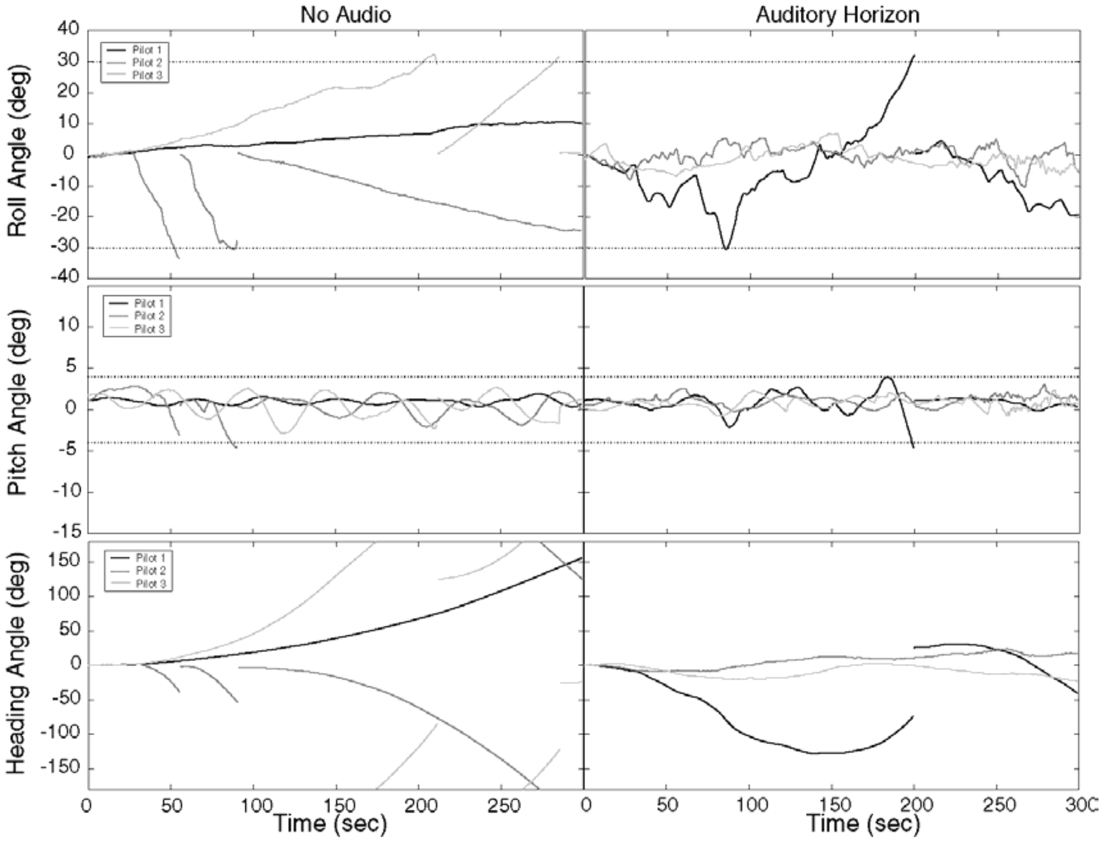

Figure 7: Flight paths travelled during the straight and level task. The left-hand panels depict the data from the trials when no audio cue was provided, and the right-hand panels depict the data when the artificial auditory horizon display was provided. The flight data are broken down into roll (top row), pitch (middle row) and heading (bottom row).

7.0 SUMMARY AND CONCLUSIONS

Overall, the results from these four experiments indicate that spatial audio displays may be useful as a means of providing information to a pilot about events in the relevant airspace and the status of the aircraft. This appears to be particularly true when displays and cues accessing other modalities are missing, obscured, or ambiguous.

The localization task served as a baseline measure of both system performance and the pilots' spatial hearing ability. However, spatial audio cues of this type may be used to provide information to the pilot regarding nearby air traffic and other obstacles. Indeed, in both basic laboratory experiments and immersive synthetic environment experiments, 3D audio has been demonstrated as an effective means of reducing visual target acquisition times by providing an audio cue that is co-located with the visual target [14,15]. However, it appears that the spatialization in such a display must be referenced to the pilot's head in order to be effective. That is, the current experiment demonstrated that a plane-referenced display (the GAINR-coupled display) resulted in larger localization errors overall and front/back ambiguities that made it more difficult for the pilot to determine the true location of the stimulus. Nevertheless, the headtracker-coupled 3D audio display still must be improved if it is to become a system upon which the pilots may truly rely. The headtracking system exhibited degraded performance over time, and thus influenced localization performance negatively. In addition, localization in the U/D dimension with 3D audio was always near chance, independent of the spatialization reference. Although it may be sufficient to cue the pilot to a general location, the ambiguity of target elevation may prove to be unacceptable to pilots. Improving the fidelity and reliability of the spectral cues (which are believed to mediate localization in this dimension) by employing better head-related transfer functions, or indeed the pilot's own head-related transfer functions, would likely lead to much better performance overall, and greater pilot acceptance.

The potential to use 3D audio cues as navigation aids has also been demonstrated. Although performance with 3D audio never reached that achieved with a verbal cue and HSI (a display system in use throughout aviation today), it is nevertheless the case that 3D audio may be used in addition to such a display, or in place of this display in the event that the pilot does not have access to the HSI to get to an approximate heading, after which other cues (e.g., out-the-window visual cues) may be employed. As in the earlier task, a 3D audio navigation aid is most useful when the sounds are displayed with respect to the orientation of the pilot's head rather than with respect to the aircraft attitude and position. Additional cues for distance, such as a speech distance cue based on vocal effort [16], would be a useful addition if the task is to navigate to specific waypoints.

The experiments have also shown that the pilots are able to use the auditory artificial horizon display for detecting attitude changes in the aircraft. It is clear that pilots were, in general, able to identify the change of attitude that took place on a given trial, and could use this information to maintain straight and level flight when no visual cues were available. This display may prove to be particularly useful for mitigating the effects of spatial disorientation by notifying the pilot that an unusual or undesirable attitude has been assumed by the aircraft. However, in both cases, the parameters of the display may need to be adjusted, or supplemented in some way in order to reduce the size of attitude change required for detection and identification.

One issue that is of great concern with displays in general, and auditory displays in particular, is pilot acceptance. The audio horizon display allows pilots to effectively choose their own stimulus – that is, select their own music. Thus, acceptance of this display was high. Because it was not perceived as a disturbance or annoyance, it could remain on continuously and effectively provide ongoing information about the attitude of the aircraft. It is important to note, however, that because the spectral content of the music varies across individual pieces of music, the display could be rendered more or less useful, or perhaps even misleading, for pitch discrimination. Additional work must be conducted to overcome this potential shortcoming. Redundant displays, such as verbal cues to indicate the crossing of a safe-flight boundary, may prove to be useful.

Although the current displays examined in these four experiments may not lead to performance in all tasks and in all situations at the same level as can be achieved with standard visual instruments and out-the-window views, the potential for 3D audio to be a useful addition to cockpit display systems is clear. 3D audio may be used in conjunction with visual displays to redundantly display critical information to the pilot, or it may be used when the standard visual displays are unavailable or unreliable. The utility of a display that can capture attention at all times and from all locations without putting a strain on the already overworked visual system cannot be overestimated. Although there are many aspects of the current display system that must be further refined, it is nevertheless the case that significant advancements have been made in this study, and it is clear that 3D audio can be a part of an overall display system to support pilot situation awareness in flight environments.

8.0 REFERENCES

[1] Veltman, J. A., Oving, A. B., & Bronkhorst, A. W. (2004). 3-D Audio in the fighter cockpit improves task performance. *International Journal of Aviation Psychology, 14,* 239-256.

[2] Bronkhorst, A. W., Veltman, J. A., & Van Breda, L. (1996). Application of a three-dimensional auditory display in a flight task. *Human Factors, 38,* 23-33.

[3] McKinley R. L., Ericson M. A., & D'Angelo W. R. (1994). 3-dimensional auditory displays: development, applications, and performance. *Aviation, Space, and Environmental Medicine, 65,* A31-38.

[4] McKinley, R. L., & Ericson, M. A. (1997). Flight Demonstration of a 3D Auditory Display. In R. H. Gilkey and T. R. Anderson (Eds.) *Binaural Hearing in Real and Virtual Environments*. (pp. 683-699). New Jersey: Lawrence Erlbaum Associates.

[5] Begault, D. R. (1993). Head-up auditory displays for traffic collision avoidance systems advisories: A preliminary investigation. *Human Factors, 35,* 707-717.

[6] Foyle, D. C., Andre, A. D., McCann, R. S., Wenzel, E. M., Begault, D. R., & Battiste, V. (1996). Taxiway navigation and situation awareness (T-NASA) system: Problem, design, philosophy, and description of an integrated display suite for low-visibility airport surface operations. *SAE Transactions: Journal of Aerospace, 105,* 1411-1418.

[7] Begault, D.R., & Pittman, M. T. (1996). Three-dimensional audio versus head-down traffic alert and collision avoidance system displays. *International Journal of Aviation Psychology, 6,* 79-93.

[8] Simpson, B. D., Brungart, D. S., Gilkey, R. H., Cowgill, J. L., Dallman, R. C., Green, R. F., Youngblood, K. Y., & Moore, T. J. (2004). 3D audio cueing for target identification in a simulated flight task. *Proceedings of the Human Factors and Ergonomics Society 48th Annual Meeting*, New Orleans, pp.1836-1840.

[9] Wenzel, E. M., Miller, J. D., & Abel, J. S. (2000). Sound Lab: A real-time, software-based system for the study of spatial hearing. *Proceedings of the 108th Convention of the Audio Engineering Society. Paris* New York: Audio Engineering Society.

[10] Kistler, D. J., & Wightman, F. L. (1992). A model of head-related transfer functions based on principal components analysis and minimum-phase reconstruction. *Journal of the Acoustical Society of America, 91,* 1637-1647.

[11] Wightman, F. L., and Kistler, D. J. (1989b). Headphone simulation of free-field listening. II: psychophysical validation. *Journal of the Acoustical Society of America, 85*, 868-878.

[12] Wenzel, E. M., Arruda, M., Kistler, D. J., & Wightman, F. L. (1993). Localization using nonindividualized head-related transfer functions. *Journal of the Acoustical Society of America, 94,* 111-123.

[13] **Lyons TJ, Gillingham KK, Teas DC, Ercoline WR, Oakley C. (1990). The effects of acoustic orientation cues on instrument flight performance in a flight simulator. *Aviation, Space, and Environmental Medicine, 61,* 699-706.**

[14] **Perrott, D. R., Saberi, K., Brown, K., & Strybel, T. Z. (1990). Auditory psychomotor coordination and visual search behaviour. *Perception and Psychophysics, 48,* 214-226.**

[15] **Bolia, R. S., D'Angelo, W. R., & McKinley, R. L. (1999). Aurally aided visual search in three-dimensional space. *Human Factors, 41,* 664-669.**

[16] **Brungart, D. S. (2000). Speech-Based Distance Cueing in Virtual Audio Displays. *Proceedings of the Human Factors and Ergonomics Society 44th Annual Meeting,* San Diego, pp.3714-3717.**