Wright State University

CORE Scholar

2014

Cued Visual Search and Multisensory Enhancement

Jordan Haggit
Wright State University

Follow this and additional works at: https://corescholar.libraries.wright.edu/etd_all

Part of the Industrial and Organizational Psychology Commons

Repository Citation

Haggit, Jordan, "Cued Visual Search and Multisensory Enhancement" (2014). *Browse all Theses and Dissertations*. 1389.
https://corescholar.libraries.wright.edu/etd_all/1389

Cued Visual Search and Multisensory Enhancement

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Science

by

Jordan M. Haggit
B.S., Wright State University, 2011

2014
Wright State University

Wright State University
GRADUATE SCHOOL

January 13, 2015

    I HEREBY RECOMMEND THAT THE THESIS PREPARED UNDER MY SUPER-VISION BY Jordan M. Haggit ENTITLED Cued Visual Search and Multisensory Enhancement BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF Master of Science.

_____
Robert H. Gilkey
Thesis Director

_____
Debra Steele-Johnson
Chair, Department of Psychology

Committee on
Final Examination

_____
Scott Watamaniuk, Ph.D.

_____
Brian Simpson, Ph.D.

_____
Robert E. W. Fyffe, Ph.D.
Vice President for Research and
Dean of the Graduate School

ABSTRACT

Haggit, Jordan. M.S., Department of Human Factors and Industrial/Organizational Psychology, Wright State University, 2014. *Cued Visual Search and Multisensory Enhancement.*

Previous research has been divided on whether or not multisensory cues can speed visual search relative to their component unisensory cues alone. Some studies (e.g., Mateo et al., 2012) found reaction times for multisensory cues were not faster than the RT of the faster component unisensory cue alone. Other studies (e.g., Oskarsson et al., 2012) found the multisensory cue to be faster than either unisensory cue alone (i.e., multisensory enhancement). This study aimed to determine whether the relative effectiveness match between auditory and tactile cues affects multisensory enhancement on a visual search task. In Experiment 1 we estimated for each subject three auditory cue inaccuracy values that corresponded to RTs equal to, 25% faster than, and 25% slower than tactile cue RTs. In Experiment 2 we combined each estimated auditory cue inaccuracy with a tactile cue to produce the multisensory conditions. We then compared RTs across the three different multisensory conditions. Our results suggest enhancement was more likely to occur when the auditory and tactile cues were closely matched in effectiveness and interference was more likely to occur when auditory and tactile cues were not closely matched. Although additional work will be needed to determine whether the interference was due to ineffective cues, poor strategies by the subjects, or some combination of these factors, our results seem to demonstrate the utility of providing two equally-matched cues as a strategy to speed visual search.

Contents

List of Figures

List of Tables

Acknowledgment

This thesis was completed thanks to a collection of efforts from a number of people. First, I would like to thank my advisor, Dr. Robert Gilkey. Certainly this would not have been possible without your continuing guidance and, often times, very necessary encouragement. I would also like to thank my other committee members, Dr. Scott Watamaniuk and Dr. Brian Simpson, whose eagerness to listen and enthusiasm for science were responsible for many productive and thought provoking discussions. Last, but most of all, I would like to thank my family. You strengthen me with your love and unwavering support of all my pursuits.

Dedicated to

Dad, Mom, Jared, and Tori

INTRODUCTION

Visual search is ubiquitous in everyday life, allowing us to detect, locate, and identify objects and events in the environment. In many cases, we engage in visual search for mundane goals (e.g., to look for a particular book on a bookshelf or for our car in a crowded parking lot). However, in some cases visual search can be critical to survival, allowing us to avoid or seek out events and objects important to our well being (e.g., staying away from a predator or finding food and water). Such critical consequences are likely in operational environments, such as the battlespace where mission effectiveness and survivability are often tied to the speed and accuracy of visual search (e.g., a ground soldier finding a sniper in a chaotic urban environment or an AWACS operator locating critical information in a cluttered display). In these situations, as well as many others, understanding the processes underlying visual search and developing strategies to increase its speed and accuracy should be of tremendous value.

By way of example, a ground soldier looking for a sniper may initially have little idea of where the sniper might be so a large region must be searched, which may surround the soldier in azimuth and range widely in elevation and distance. The search region may be filled with numerous distractors that are similar in appearance to the target (e.g., civilians, friendly soldiers, non-sniper enemy soldiers). In order to correctly identify the target and avoid catastrophic errors, the soldier must quickly foveate the sniper. Any strategy that would limit the size of the search region and/or help to foveate the target more quickly is likely to significantly increase mission effectiveness and survivability. It is not difficult to

imagine a third party who has some knowledge of the sniper's location, but is unable to address the threat themselves (e.g., another soldier without a clear shot or a UAV operator). So, they must try to communicate their knowledge so that the ground soldier can act. In many cases text, map, or other visual displays may be too slow, too distracting, or too complicated. Instead Ephrem et al. (2008) suggested that the third party could use a spatialized auditory display to create an audio signal that would be heard by the ground soldier as arising from the location of the sniper (audio annotation). Although such displays have been shown to reduce search times in ideal environments, in the real world, performance is likely to depend on noise level, cue accuracy, and display quality. Mateo et al. (2012) examined the effectiveness of a tactile display that delivered a vibration to the torso in order to cue the location of the target. Again, this display was effective (less effective than an auditory display) in the ideal laboratory setting, but would likely be less effective in a real world environment in which other vibrations and mechanical stimuli are encountered by a moving and twisting torso. This thesis will further examine the effectiveness of auditory and tactile displays to aid visual search, with particular emphasis on circumstances in which combining auditory and tactile displays might help to mitigate their individual limitations.

1.1 Visual search

Although the spatial resolution of the visual system is superior to that of other sensory systems (Welch and Warren, 1986), the best visual acuity is limited to the 1° to 2° of the visual field that falls on the fovea. Thus, although the gist of the visual scene can often be captured in a single glance, in order to process the fine details of particular objects, eye movements (and perhaps head movements) are needed to bring objects into the foveal region. Because the location of the relevant objects is often not known a priori, visual search is required.

In the laboratory, visual search tasks are generally implemented by simultaneously

displaying a target stimulus and a number of similar distractor stimuli on a computer monitor. The speed of target acquisition (i.e., RT) is examined as search parameters are varied. For example, search speeds depend on the phenomenological differences between the target and distractors, the number of distractors in the environment, and the eccentricity of the visual target. Typically, if the differences between the target and distractors are relatively large then RTs will be faster than if those differences are relatively small (e.g., when the tilt of a target line is different from the tilt of the distractors lines RTs are generally faster relative to when the target and distractors have the same tilt; Foster and Ward, 1991). If the differences between the target and distractors are small then RTs also tend to increase with the number of distractors. For example, Treisman and Gelade (1980) found an approximately 200% (800 ms) increase in RTs for 30 distractors compared to 1 distractor. Finally, RTs will be shorter for visual targets with small eccentricities (i.e., located nearer to the initial fixation point) and longer for visual targets with larger eccentricities (Wolfe, 1998). For stimuli with eccentricities beyond 60° elevation and/or 100° azimuth RTs can increase up to 150% (1000 ms) over stimuli presented at the initial fixation point (see Figure 1.1).

In the real world the size and complexity of visual scenes tends to make search more difficult than in the laboratory. Operators must often process large amounts of data, monitor multiple displays, or search through a large environment, which can easily overload the visual system.

1.2 Auditory cueing

Researchers (e.g., Perrott et al., 1990) have suggested that one of the most important functions of the auditory system is to direct the eyes to events or objects in the environment. Unlike vision, audition is a 360° sense. The ears can receive and localize acoustic stimuli that emanate from any location around the head, whereas visual stimuli must be located within the field of view (approximately 95° in azimuth and -80° to +60° in elevation; Wolfe
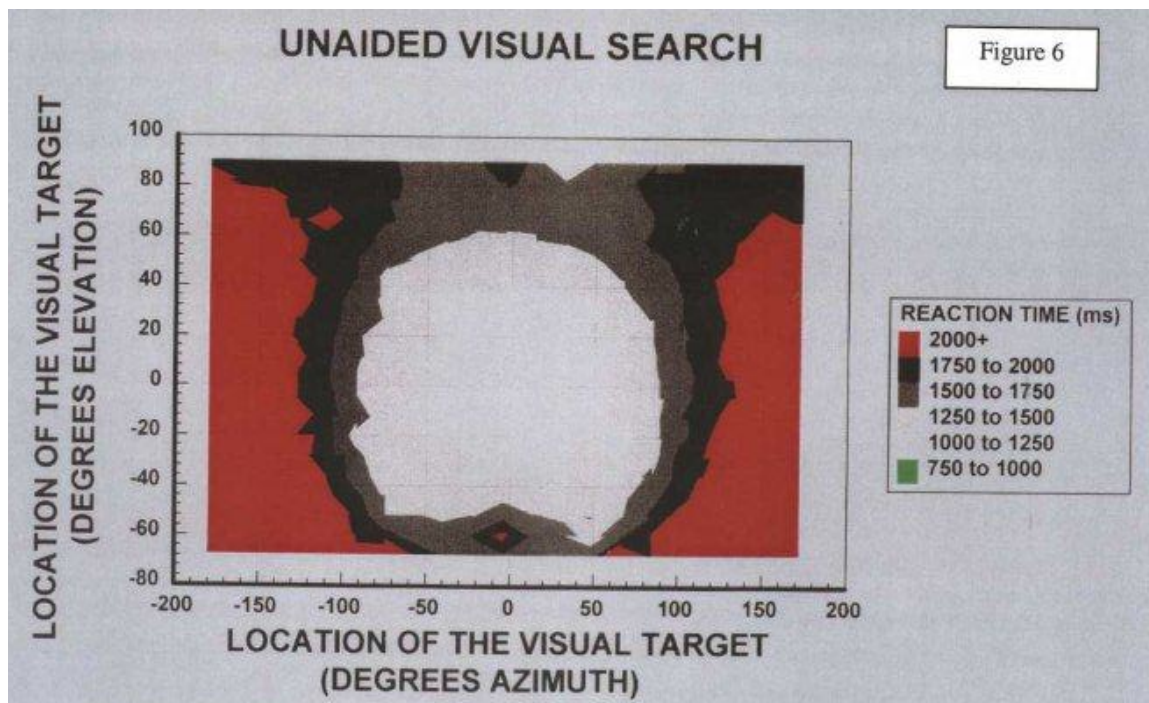
Figure 1.1: A contour plot showing response times as a joint function of target azimuth and elevation for the condition in which no cue was presented. The central light gray area represents response times ranging from 1000 to 1500 ms, the dark gray area represents response times ranging from 1500 to 1750 ms, the outer, black area represents response times ranging from 1750 to 2000 ms, and the red area represents response times above 2000 ms. (Used with verbal permission from Richard McKinley).

4

et al., 2009). So, a spatialized auditory display seems like an obvious strategy to speed search and help get the eyes to the region of the visual target.

A number of studies (e.g., Perrott et al., 1991, 1996; Flanagain et al., 1998; Bolia et al., 1999; McIntire et al., 2010; Nelson et al., 1998; Mateo et al., 2012) demonstrate the benefit of providing auditory cues to aid visual search. In many of these studies, subjects search for a visual target among distractors with the aid of a spatialized auditory cue that is either presented from a speaker that is co-located with the target (free field) or presented over headphones and spatialized using head-related transfer functions (virtual), in which a sound is generated so that it seems to arise at or near the location of the visual target. Typically, free-field auditory cues lead to faster search times than virtual cues. Nevertheless, in general, both types of spatialized auditory cues reduce search times compared to conditions where no auditory cue or a non-spatialized auditory cue is presented. Their greatest benefit tends to be when unaided search times are slow (e.g., for large eccentricities and large numbers of distractors). For example, when the target is behind the subject, auditory cues can reduce search times by 90% (700 ms) over uncued conditions (see Figure 1.2; Perrott et al., 1990). Importantly, the benefits of spatialized auditory cues are not only in areas outside the field of view. Even for targets in the central visual field (within a few degrees of the initial fixation point), spatialized auditory cues can still reduce search times by 40% (150 ms; Perrott et al., 1990).

Although auditory cues have been shown to improve search times in the laboratory, in real world environments environmental noise, technological issues, etc. can affect the rendering of auditory displays. Therefore, the cue may be inaccurate and not heard as arising from the intended location. Few studies have examined the impact of auditory cue accuracy on search performance and those that have either changed cue accuracy by a small amount (e.g., 6°; Rudmann and Strybel, 1999) or had subjects search over a small region of space (e.g., only within their current field of view; Vu et al., 2006).

DIFFERENCE IN REACTION TIME (MS) BETWEEN
THE UNCORRELATED AND CORRELATED
CONDITIONS
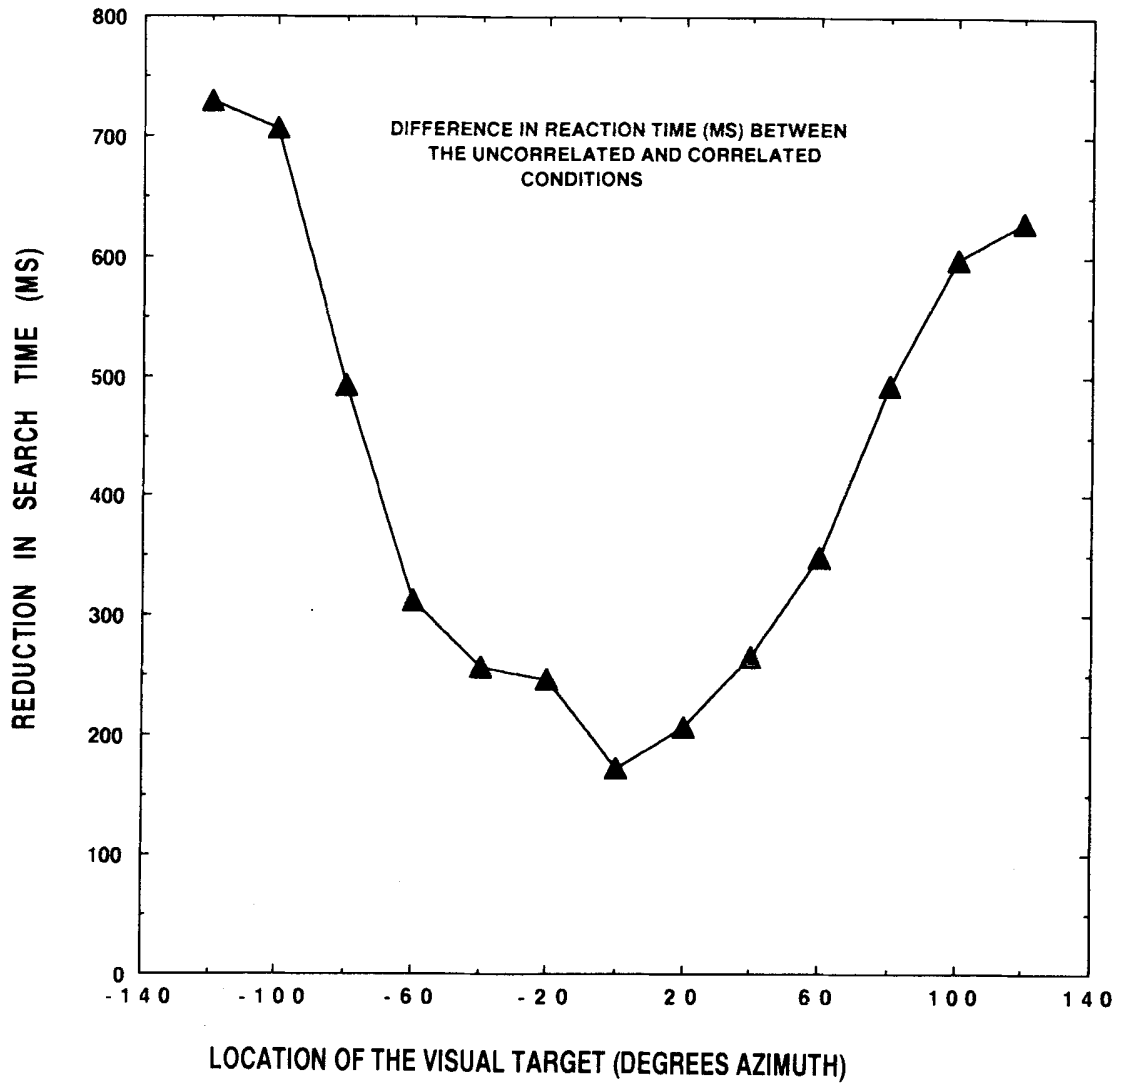
Figure 1.2: The mean reduction in visual search time when a spatialzed auditory cue is added as a function of target azimuth (i.e., eccentricity). (From Perrott et al., 1990, reprinted with permission from Sage publications.)

1.3   Tactile cueing

Although perhaps less straight forward to implement than auditory cues, spatialized tactile cues can also be used to help direct the eyes to the vicinity of a visual target. It is not obvious what type of stimulator or which body location is best suited for such a tactile display, but recent studies (e.g., Mateo et al., 2012; Hancock et al., 2013; Rupert et al., 2003) have generally employed clusters of vibrotactile stimulators on the torso. Despite the torso being less sensitive than other body parts, (e.g., the fingers or arms) it provides a relatively stable frame of reference to the 3-D space around the body as compared to the more sensitive, but mobile, limbs (Karnath et al., 1991). A number of studies (e.g., Lindeman et al., 2003; Tan et al., 2003; Mateo et al., 2012; Hancock et al., 2013) have demonstrated the effectiveness of tactile cues. Typically performed in the laboratory, these studies show tactile cues can speed search not only for small regions of space (e.g., targets on displays; Lindeman et al., 2003; Hancock et al., 2013) but also in omnidirectional environments (e.g., an area that covers 360° azimuth and -70° to +90° elevation; Mateo et al., 2012). The usefulness of tactile cues is not limited to "ideal" laboratory environments, but also extends to high-fidelity simulators and real world tasks. For example, Van Erp et al. (2006) found tactile cues reduced search times compared to no cue conditions when subjects had to detect threats in a flight simulator and McGrath et al. (2004) found that a tactile display presented to a pilot's torso could decrease workload, increase situational awareness (SA), and allow pilots to better maintain aircraft stability during a real-world flight task.

Although tactile cues can reduce the time required to find a visual target compared to no-cue conditions, typically search times with tactile cues are not as fast as with auditory cues (Ngo and Spence, 2010). Compared to auditory stimuli, it is more difficult to present tactile stimuli that naturally and consistently map to remote stimulus locations. Because the tactile sense is a proximal (near) modality tactile stimuli need to touch the skin, are perceived as on the skin, and thus cannot be co-located with a remote visual target. In contrast, auditory stimuli produce sound waves that can reach the ear from a distance, are

perceived as remote from the body at their source, and thus can readily be co-located with a visual target.

1.4 Multisensory Cueing

Multisensory cues (e.g., auditory + tactile cue) can also help direct attention to locations in the environment. In fact, some research suggests that they can do so more effectively than their component unisensory cues alone (Santangelo and Spence, 2007). For example, Hancock et al. (2013) had subjects simultaneously monitor 3 displays while searching for a visual target. In some conditions subjects received unisensory or multisensory cues to the screen where the target was located. Subjects found the target fastest and most accurately when they received the multisensory (audio-tactile) cue. Hancock et al. (2013) suggest that the benefits of multisensory cues may be attributed to redundant sensory information. That is, performance is improved when multiple sensory cues provide similar information (i.e., cue the same location in space).

Although some studies have found multisensory cues to be advantageous, in others they provided no benefit beyond the component unisensory cues. For example, Mateo et al. (2012) evaluated visual search performance under various conditions of auditory and tactile cueing. Among other conditions, they presented subjects with spatialized auditory and tactile cues together. They found that when the auditory and tactile cues were presented together RTs were not faster than when the auditory cue was presented alone. They explained their results in terms of the inverse-effectiveness principle of multisensory processing (see Stein et al., 2001). This principle suggests that there will be no multisensory benefit if one of the unimodal cues is very effective relative to the other. The auditory stimuli in Mateo et al. (2012) were much more effective than the tactile stimuli. As mentioned previously, the poorer performance with tactile cues may, in part, be due to the fact that touch is a near sense and the mapping to distal target locations was thereby somewhat artificial. In

addition, in Mateo et al. (2012) the auditory cue came from the exact location of the visual target, whereas the tactile cue only indicated a region of space that contained several possible target locations. In any case, the auditory cue may have been so effective that there was no apparent benefit from the addition of a less effective tactile cue and thus no multisensory enhancment. That is, multisensory enhancment might have been seen if the auditory cues had been less accurate.

So, we might expect that in operational environments in which one of the unisensory cues was reliably better than the other, there would be little advantage to a multisensory display. However, an auditory cue that in many situations may be very effective could be degraded by environmental noise, technological malfunctions, etc. So, also presenting a redundant tactile cue may lead to better performance.

1.5 Proposed Research

This thesis builds on the results of Mateo et al. (2012) and investigates whether multisensory enhancment is a function of the relative quality of the auditory and tactile cues. We implement a visual search task that includes auditory and/or tactile cues to the proximity of the visual target location. Like Mateo et al. (2012) we employ an omnidirectional search in which a target could be located anywhere around the subject (360° azimuth and -70° to +90° elevation). Under some conditions, spatial auditory and/or tactile cues to the visual target location are provided. Of particular interest are conditions that combine auditory and tactile cues, which might be expected to lead to better performance than either cue alone. In the first experiment, we systematically vary the inaccuracy of the auditory cue to the target's location in order to find auditory-only conditions that produce search times comparable to the search time for the tactile-only condition. Three different auditory performance levels are determined: one that produces performance equal to the tactile cue, one that produces performance approximately 25% faster than the tactile cue, and one that produces

performance approximately 25% slower than the tactile cue. The second experiment then compares performance for these auditory cue conditions to performance with these same auditory cues when combined with a tactile cue.

Based on the inverse effectiveness principle as well as data from Mateo et al. (2012) we do not expect to find multisensory enhancement for the multisensory cues that are comprised of unequal (in RT performance) unisensory cues (e.g., the tactile cue paired with the auditory cue that leads to performance 25% better than the tactile cue or the tactile cue paired with the auditory cue that leads to performance 25% worse than the tactile cue). We expect to find multisensory enhancment for the above conditions that comprise of equal (in RT performance) unisensory cues (e.g., the tactile cue paired with the auditory cue that leads to equivalent performance).

GENERAL METHOD

2.1 Participants

The same 4 paid subjects (all male) from the research subject panel at AFRL/RHCB at Wright Patterson Air Force Base and the author participated in each of the two experiments. Subjects were compensated for their time based on their normal hourly pay. All subjects had normal hearing, as well as normal or corrected-to-normal vision.

2.2 Apparatus

All testing was done at the Auditory Localization Facility (ALF; Figure 2.1) at Wright Patterson Air Force Base in Dayton, Ohio. ALF is a floating-floor anechoic chamber. The ceiling, floor, and walls are covered in 4-foot fiberglass wedges. In the middle of the anechoic chamber sits a geodesic sphere with a radius of 2.3 m. The sphere is equipped with 277 Bose (11-cm full-range) loudspeakers, which are situated at the vertices and are approximately 15° apart. Mounted on each loudspeaker is a four-element square cluster of light emitting diodes (LEDs), which presented the visual stimuli. Each LED cluster subtends 15' of visual angle. A 0.6m x 0.9m platform is located in the center of the sphere, upon which subjects stood during the experiment. The height of the platform was adjusted individually so that each subject's interaural axis aligned with 0° elevation in the sphere.

Figure 2.1: Auditory Localization Facility (ALF), Wright Patterson Air Force Base, Dayton, Ohio.

ALF is also equipped with an Intersense IS900 six-degrees-of-freedom tracking system that provides real-time location and orientation information about the subjects' head, via a headtracker, as well as a hand-held response wand, which is used to make and track the location of responses.

2.3   Visual Search Task

Subjects searched for a target among a number of similar distractors. The target was an LED cluster in which two or four of the four LEDs were illuminated, whereas the distractors were LED clusters in which one or three of the four LEDs were illuminated. Of the total 277 LED clusters spaced around the sphere, only 269 were used; some were not included due to obstructed visibility (below -70°elevation) and others due to technical difficulties with the associated speakers. The remaining 269 LED cluster locations were divided into 61 regions around the sphere (see Figure 2.2) where each region consisted of 2-7 LED clusters. Unbeknownst to subjects (with the exception of Subject 5, the author) only two LED clusters in each region were designated as possible target locations. This corresponded to 122 possible target locations approximately evenly spaced around the sphere. Distractors could be located at any of the possible 269 LED cluster locations. At the beginning of each block the experimenter assisted the subject into ALF and equipped them with the headtracker, wand, and a tactile vest (described in Section 2.4.3). To begin each trial, each subject pointed their head at the speaker/LED cluster at 0°azimuth and 0°elevation (i.e., the trial would only start if the subject's head was aligned within approximately 3.5°azimuth and 7.5°elevation of the speaker/LED cluster at 0°azimuth and 0°elevation) and then depressed the trigger on the bottom of the wand (to facilitate the alignment, during this period the LED cluster closest to where the subject's head was pointing at any given instant was illuminated). The trial did not begin unless the alignment and trigger press occurred together. The trial began 250 ms after the trigger was pressed, at which time the target and
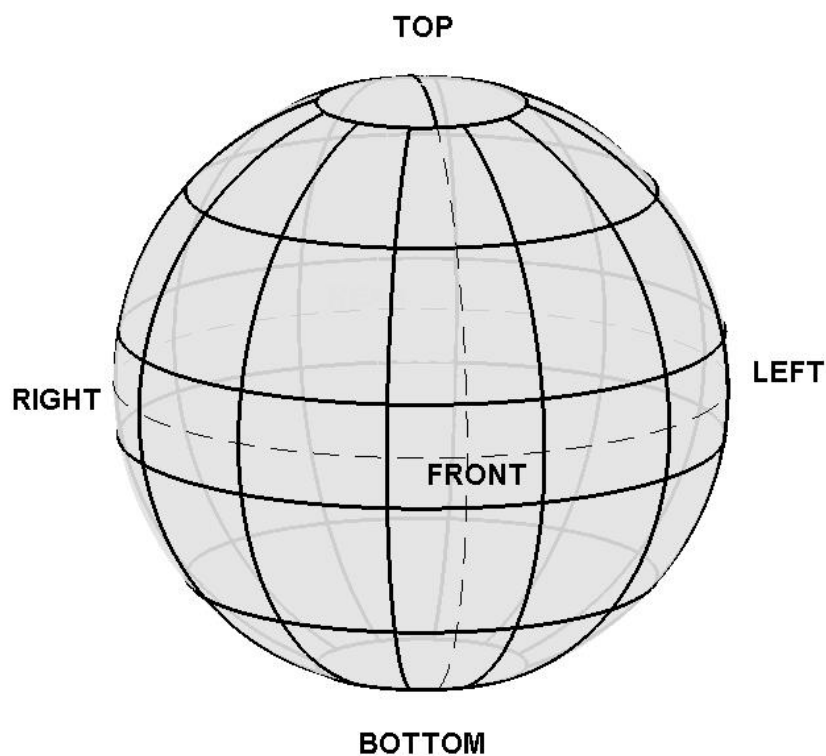
Figure 2.2: Depiction of the 62 stimulus regions above -70° elevation. Each region includes two, randomly selected, target locations. The regions were restricted in the following ways: no speakers were chosen below -70° in elevation because view of them was obstructed; above +70° in elevation was divided into two regions; the region centered at -90° azimuth that ranges from -70° to -40° in elevation was not used because the speaker hardware was unreliable; the remaining portion of the sphere was divided into 59 regions, which were 30° wide in azimuth. The regions centered at 0° elevation were 20° high in elevation, all other regions were 30° high in elevation.

distractor LEDs were illuminated and remained on until a response was made. Responses were based on the characteristics of the target LED cluster. For targets that contained 2 LEDs, subjects should have responded with either of the left 2 buttons on the wand and for targets that contained 4 LEDs, they should have responded with either of the right 2 buttons. Subjects were told to respond as quickly as possible while keeping the percentage of correct target identification responses above 95%.

2.4   Cueing

Depending on condition, a spatialized auditory and/or tactile cue was provided to help the subject locate the visual target.

2.4.1   Auditory stimuli

Auditory cues were bursts of 65-dB SPL white noise. They were gated on simultaneously with the visual target and remained on for a duration of 250 ms (0-ms rise/fall ramps).

2.4.2   Auditory cue inaccuracy

Mateo et al. (2012) suggested that they had failed to observe multisensory enhancement because their auditory cue was too effective. To vary the effectiveness of our auditory cue, we manipulated the spatial inaccuracy. In the auditory cue conditions, the cue was presented from a loudspeaker in ALF that was at or near (depending on auditory cue inaccuracy level) the location of the target. A given level of auditory cue inaccuracy refers to the maximum distance (in degrees) that the cue location could be from the actual target location. In this way, the auditory cue inaccuracy defines a selection set of speakers around the target LED cluster and the actual cue location on a given trial is randomly selected from this selection set. For example, if the target LED cluster is at X azimuth and Y elevation, an auditory cue

inaccuracy level of 30° defines a selection set, which includes 19 speakers, with the furthest speaker being no more than 30° from the target location. On a given trial the cue location was randomly selected from this selection set. (Note, if the auditory cue inaccuracy is 0°, then the selection set always contains exactly one speaker and the auditory cue is always collocated with the target LED cluster).

2.4.3   Tactile stimuli

The tactile cue was presented through a Tactile Torso Display "vest," which was designed and developed by The Netherlands Organization for Applied Scientific Research (TNO). This tactile vest consists of 62 individual tactors that are arranged so that there is 1 tactor on each shoulder and the remaining 60 tactors are situated in a 5-row by 12-column array around the torso (note, in one column one of the tactors was not used because the corresponding region had unreliable speaker hardware, see Figure 2.2). The vest was always fitted so that one column of tactors was aligned with the navel and one column of tactors was aligned with the spine. The remaining 10 columns were adjusted so that they were approximately evenly spaced around the torso of the subject. During each trial one of the 61 tactors was activated based on the location of the target in the sphere. Each tactor on the vest was mapped, through training, to one of the 61 regions (see Figure 2.2) in the ALF sphere. Roughly, the mapping was such that if you imagine a point in the center of the body, about navel high, a vector extending from that point out through a tactor would hit that tactor's region. When the target is presented in a given region the tactor assigned to that region was activated at the same time. All tactile stimuli vibrated at 250 Hz. They were presented simultaneously with the visual target and remained on for 250 ms.

2.5 Procedure

2.5.1 Training

Subjects were trained in the basic visual search task (with no auditory or tactile cues) to familiarize them with the trial timing and response procedure. Then subjects received two additional types of training: auditory training and tactile localization training. The purpose of the auditory training session was to familiarize subjects with the auditory cue only conditions. For this, subjects completed one 61-trial block for both 30- and 100-distractor cases in the basic visual search task with the addition of an auditory cue that indicated the location of the target.

The purpose of the tactile localization training was to familiarize subjects with the intended tactor-to-LED cluster mapping (see Figure 2.2), since there was not necessarily a natural mapping between the tactors and the distal visual stimuli. Following the procedure used by Mateo et al. (2012) the tactile localization training was divided into two parts. First, the subjects became familiar with all possible LED clusters assigned to each tactor's region. For the first 5 minutes of training each tactor as well as all LED clusters in that tactor's region of the sphere were sequentially activated. No response was required from the subjects but they were instructed to pay attention to the mapping and to try and keep their torso forward facing, only rotating when necessary (e.g., to see LED clusters in the rear hemifield). Next, subjects performed the localization task. For this, one tactor was activated on a given trial and subjects were to use the response wand to try and point to any LED location within the region that correctly corresponded to the learned mapping for the activated tactor. In order to allow subjects to see where the wand was pointing prior to making a selection the response wand was tracked so that it illuminated whichever LED cluster was closest to where the subject was pointing at a given time. After a response was made, feedback was given by illuminating all LED clusters in the correct region for the activated tactor on that trial. Once subjects pressed a button on the wand indicating they

had seen the feedback they were allowed to continue to the next trial. Subjects performed 2 61-trial blocks of the tactile localization training. After the tactile localization training subjects completed 2 61-trial blocks in the basic visual search task with the aid of tactile cues to the location of the visual target.

In addition, prior to data collection in each experiment, subjects completed 2 practice blocks under each condition in that experiment to familiarize them with each type of cue. For all types of training and practice tasks subjects were required to maintain target identification accuracy above 95%.

2.5.2 Experimental trials

Once subjects completed training they began experimental trials. Experimental trials consisted of the visual search task with the addition of an auditory cue and/or tactile cue to the target's location.

Experiment 1

3.1 Purpose

The purpose of Experiment 1 was to estimate appropriate parameter values for use in Experiment 2. For each subject we estimated one auditory cue inaccuracy value that would lead to an average RT that was equal to the RT with the tactile cue (A_T), one that was 25% faster than with the tactile cue (A_{T-25}), and one that was 25% slower than with the tactile cue A_{T+25}).

3.2 Conditions

We manipulated cue-type and number of distractors using a 4 (cue-type) by 2 (number of distractors) within subjects design, for a total of 8 conditions. The 2 levels of number of distractors were 30 or 100. The 4 levels of cue-type were 0° auditory cue inaccuracy, 30° auditory cue inaccuracy, 60° auditory cue inaccuracy, and tactile cue only condition (T). For the auditory cue inaccuracy conditions, subjects were instructed that auditory cues would always be presented in the vicinity of the target but that the distance from the target would vary from trial to trial and block to block. We chose these values for auditory cue inaccuracy to span tactile performance based on tactile data from Mateo et al. (2012) and our pilot data.

In Experiment 1, subjects first completed the training described in Section 2.5.1. Then each subject completed 32, 61-trial experimental blocks. The training and experimental trials took approximately 12 hours to complete for each subject.

3.3 Results and Discussion

The raw data for all conditions and subjects from Experiment 1 are shown in Table 3.1 and Figure 3.1 (30-distractor case) and in Table 3.2 and Figure 3.2 (100-distractor case). In the figures each subplot represents data from one subject and shows mean RT plotted as a function of auditory cue inaccuracy for the auditory conditions ($0°$ inaccuracy, $30°$ inaccuracy, $60°$ inaccuracy). As expected, RTs for each subject increased monotonically with the inaccuracy of the auditory cue (e.g., RTs were shorter in the $0°$ inaccuracy condition than the $30°$ inaccuracy condition, which were shorter than the $60°$ inaccuracy condition). These results are similar to data from previous research that suggests increasing the distance between an auditory cue and visual target leads to longer search times (e.g., Rudmann and Strybel, 1999; Vu et al., 2006).

For each subject, the three auditory cue inaccuracy levels (A_T, A_{T-25}, and A_{T+25}) to be utilized in Experiment 2 were estimated as follows. First, straight lines were fit to the RT vs. inaccuracy data for each subject (the dashed lines in Figure 3.1 and 3.2). Then initial estimates of the auditory cue inaccuracy levels that produced RTs equal to the RT obtained with the tactile cue (the solid lines) and that produced RTs 25% faster and 25% slower than the RT obtained with the tactile cue (the dotted lines) were read off the lines. For each subject this method yielded 6 inaccuracy values, which corresponds to one value for each auditory condition (A_T, A_{T-25}, A_{T+25}) for each number of distractors [these values are shown

Table 3.1: RT data (in ms) for all subjects in the T, 0°, 30°, and 60° auditory cue inaccuracy conditions for the 30-distractor Case.

Subject	T	0°	30°	60°
1	3849	2137	2904	6069
2	2342	1416	2586	4801
3	3946	1839	2657	4284
4	2713	1769	2224	4734
5	1779	1162	1755	3417
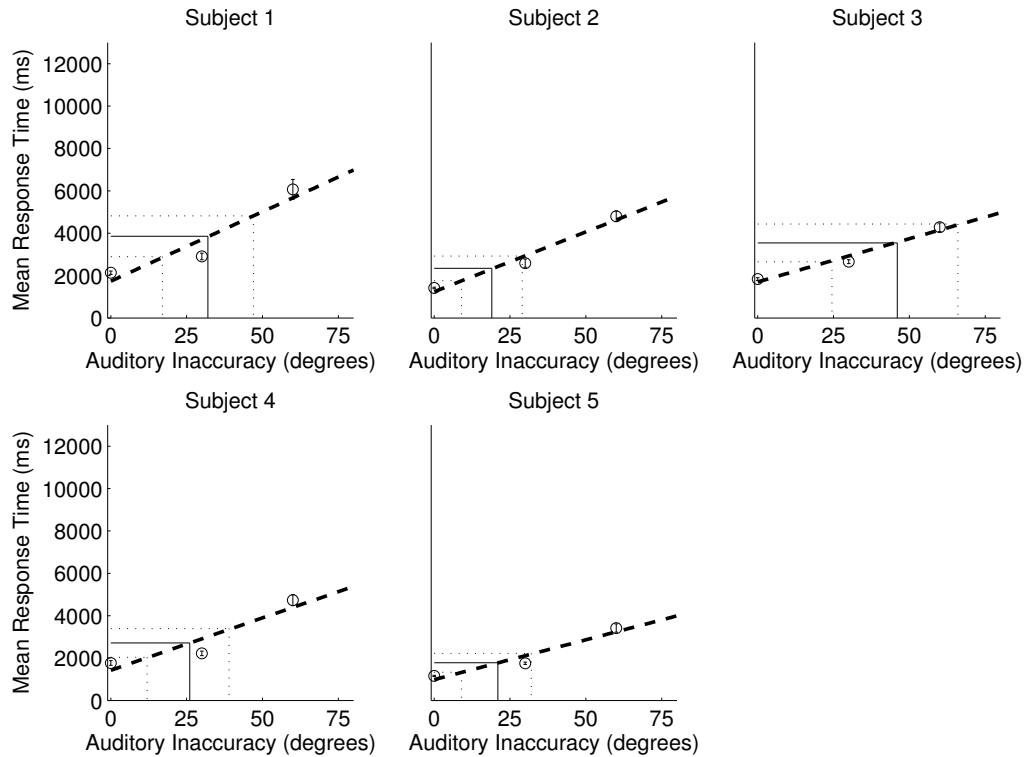
Figure 3.1: In each plot RT is plotted as a function of auditory cue inaccuracy. Each subplot represents data from one subject in the 30-distractor case. The circles represent RTs for the three auditory conditions (0°, 30°, 60°) to which the dashed line was fit. The solid horizontal line shows RT for the T conditions and the dotted horizontal lines show RTs for the T + 25% and T - 25% estimates. The vertical lines show estimated inaccuracies for the A_T, A_{T-25}, and A_{T+25} conditions.

Table 3.2: RT data (in ms) for all subjects in the T, 0°, 30°, and 60° auditory cue inaccuracy conditions for the 100-distractor Case.

Subject	T	0°	30°	60°
1	7282	3159	4675	10772
2	3149	1875	4845	10338
3	5796	2333	5178	9731
4	4381	2062	4253	12845
5	2603	1319	2585	6779
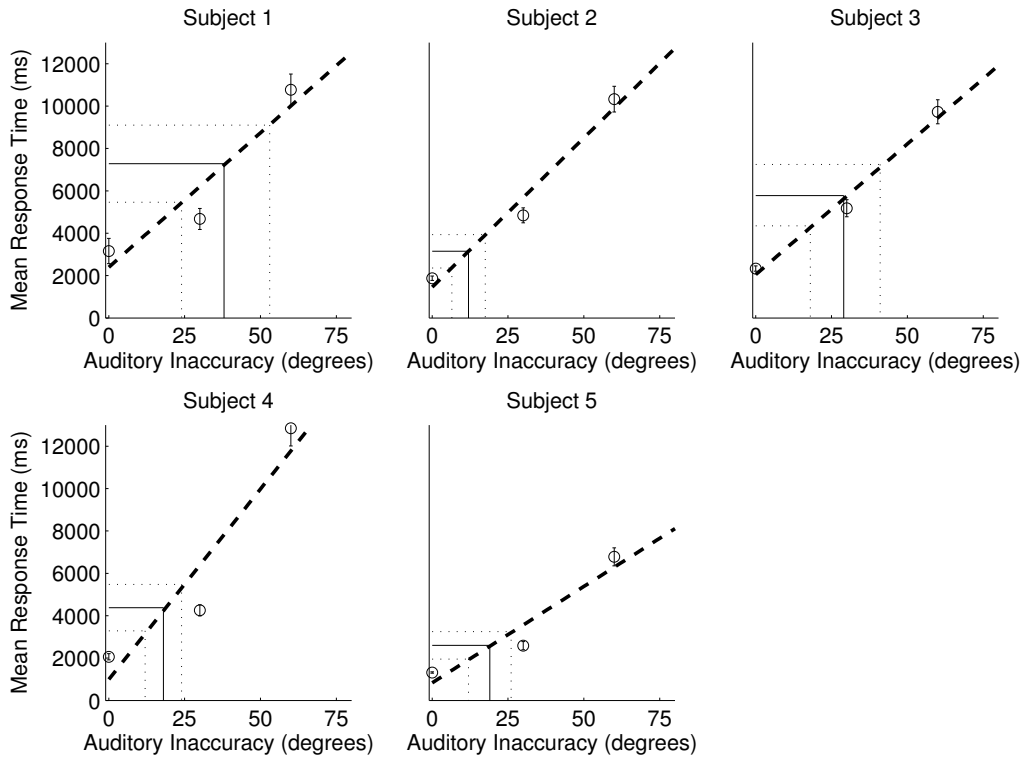
Figure 3.2: In each plot RT is plotted as a function of auditory cue inaccuracy. Each subplot represents data from one subject in the 100-distractor case.The circles represent RTs for the three auditory conditions (0°, 30°, 60°) to which the dashed line was fit. The solid horizontal line shows RT for the T conditions and the dotted horizontal lines show RTs for the T + 25% and T - 25% estimates. The vertical lines show estimated inaccuracies for the A_T, A_{T-25}, and A_{T+25} conditions.

Table 3.3: Estimated and Adjusted Auditory Cue Inaccuracy Values for 5 Subjects Corresponding to the A_{T-25}, A_T, and A_{T+25} Conditions for the 30-distractor Case.

Subject	A_{T-25}		A_T		A_{T+25}	
	Estimated	Adjusted	Estimated	Adjusted	Estimated	Adjusted
1	17	—	32	—	47	—
2	9	—	19	—	29	—
3	26	—	46	—	66	—
4	12	—	26	—	39	—
5	9	—	21	—	32	—

Note: No values needed to be adjusted.

Table 3.4: Estimated and Adjusted Auditory Cue Inaccuracy Values for 5 Subjects Corresponding to the A_{T-25}, A_T, and A_{T+25} Conditions for 100-distractor Case.

Subject	A_{T-25}		A_T		A_{T+25}	
	Estimated	Adjusted	Estimated	Adjusted	Estimated	Adjusted
1	24	—	38	—	53	—
2	6	—	12	15	17	25
3	18	—	29	—	41	—
4	12	—	18	—	24	—
5	12	—	19	—	26	—

Note: The estimated values for Subject 2 needed to be adjusted.

in Table 3.3 (30-distractor case) and Table 3.4 (100-distractor case)].

Note that although this method of estimation worked in most cases, in one case the estimated levels were adjusted to avoid redundant measurements in Experiment 2. Although inaccuracy varied continuously, the number of speakers in the selection set varied discretely and so different inaccuracy values could lead to identical selection sets and thus are not functionally different. So for example inaccuracy levels of 0° and 5° would both include only one speaker in the selection set. In order to avoid running a subject under 2 nominally different but functionally identical conditions in Experiment 2, we adopted the following rule. The three auditory cue inaccuracy values for a given subject had to define selection set sizes that differed from each other by at least 5 speakers. Under this rule, inaccuracy values for one subject were adjusted in the 100-distractor case and none were adjusted in the 30-distractor case (see Table 3.4, Subject 2).

Experiment 2

4.1 Purpose

The purpose of Experiment 2 was to use the inaccuracy values obtained from Experiment 1 to further examine multisensory cueing and determine the conditions necessary for multisensory enhancement.

4.2 Conditions

We manipulated cue-type and number of distractors using a 7 (cue-type) by 2 (number of distractors) within subjects design, for a total of 14 conditions. Visual targets and distractors were the same as in Experiment 1. Cue-type included 4 unisensory cue conditions and 3 multisensory cue conditions. The 4 unisensory cues consisted of the tactile cue only condition (T) and the 3 auditory cue only conditions based on the estimated inaccuracies for each subject from Experiment 1 (A_{T-25}, A_T, A_{T+25}, from Table 3.3 and Table 3.4). The 3 multisensory conditions combined the tactile cue and an auditory cue with one of the three levels of auditory cue inaccuracy (A_T +T, A_{T+25} + T, and A_{T+25} + T). For the unisensory conditions subjects were instructed in the same way as in Experiment 1 and for the multisensory conditions subjects were instructed to try and utilize information from both cues in order to find the target.

In Experiment 2, prior to experimental trials, subjects completed two 61-trial blocks of training in the visual search task for each of the tactile cue only conditions, auditory cue only conditions, and multisensory cue conditions under the 30- and 100-distractor cases. The experimental trials consisted of 56, 61-trial blocks per subject. The training and experimental trials took approximately 14 hours to complete for each subject.

4.3   Results and Discussion

The data for the unisensory conditions (tactile only and auditory only) from Experiment 2 are shown in Figure 4.1 (30-distractor case) and Figure 4.2 (100-distractor case). Each cluster of three bars shows the results for one subject. The dashed line in each cluster shows RTs under the T condition. The white, light gray, and dark gray bars show RTs under the A_{T-25}, A_T, and A_{T+25} conditions, respectively. Although the ordering of these auditory cue only conditions was as anticipated (i.e., A_{T-25} was faster than A_T, which was faster than A_{T+25}, in all cases), we typically did not hit our performance targets. In general, RT performance under the T condition was comparable to the performance observed in Experiment 1, but RTs under the auditory cue only conditions were better than we had estimated based on performance in Experiment 1. For example, had our estimates been accurate, the height of the light gray bars for each subject would match the height of the dashed line. However this is typically not the case. The A_T condition was faster than the T condition for 4 out of 5 subjects with 30 distractors and for all 5 subjects with 100 distractors. Indeed, in several cases, even the A_{T+25} condition was faster than the T condition.

Underestimating the RT performance with the auditory cues was likely the result of estimating from the straight lines fit to the data of Experiment 1. The true relationship is more likely concave upward. So, for example, in Figures 3.1 and 3.2 the data points for the 30° inaccuracy condition consistently fall below the fitted line (i.e., observed perfor-

Figure 4.1: RT plotted as a function of subject number for the A_{T-25} (white bars), A_T (light gray bars), and A_{T+25} (dark gray bars) conditions in the 30-distractor case. The dashed line with each group of bars represents performance in the T condition for that subject.
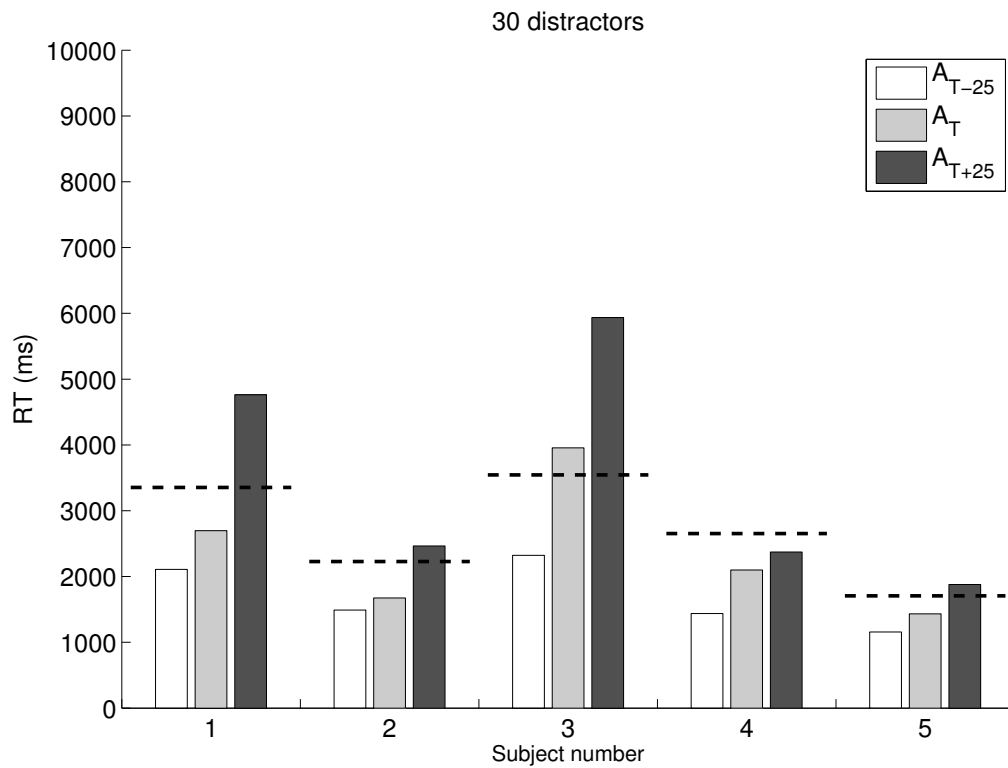
Figure 4.2: RT plotted as a function of subject number for the $A_{T\text{-}25}$ (white bars), A_T (light gray bars), and A_{T+25} (dark gray bars) conditions in the 100-distractor case. The dashed line with each group of bars represents performance in the T condition for that subject.
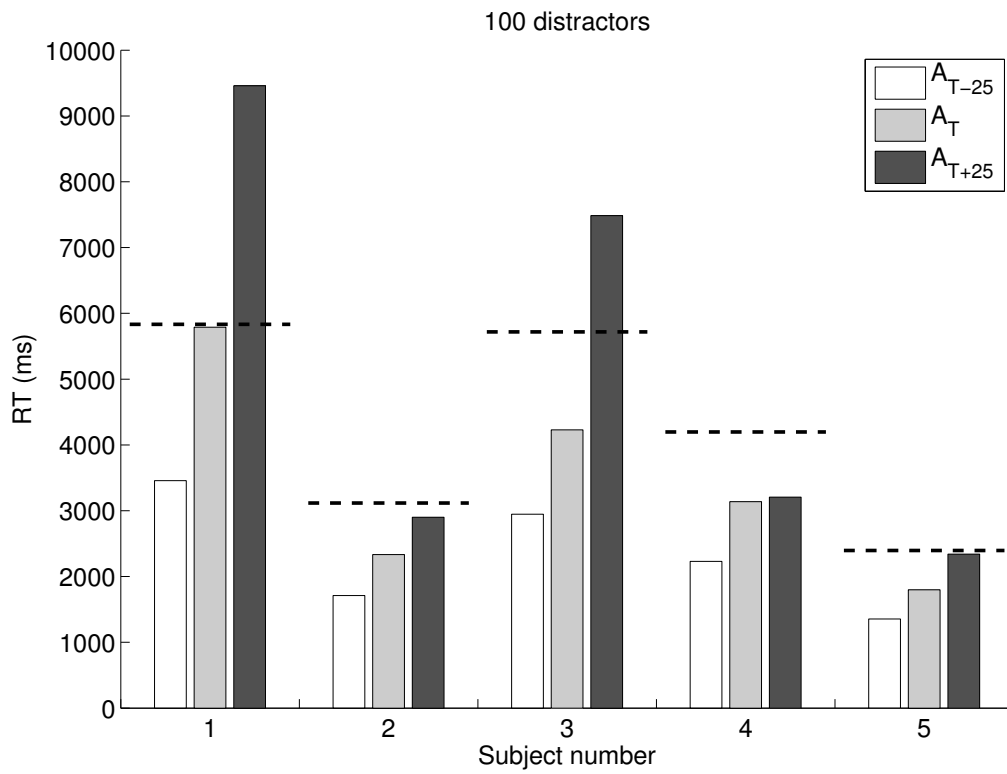
mance was better than performance predicted by the linear function). Because the target RTs (based on the tactile cue only performance) fall in this region, the estimated auditory inaccuracies were too low.

For the multisensory conditions, we had predicted that the observed enhancement (the better unisensory RT minus the multisensory RT) would be larger under the A_T+T condition than under either A_{T-25}+T or A_{T+25}+T condition (i.e., because the cue in the unisensory T and A_T conditions were expected to be equally effective). However, this type of comparison was not viable because in most cases none of the auditory cue only conditions were well matched to the T condition. That said, some auditory cue only conditions were better matched to the T condition than other auditory cue only conditions. In Figure 4.3 (30-distractor case) and Figure 4.4 (100-distractor case), we examined how the amount of multisensory enhancement changed as a function of the mismatch in the effectiveness of the unisensory cues (positive values on the abscissa indicate the auditory cue RT was faster than the tactile cue RT; negative values indicate the tactile cue RT was faster than the auditory cue RT). Common symbols represent data of one subject for the three multisensory cue conditions. The two dashed lines were simultaneously fit to the data using *fminsearch* (MATLAB R2012a Student, The MathWorks), under the constraint that the lines fit to the positive and negative mismatches shared a common y intercept but had separate slopes. These fitted functions account for 59% and 43% of the variance for the 30-distractor and 100-distractor cases, respectively. Based on these fits, we expect that we would have observed enhancement of approximately 397 ms and 753 ms for the 30- and 100-distractor cases, respectively, if the auditory and tactile cues had been equally effective. Note, however, that when the mismatch is large (greater than approximately +720 ms or less than -1170 ms for the 30-distractor case; greater than approximately +1400 ms or less than -3280 for the 100-distractor case), "negative" enhancement (i.e., interference) is often observed (in total, 8 out of 30 data points suggested interference – 5 for the 30 distractor case; 3 for the 100 distractor case). That is, when the mismatch was large subjects seemed unable

to simply focus on the more effective cue and responded more slowly than they did in the corresponding unisensory condition. In retrospect, this should not be have been surprising because no instructions were given to the subjects about which cue would be more effective in the coming block of trials and they could only learn this through experience during the block. Most often, interference occurred for Subjects 1 and 3 (7 out of 8 cases) who had very long RTs under the tactile-only condition in Experiment 1; this led to high inaccuracy values for the auditory cues that were chosen for Experiment 2. That is, in some conditions, these subjects received an ineffective auditory cue as well as an ineffective tactile cue. This may have confused the subjects and made it more difficult to determine which cue was more accurate if they were trying to focus on a single cue, or made it more difficult to combine (e.g., average) the two locations if they were trying to form a single multisensory cue.

In any case, it seems likely that subjects may have processed the cues differently in the conditions that led to enhancement compared to those that led to interference. If so, it may be unreasonable to represent both processes with a single linear function. Because we are primarily interested in enhancement for this thesis we reanalyzed the data, eliminating points that showed negative enhancement (i.e., interference). The resulting fitted lines can be seen in Figure 4.5 and Figure 4.6. The slopes of the lines are shallower and the y intercepts are lower (suggesting less enhancement, 317 ms and 573 ms, with the 30 and 100 distractors, respectively) than those in previous figures. That is, the slopes, and thereby the intercepts, of the fitted lines in Figures 4.3 and 4.4 were substantially influenced by the points showing interference and so the fitted functions may have over estimated the amount of enhancement that should be expected when the cues are well matched.

Because of the small number of data points, the substantial performance differences among subjects, and the uncertainty about how best to combine data showing enhancement with data showing interference, it is difficult to be confident about the exact level of en-

30

Figure 4.3: Multisensory enhancement plotted as a function of the difference between the auditory cue only RT and the tactile cue only RT for the 30-distractor case. Similar symbols represent data for one subject. On the ordinate positive values show enhancement and negative values show interference. Dashed lines were fit using *fminsearch* in MATLAB R2012a Student.
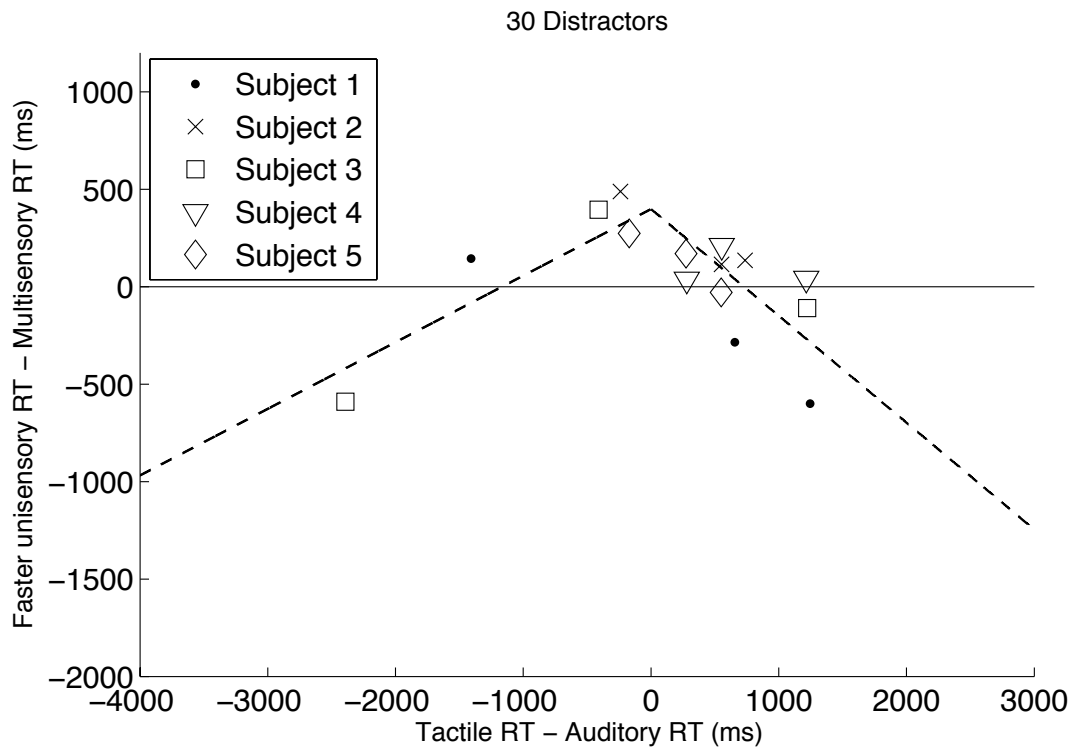
Figure 4.4: Multisensory enhancement plotted as a function of the difference between the auditory cue only RT and the tactile cue only RT for the 100-distractor case. Similar symbols represent data for one subject. On the ordinate positive values show enhancement and negative values show interference. Dashed lines were fit using *fminsearch* in MATLAB R2012a Student.

hancement that should be expected when the tactile and auditory cues are equally effective. That said, the data indicate that enhancement should be anticipated when the difference in effectiveness between tactile and auditory cues is small, that greater levels of enhancement are more likely when the visual search is more difficult (i.e., when there are more distractors), and that for our task enhancement of between 300 ms and 800 ms is expected when the two cues are equally effective.

Figure 4.5: Multisensory enhancement (positive only values) plotted as a function of the difference between the auditory cue only RT and the tactile cue only RT for the 30-distractor case. The symbols are the same and the lines were fit in the same way as in Figures 4.3 and 4.4. Dashed lines were fit using *fminsearch* in MATLAB R2012a Student.

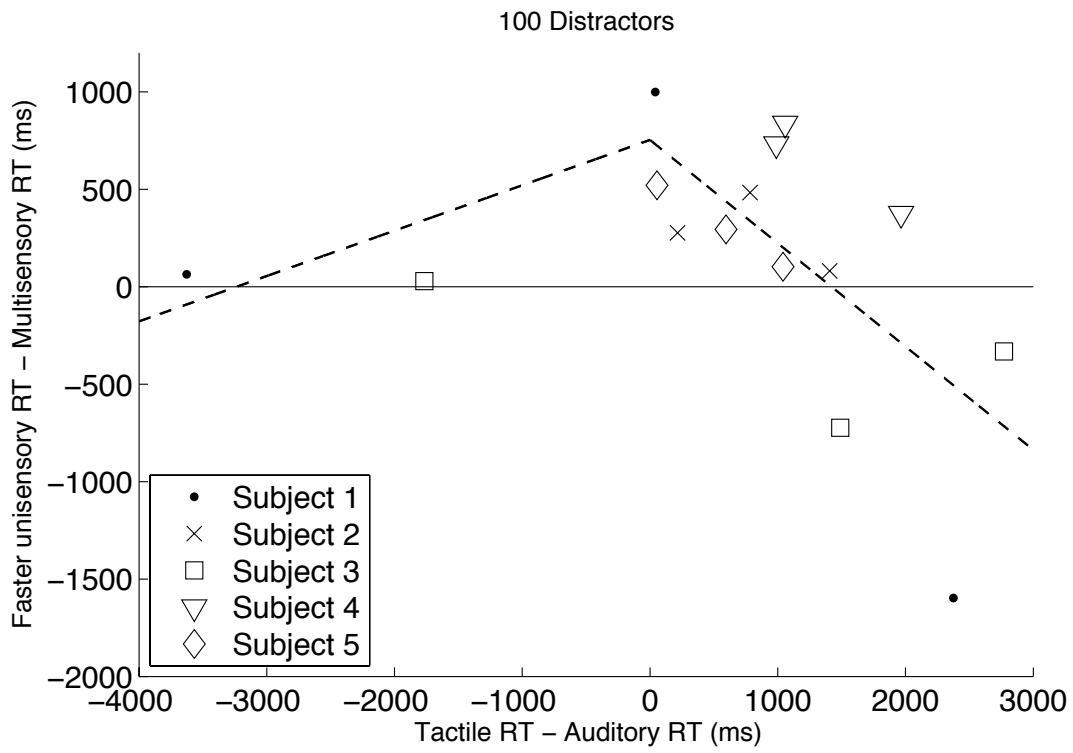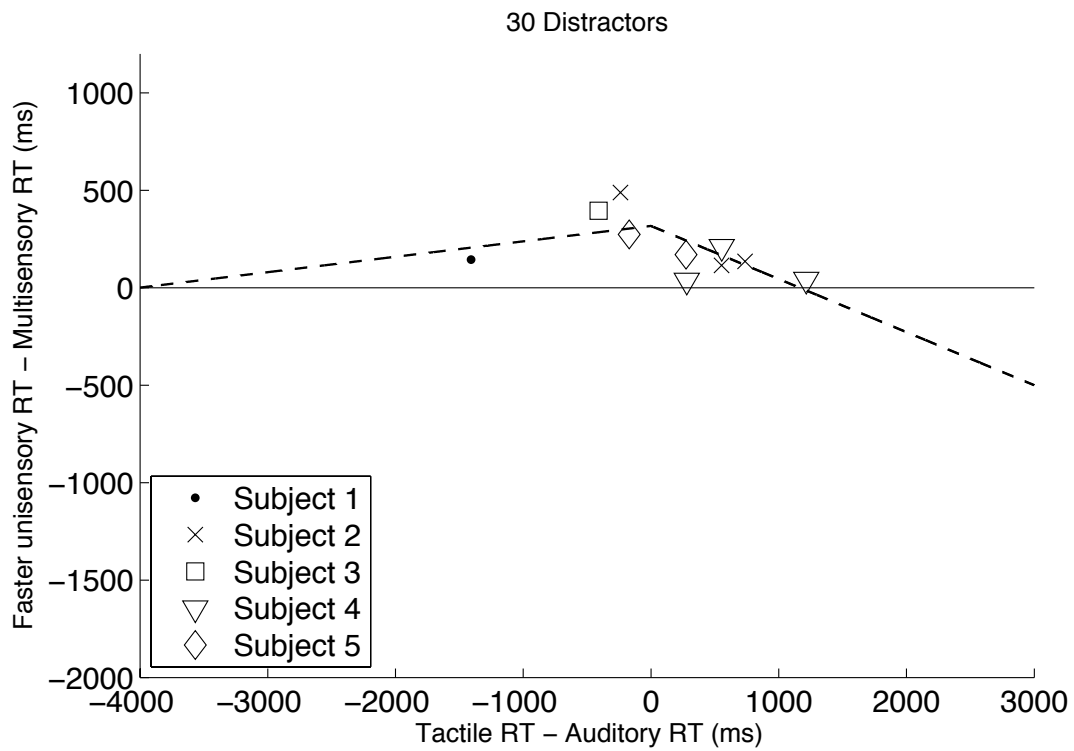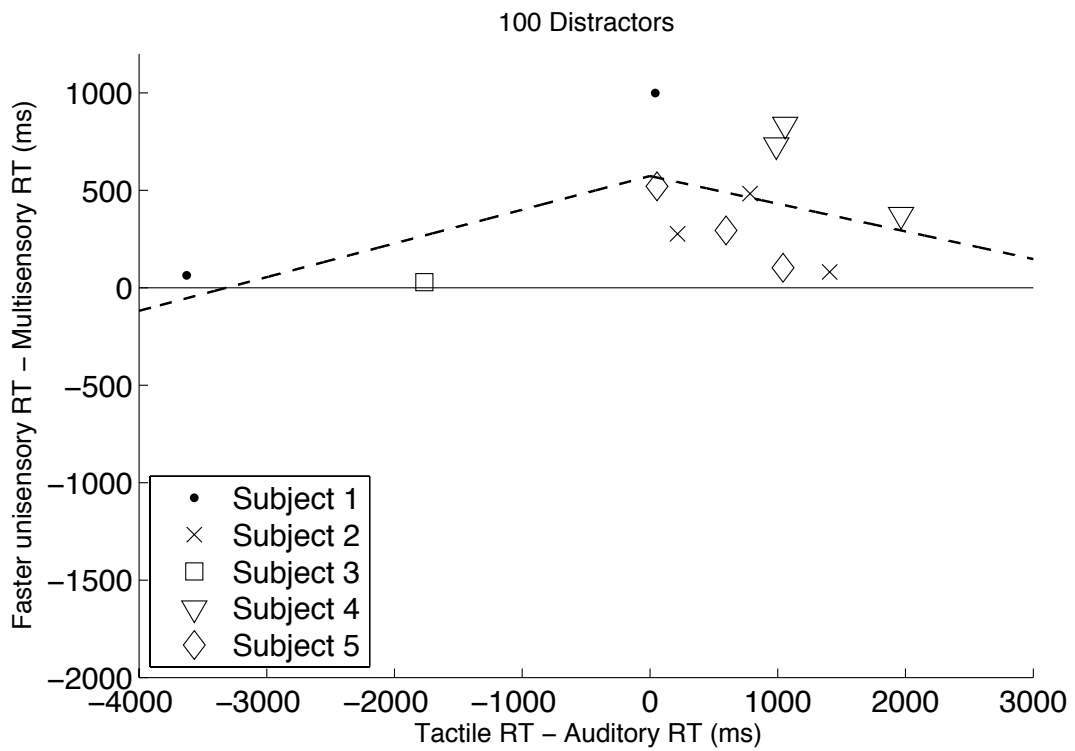Figure 4.6: Multisensory enhancement (positive only values) plotted as a function of the difference between the auditory cue only RT and the tactile cue only RT for the 100-distractor case. The symbols are the same and the lines were fit in the same way as in Figures 4.3 and 4.4. Dashed lines were fit using *fminsearch* in MATLAB R2012a Student.

GENERAL DISCUSSION

This study examined whether the simultaneous presentation of spatialized auditory and tactile cues that were well-matched in terms of effectiveness would increase the speed of visual search relative to the speed when either cue was presented alone. We originally planned to compare enhancement in the A_T + T condition to the A_{T-25} + T, and A_{T+25} + T conditions, but this was not viable because the A_T and T conditions were not well matched. However, because some auditory cue only conditions were better matched to the T condition than others, we sought to determine whether or not the level of multisensory enhancement changed as a function of the effectiveness mismatch between the auditory-only and tactile cues. Overall our results suggest that when the effectiveness mismatch was small we were more likely to see multisensory enhancement was likely and when the effectiveness mismatch was large multisensory interference was more likely.

Recall that Mateo et al. (2012) did not find multisensory performance to be better than performance with the auditory cue alone. This is not surprising because they did not try to match their auditory and tactile cues in effectiveness and in fact their auditory cue was much more effective than their tactile cue. They suggest that had their cues been matched in effectiveness they would likely have seen multisensory enhancement. Our 0° auditory cue inaccuracy condition and tactile cue condition in Experiment 1 were equivalent to their auditory and tactile conditions and yielded similar results in the 30-distractor case. In our study the average RTs for the tactile cue only and 0° auditory cue inaccuracy conditions were 2928 ms and 1665 ms, respectively; in Mateo et al. (2012) they were 2689

ms and 1441 ms, respectively. So, although we did not consider the multisensory condition that combined these two cues, our results suggest that no enhancement would have been observed. However, our results also indicate that had Mateo et al. (2012) included multisensory conditions with auditory and tactile cues that were well matched, they would have observed enhancement. Other studies (e.g., Hancock et al., 2013; Oskarsson et al., 2012) that presented better matched cues than Mateo et al. have also found multisensory enhancement. For example, Hancock et al. (2013) provided cues that directed subjects to a small region of space that subjects then had to search through in order to find the target. Although neither the auditory nor tactile cue was very effective in the sense that it did not direct subjects to the precise location of the target, both were similarly effective and so when combined resulted in a performance benefit.

The enhancement observed in this study and Hancock et al. (2013), as well as the lack of enhancement found in Mateo et al. (2012), are consistent with the inverse-effectiveness principle of multisensory processing. Recall that the inverse-effectiveness principle suggests that there will be little enhancement for a multisensory cue if one of the cues is very effective on its own (Stein et al., 2001). Physiological studies (see Stein et al., 2001) have shown that highly effective stimuli (similar to the auditory cues in Mateo et al., 2012) result in large neural responses, which are not changed much by information provided by an additional stimulus. However, two stimuli that are not highly effective on their own (similar to the tactile cues and some auditory cues in this study and in Hancock et al., 2013) can combine to produce a larger neural response than the sum of their individual responses (Stein et al., 2001).

As the effectiveness mismatch increased we expected there would be a point where there would no longer be enhancement, at which point performance would be the same as the more effective of the two cues. We further expected that greater mismatches would also yield multisensory performance equal to performance with the better of the two unisensory cues. However, for two subjects this was not true; with greater mismatches performance

continued to decline, showing negative enhancement (i.e., interference). In some cases, this may have been due to subjects receiving auditory and tactile cues that were both relatively ineffective (even though with these large mismatches one of the cues, usually the auditory cue, was substantially more effective). In these situations, subjects may have been confused, unable to determine which cue was more effective or uncertain about how best to combine cues that would separately cause the subject to orient to quite disparate locations. Recall that the instructions simply told the subjects to utilize both sources of information, without telling them how to combine the information or which cue was likely to be more accurate. The subjects had to learn these things through experience within and between blocks. So, it is not surprising that they may have processed the two cues in a non-optimal way at least some of the time.

5.1 Practical Significance

In addition to the basic science implications of this work, it is not difficult to imagine the utility of a multisensory display in the context of operational environments in which speeding visual search is critical (e.g., visually acquiring an incoming aircraft; finding a particular icon on an AWACS display, determining the location of friendly or hostile personnel in the battlefield, etc.) and an effective unisensory cue is likely to be degraded. Although auditory cues alone have been shown to be very effective at speeding search in laboratory environments it is likely that in the real world environmental noise or unreliable communication channels will make them less effective. In situations where an auditory cue might be degraded, our results suggest that also presenting a similarly effective tactile cue can benefit search times. However, in some cases in which the effectiveness of the cues is substantially different, presenting a tactile cue can result in interference. Based on our data it is unclear if interference occurred for some subjects because they were relatively poor unisensory cue performers, because they had high auditory cue inaccuracy values, because

they had cues mismatched in effectiveness, or some combination of these factors. Note however, in a few cases one of the three better performing subjects had larger mismatches than one of the two poorer performing subjects and in those cases the better performing subject did not experience interference and the poorer performing subject did. So it may be that the interference we observed is associated with the subject and their abilities or their strategy that they use, rather than with the stimulus conditions per se.

In a situation where the speed of visual search is critical, experiencing interference could be devastating. One solution to the problem of interference may be simply to not present the worse unisensory cue. However, it is possible that an auditory cue could be degraded by environmental noise or technological malfunctions, and so the soldier might not receive useful spatial information about the sniper's location. Another solution may be to present the tactile cue only when the primary cue might be degraded (e.g., only if the auditory signal-to-noise ratio is low or front/back confusion is likely). This could be achieved by a "smart" multisensory display that would activate a tactile cue when the signal-to-noise ratio reached a certain level or it could be achieved by a soldier activating the tactile cue only in situations in which they feel it might be useful. Because auditory cues are likely to be more effective than tactile cues for left/right and up/down cueing in general and tactile cues are likely to be more effective for front/back cueing, another solution may be to present a tactile cue that only cues the front or rear hemispheres. Last, it may be possible simply to train away the interference. It could be the case that poor performing subjects could be trained to develop better strategies.

5.2   Future Work

In order to design effective multisensory displays, there are a number of outstanding questions that should be answered. Most importantly, a better understanding of the nature of interference is needed. From our results it is unclear whether interference is associated with

poor performing subjects, poorly matched cues, or high auditory inaccuracy values. In order to further examine these questions future work should include a larger sample, which would likely include "good" and "bad" unisensory performers. Then the relative effectiveness of the cues could be manipulated to determine if good performers would experience interference with poorly-matched cues and if bad performers would not experience interference with well-matched cues.

It would also be important to determine if good performers and bad performers use different strategies. This could be done by looking at head movement data within a trial, which might determine whether the subject tries to find the target by first moving towards one of the cues over the other, the subject switches between the two cued locations, or the subject combines the two cued locations into a different, third, location.

Understanding the strategies that good performers and bad performers use could also be important if interference can be trained away. Bad performers may be able to overcome interference with additional training, in which they would develop more optimal search strategies. Additionally, if good performers and bad performers use different strategies an experimenter could suggest that bad performers use the strategies of good performers in order to improve the speed of search.

If it is not possible or feasible to train away interference, it may be important to evaluate other types of multisensory displays. For example, it is unclear whether a multisensory display that consists of an auditory cue as well as a front/back tactile cue would benefit search times and/or not result in interference. Additionally, an adaptive multisensory display that activates a tactile cue in response to environmental noise or the user's preference should be evaluated.

Last, it should be noted that future experiments that attempt to estimate cue inaccuracy levels as we did should measure more levels of inaccuracy, because the relation between inaccuracy and RT appears to be curvilinear.

CONCLUSIONS

In this study we observed both enhancement and interference for multisensory cues relative to their component unisensory cues alone. Enhancement was more likely to occur when the auditory and tactile cues were closely matched in effectiveness and interference was more likely to occur when auditory and tactile cues were not closely matched. Because the nature of the interference is not clear, future work is needed to determine whether the interference found in this study was a result of poor performing subjects, poorly matched cues, high inaccuracy values, or some combination of these factors. Although more questions remain it seems that effectively matched auditory and tactile cues would benefit search in operational environments.

Bibliography

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

Ephrem, A. J., Finomore, V. S., Gilkey, R. H., Newton, R. M., Romigh, G. D., Simpson, B. D., Brungart, D. S., and Cowgill, J. L. (2008). An "audio annotation" technique for intuitive communication in spatial tasks. *In Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 52:2127–2131.

Flanagain, P., McAnally, K. I., Martin, R. L., Meehan, J. W., and Oldfield, S. R. (1998). Aurally and visually guided visual search in a virtual environment. *Human Factors*, 40:461–468.

Foster, D. H. and Ward, P. A. (1991). Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 243:75–81.

Hancock, P. A., Mercado, J. E., and Van Erp, J. B. (2013). Improving target detiction in visual search through the augmenting multi-sensory cues. *Ergonomics*, 56:729–738.

Karnath, H. O., Schenkel, P., and Fischer, B. (1991). Trunk orientation as the determining factor of the 'contralateral' deficit in the neglect syndrome and as the physical anchor of the internal representation of body orientation in space. *Brain*, 114:1997–2014.

Lindeman, R. W., Yanagida, Y., Sibert, J. L., and Lavine, R. (2003). Effective vibrotactile cueing in a visual search task. *Proceedings of the ninth IFIP TC13 international conference on Human-computer interaction*, (89-96).

Mateo, J. C., Simpson, B. D., Gilkey, R. H., Iyer, N., and Brungart, D. S. (2012). Spatial multisensory cueing to support visual target-acquisition performance. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 56:1312–1316.

McGrath, B., Estrada, A., Braithwaite, M. G., Raj, A. K., and Rupert, A. H. (2004). Tactile situation awareness system flight demonstration final report. Technical report, Naval Aerospace Medical Research Lab, Pensacola, Fl.

McIntire, J. P., Havig, P. R., Watamaniuk, S. N., and Gilkey, R. H. (2010). Visual search performance with 3-d auditory cues: Effects of motion, target location, and practice. *Human Factors*, 52:41–53.

Nelson, W. T., Hettinger, L. J., Cunningham, J. A., Brickman, B. J., Haas, M. W., and McKinley, R. L. (1998). Effects of localized auditory information on visual target detection performance using a helmet-mounted display. *Human Factors*, 40:452–460.

Ngo, M. K. and Spence, C. (2010). Auditory, tactile, and multisensory cues facilitate search for dynamic visual stimuli. *Attention, Perception, and Psychophysics*, 72:1654–1665.

Oskarsson, P. A., Eriksson, L., and Carlander, O. (2012). Enhanced perception and performance by multimodal threat cueing in simulated combat vehicle. *Human Factors*, 54:122–137.

Perrott, D. R., Cisneros, J., McKinley, R. L., and D'Angelo, W. R. (1996). Aurally aided visual search under virtual and free-field listening conditions. *Human Factors*, 38:702–715.

Perrott, D. R., Saberi, K., Brown, K., and Strybel, T. Z. (1990). Auditory psychomotor coordination and visual search performance. *Perception and Psychophysics*, 48:214–226.

Perrott, D. R., Sadralodabi, R., Saberi, K., and Strybel, T. Z. (1991). Aurally aided visual search in the central visual field: Effects of visual load and visual enhancement of the target. *Human Factors*, 33:389–400.

Rudmann, D. S. and Strybel, T. Z. (1999). Auditory spatial facilitation of visual search performance: Effect of cue precision and distractor density. *Human Factors*, 41:146–160.

Rupert, A. H., McGrath, B., and Griffin, M. (2003). A tool to maintain spatial orientation and situation awareness for operators of manned and unmanned aeriel vehicles and other military motion platforms. Technical report, NAVAL AEROSPACE MEDICAL RESEARCH LAB, Pensacola, Fl.

Santangelo, V. and Spence, C. (2007). Multisensory cues capture spatial attention regardless of perception load. *Journal of Experimental Psychology: Human Perception and Performance*, 33:1311–1321.

Stein, B. E., Wallace, M. T., and Stanford, T. R. (2001). *Brain mechanisms for synthesizing information from different sensory modalities*. Blackwell Publishers.

Tan, H. Z., Gray, R., Young, J. J., and Traylor, R. (2003). A haptic back display for attentional and directional cueing. *Haptics-e*, 3(1):1–20.

Treisman, A. M. and Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12:97–136.

Van Erp, J. B., Groen, E. L., and van Veen, H. A. (2006). A tactile cockpit instrument sup-

ports the control of self-motion during spatial disorientation. *Human Factors*, 48:219–228.

Vu, K. P. L., Strybel, T. Z., and Proctor, R. W. (2006). Effects of displacement magnitude and direction of auditory cues on auditory spatial facilitation of visual search. *Human Factors*, 48:587–599.

Welch, R. B. and Warren, D. H. (1986). *Intersensory interactions*. John Wiley and Sons, New York.

Wolfe, J. M. (1998). Visual search. *Attention*, 1:13–73.

Wolfe, J. M., Kluender, K. R., Levi, D. M., Bartoshuk, L. M., Herz, R. S., Klatzky, R. L., Lederman, S. J., and Merfeld, D. M. (2009). *Sensation and Perception*. Sinauer Associates, Sunderland, Ma.