Active heading control in simulated flight based on vertically extended contours

ROBERT PATTERSON

Washington State University, Pullman, Washington

GEORGE A. GERI

Link Simulation and Training, Mesa, Arizona

BRIAN P. DYRE

University of Idaho, Moscow, Idaho

BYRON J. PIERCE

Air Force Research Laboratory, Mesa, Arizona

SHAMA C. AKHTAR

The Boeing Company, Mesa, Arizona

CHRISTINE M. COVAS

Link Simulation and Training, Mesa, Arizona

and

WILLIAM MORGAN

The Boeing Company, Mesa, Arizona

In two experiments, we manipulated the properties of 3-D objects and terrain texture in order to investigate their effects on active heading control during simulated flight. Simulated crosswinds were used to introduce a rotational component into the retinal flow field that presumably provided the visual cues used for heading control. An active control task was used so that the results could be generalized to real-world applications such as flight simulation. In Experiment 1, we examined the effects of three types of terrain, each of which was presented with and without 3-D objects (trees), and found that the presence of 3-D objects was more important than terrain texture for precise heading control. In Experiment 2, we investigated the effects of varying the height and density of 3-D objects and found that increasing 3-D object density improved heading control, but that 3-D object height had only a small effect. On the basis of these results, we conclude that the vertical contours improved active heading control by enhancing the motion parallax information contained in the retinal flow.

One of the most important tasks that a pilot performs is the control of heading. Of the various visual cues that determine heading, one important cue is the optic flow produced when an observer moves through a textured environment (Gibson, 1950). When an observer moves forward and maintains fixation, the pattern of relative motion in the retinal images contains a point from which the relative motion vectors radiate outward. This point is called the *focus of expansion* and corresponds to the location

in the environment toward which the observer is heading. The pattern of relative motion just described may be used by individuals to steer vehicles, direct locomotion, and control postural balance. For example, this pattern can support accurate heading judgments when observer motion is simulated through a cloud of dots that provides minimal object-based cues (Warren, 1998). It can also be used for postural control during walking and can be sensed from retinal regions up to an eccentricity of 90° (Bardy, Warren, & Kay, 1996, 1999). Finally, the pattern of relative motion in the retinal images has been shown to selectively activate specialized regions in the extrastriate cortex (see, e.g., Vanduffel et al., 2002).

As Warren (1998) discussed in a review paper, when an observer is traveling on a straight path and does not make eye or head movements, translational information contained in the retinal flow is sufficient for recovering heading direction. However, when an observer moves on

The research described here was performed at the Air Force Research Laboratory, Mesa, AZ, under Contract F41624-97-D-5000 to Link Simulation and Training, a division of L-3 Communications Corp. R.P. was supported by a Summer Faculty Fellowship from AFOSR. Correspondence concerning this article may be addressed to R. Patterson, Department of Psychology, Program in Neuroscience and Center for Integrated Biotechnology, Washington State University, Pullman, WA 99164-4820 (e-mail: rpatter@mail.wsu.edu).

a curved path, or on a straight path while making eye or head movements (as occurs when fixating a spot on the ground while translating forward), the flow pattern on the retina will be a combination of a translational component and a rotational component that interferes with the recovery of heading from retinal flow. In order for heading to be recovered from the translational component, an observer must decompose the retinal flow into its rotational and translational components. Warren and colleagues (e.g., Li & Warren, 2000; Warren, 1998) have discussed several theoretical approaches for understanding how the visual system performs such a decomposition. It now seems likely that both extraretinal information, such as proprioceptive information about eye and head movements (see, e.g., Royden, Crowell, & Banks, 1994), and retinal flow information, such as motion parallax (see, e.g., Li & Warren, 2000, 2002), are used by the visual system to recover the translational component of optic flow. Concerning retinal flow, an analysis by Longuet-Higgins and Prazdny (1980; see also Rieger & Lawton, 1985) suggests that motion parallax can be used to recover heading via a process of image velocity subtraction, which serves to eliminate the rotational component and produce a radial pattern from which heading can be determined.

Most of the heading studies have involved the passive perception of heading rather than its active control. In many heading studies (e.g., Li & Warren, 2000; Royden et al., 1994), optic flow is presented briefly and then observers are required to judge the future path they would have taken relative to a probe stimulus. Typically, two conditions are compared: an actual-rotation condition, wherein a moving fixation point is used to induce an eye rotation, and a simulated-rotation condition, wherein the optical information consistent with rotation is presented in the retinal flow while a stationary fixation point is used to minimize eye rotations. If the actual-rotation condition produces more accurate heading judgments than the simulated-rotation condition, then it is inferred that extraretinal signals are involved in heading perception under the actual-rotation condition. In many studies, more accurate heading perception is found in the actual-rotation condition than in the simulated-rotation condition, especially when optical information for cues such as motion parallax is sparse (see, e.g., Li & Warren, 2000; Royden et al., 1994).

Although the studies just described provide important information about heading perception, their results do not necessarily generalize to the active control of heading. This would be true if perceptual judgments and control of action were based on different information or implemented by different streams of processing, and there is a body of evidence that just such anatomically and functionally distinct cortical systems may exist (e.g., Aglioti, DeSouza, & Goodale, 1995; Bridgeman, Gemmer, Forsman, & Huemer, 2000; Goodale, Milner, Jakobson, & Carey, 1991; see the review by Milner & Goodale, 1995). This work suggests that one cortical system (the parvocellular-dominated ventral stream) mediates conscious perception that is linked to memory and another system (the magnocellular-dominated dorsal stream) mediates visuo-

motor control. Thus, it is possible that heading perception may be linked to activation of the ventral stream, whereas heading control may be linked to activation of the dorsal stream. Although these anatomical and functional claims remain under active investigation, they provide impetus not to assume that active heading control behaves in the same way as heading perception.

Despite the importance of examining active heading control when a rotational component is present in the retinal flow pattern, few studies have done so. In one of the few, Rushton, Harris, and Wann (1999) reported that depth information (provided either by static perspective cues or by binocular disparity) did not improve active heading control. In another, Frey and Owen (1999) found that adding motion parallax cues produced little or no improvement in active heading control. However, as discussed by Li and Warren (2002), the simulated-rotation rates in those studies were relatively low (less than 0.6°/sec). Li and Warren (2002) investigated active heading control under relatively high rates of simulated rotation (up to 5°/sec) and found that the presence of motion parallax cues, together with vertical reference objects, improved heading control. On the basis of these results, they proposed that instantaneous heading can be determined from motion parallax in the retinal flow alone, whereas the path of simulated self-motion is determined by updating over time the heading relative to objects in the scene.

Given that motion parallax information is an important visual cue for the control of heading during the presence of eye or head movements, the presence or enhancement of such information may be an important consideration in the design of databases used to simulate navigation during low-altitude flight. At low altitude, many environmental features would be present for creating cues for motion parallax. Thus, it is important to determine which features of the flight environment convey useful motion parallax information to pilots.

The purpose of the present study was to investigate active heading control under conditions for which a rotational component was introduced into the retinal flow field. This component was added by using simulated crosswinds that acted on a simulated aircraft as it was steered toward a distant target. The crosswinds, together with the observer's steering responses, resulted in a curved path whose direction and radius of curvature were constantly changing and which together defined the rotational component of the retinal flow. Across two experiments, we manipulated the properties of 3-D objects and terrain texture to examine their effects on heading control during simulated translation and rotation.

GENERAL METHOD

Observers

Seven observers served in each experiment, and only 1 observer participated in both experiments (the results of the second experiment look the same with this observer's data removed). All the observers (age range, 21–51 years) had normal or corrected-to-normal acuity in both eyes, normal binocular vision, normal color vision, and normal phoria (Optec Vision Tester, Stereo Optical, Chicago). All but

2 observers were paid for their participation, and all gave informed consent. None of the observers had actual flying experience.

Stimuli and Apparatus

The terrain imagery was generated using commercial database development software (World Perfect 2.0, MetaVR, Brookline, MA). To create the terrain images, a flat database was constructed whose dimensions were 10×15 km. The database was subdivided into different terrain sections, each measuring 5×5 km. The different sections corresponded to the various experimental conditions described in the Method section of each of the two experiments.

Flight over the terrain was simulated using a high-performance PC-based runtime system (Virtual Reality Scene Generator, MetaVR). The terrain imagery was displayed on three channels of a wide-field, rear-projection system (see Best, Wight, & Pepler, 1999) using full-color CRT projectors (Barco, Model 808). The imagery subtended 240° horizontally $\times$ 63° vertically at a viewing distance of 94 cm. Each display channel provided $1,600 \times 1,200$ pixels at an update rate of 60 Hz. A joystick interfaced to the PC was used by the observer to control heading. The joystick could be rotated up to $\pm 15^\circ$; full rotation corresponded to a $\pm 25^\circ/\text{sec}$ simulated rotation of the visual imagery.

At the beginning of each trial, a target (a red building) was positioned near the horizon and to either the left or the right of the center of the scene. The size of the target was $23 \times 23 \times 60$ m, which rendered it clearly visible to the observer at the start of each trial. The distance from the observer's starting point to the target was 1.8 km.

Control Task

Figure 1 shows the relationships among the observer's aircraft, the target (T), and the crosswind at various times during a hypothetical trial. Figure 1A depicts the starting position of an aircraft flying toward T. The arrow at the front of the aircraft represents its velocity (i.e., its direction and speed). The arrow to the left of the aircraft represents the velocity of a crosswind that is at a right angle to the aircraft's velocity. The letter X to the front and left of the aircraft shows a combined velocity vector that represents the sum of the aircraft and crosswind velocities. If the observer keeps the aircraft pointed at the target but otherwise does not adequately compensate for the crosswind, the aircraft will eventually reach the target, but it will travel along a curved and indirect path, which will introduce rotation into the retinal flow (see Figure 1B).

Figure 1C depicts a situation in which the observer, midway through the flight, compensates fully for the crosswind by pointing the aircraft to the right of the target (i.e., into the crosswind), which directs the combined velocity vector toward the target. As a result, the aircraft moves in a more direct path toward the target, thus reducing the rotational component of the retinal flow. The control task summarized in Figure 1C is referred to here as the *constant-crosswind task*. Figure 1D depicts the situation in which the magnitude and direction of the crosswind is varied over time, whereas the angle of the crosswind relative to the long axis of the aircraft remains at 90° . This causes the direction of travel to shift continually to the left or right as the observer translates forward. The control task summarized in Figure 1D is referred to here as the *variable-crosswind task*.

For the variable-crosswind task, the magnitude of the simulated crosswind varied over the course of each trial according to a sum of three sinusoids. The first of the three sinusoids had a frequency (F1) of 0.2–0.5 Hz and an amplitude (A1) of 20–25 m. The second sinusoid had a frequency that ranged from $1.6 \times F1$ to $4.0 \times F1$ and an amplitude that ranged from $0.32 \times A1$ to $0.75 \times A1$. The third sinusoid had a frequency that ranged from $2.0 \times F1$ to $12.5 \times F1$ and an amplitude that ranged from $0.24 \times A1$ to $0.80 \times A1$. The phase angle of each waveform was varied from 0° to 359° . The actual frequencies, amplitudes, phase angles, and direction (i.e., leftward or rightward) used on any given trial were pseudorandomly determined by computer (the direction of the crosswind also varied randomly within each trial).

Figure 1. A depiction of the heading tasks used in the present study. (A) The starting position of an aircraft flying toward a target (T). (B) With a leftward crosswind present, if the aircraft is pointed at the target it will drift to the left as it travels forward. (C) The aircraft will travel a direct path to the target if it is steered into the crosswind. Pointing the aircraft into the crosswind makes the vector sum of the aircraft and crosswind velocities point toward the target. The control task summarized in panel C is called the *constant-crosswind task*. (D) The situation in which the magnitude of the crosswind varies over time. In this case, the observer must continuously vary the direction of the velocity of the aircraft to remain on a direct path to the target. The control task summarized in panel D is called the *variable-crosswind task*.

For the constant-crosswind task, the observer could have initially detected heading direction on the basis of motion parallax information, noted the position of the image of the target relative to the center of the display screen, and then maintained that position throughout the remainder of the trial. For the variable-crosswind task, however, the observer could not have maintained a direct path to the target simply by aligning the target with a position on the display screen. Rather, due to the continuous rotation of the scene produced by the varying crosswind, the observer had to continuously vary the directional component of the aircraft's velocity (i.e., its rotation about the vertical axis, also known as *yaw*) to keep the aircraft moving in a direct path toward the target. In order to do this, the observer would have to rely on something other than position information, such as motion parallax. Note that the rotation rates to which our observers were exposed depended on their responses on our steering task because it was a first-order negative feedback control task over which the observers had control.

In Experiment 1, the observers performed both the constant-crosswind and variable-crosswind heading tasks. In Experiment 2, the observers performed only the variable-crosswind task. Both tasks

are simplified versions of the control tasks actually performed by pilots. For example, flying an aircraft involves third- and fourth-order control responses (Roscoe, Eisele, & Bergman, 1980). By using a simplified first-order control task, we were able to focus on the perceptual cues that affect performance with synthetic vision displays, and we could use observers with no prior flight experience.

Procedure

Prior to each trial, for several seconds the observer viewed the terrain area over which flight would be simulated. Motion was then simulated at an altitude of 30 m and a ground speed of 90 m/sec (175 knots), and the observer was instructed to fixate the target and steer toward the target by the most direct (i.e., shortest) path. The simulated wind gusts required the observer to compensate by changing yaw in order to sideslip, or "crab," the vehicle toward the target. The duration of each trial was 20 sec. At the end of each trial, the simulated flight stopped short of the target by several meters. Root mean square (RMS) error was computed from the deviation in degrees to the right or left of a straight path toward the target as sampled from 80 measurements obtained over the course of each 20-sec trial (i.e., RMS error sampled every 250 msec).

In each experiment, data were collected from each observer in four to six experimental sessions. During each session, the observer first performed one block of warm-up trials (identical to the data trials; see below). Next, one or two blocks of data trials were run. In Experiment 1, each block consisted of 6 repetitions of each of 6 conditions, for a total of 36 trials per block. In Experiment 2, each block consisted of 5 repetitions of each of 21 conditions, for a total of 105 trials per block. The order in which conditions were tested within each block was randomized. For each experiment, four blocks of data trials were run over the course of several experimental sessions. Thus, a total of 24 trials per condition per observer was performed in Experiment 1 (except for observer G.S. in Experiment 1, for whom a total of 12 trials was run for each condition), and a total of 20 trials per condition per observer was performed in Experiment 2. Prior to data collection for each experiment, each observer received sufficient practice trials (four to six blocks of trials, as described above) to reach asymptotic performance for each condition.

EXPERIMENT 1

In this experiment, we investigated three types of terrain: gray untextured, random noise, and geospecific (simulated real world). Each type of terrain was presented with and without 3-D objects in the form of trees (see Figure 2).

Method

The gray untextured terrain was a homogeneous gray field. It contained no static relative depth cues based on perspective, no flow field, and no motion parallax. This terrain was included in order to help assess whether a simple lateral displacement of the target could function as a heading cue. The random-noise terrain was derived from a random-noise texture consisting of 512 texture elements/m², which had a low-pass spatial frequency spectrum. The geospecific terrain was a rendering of a satellite image of an area of Nellis Air Force Base, NV. Both the random-noise and the geospecific terrain types contained static perspective cues for relative depth and provided both a flow field and motion parallax information. All three terrain types had an average luminance of about 3 fL.

Each terrain type was presented both with and without 3-D objects in the form of trees. The simulated trees were 7 m tall and 1 m wide, and their density was 100 trees/km². The trees provided object perspective, increased motion parallax closer to eye level, dynamic occlusion and disclosure, and multiple reference objects. Trees were included in order to examine whether objects extending above the ground plane would provide enhanced motion parallax cues and/or reference cues that would improve performance.

Results

Figure 3A shows the results for the constant-crosswind task. Average RMS error in heading control is plotted for the three terrain types presented with and without trees. It can be seen that the gray untextured terrain with no trees produced the worst performance (greatest RMS error),

Figure 2. Photographs of the six terrain images used in Experiment 1. Top row, from left to right: gray untextured, low-frequency random noise, and geospecific (real world), all with no trees present. Bottom row, from left to right: gray untextured, low-frequency random noise, and geospecific (real world), all with trees present. The target, which was a red building, is not shown in the images.

whereas the other conditions produced better, and nearly equivalent, performance.

Figure 3B depicts the results for the variable-crosswind task. Average RMS error is once again plotted for the three terrain types, each again tested with and without trees. As was the case for the constant-crosswind task, the gray untextured terrain with no trees produced the worst performance. The addition of trees improved performance for all terrain types.

A within-subjects ANOVA revealed significant main effects of terrain [$F(2,10) = 13.3, p < .02$], trees [$F(1,5) = 48.6, p < .002$], and task [$F(1,5) = 284.0, p < .001$]. Moreover, the interaction between terrain and trees was significant [$F(2,10) = 10.7, p < .03$], as was the interaction between trees and task [$F(1,5) = 36.6, p < .01$].

To probe the interaction between terrain and trees, we performed post hoc tests, based on Tukey's honestly significant difference (HSD) statistic, of the effect of terrain both with and without trees. Without trees, the effect of terrain was significant [$F(2,12) = 13.8, p < .01$]. An a posteriori t test showed that the gray untextured terrain produced significantly worse performance than both the low-frequency noise terrain [$t(6) = 3.7, p = .01$] and the geospecific (real-world) terrain [$t(6) = 3.8, p < .01$], but there was no significant difference between the low-frequency noise terrain and the geospecific terrain. With trees present, the effect of terrain was not significant ($p > .05$). We also performed post hoc tests of the effects of trees for the different kinds of terrain. This analysis indicated that the presence of trees produced significantly better performance than the absence of trees for the untextured terrain [$F(1,6) = 29.9, p < .01$]. A similar but much smaller trend was observed for the low-frequency noise terrain [$F(1,6) = 42.1, p < .01$] and for the geospecific terrain [$F(1,6) = 21.1, p < .01$].

To probe the interaction between trees and task, we performed post hoc tests on the difference in heading control between the constant-crosswind task and the variable-crosswind task, each performed with and without trees. This analysis indicated that the variable-crosswind task resulted in significantly worse performance than the constant-crosswind task with trees absent [$F(1,6) = 488, p < .001$]; a similar but smaller trend was observed with trees present [$F(1,6) = 37.8, p < .01$]. Finally, we performed post hoc tests of the effect of trees for the two kinds of task. For the constant-crosswind task, the presence of trees produced a small but significant improvement in performance in comparison with the absence of trees [$F(1,6) = 16.4, p < .01$]. For the variable-crosswind task, the presence of trees produced a significant and much larger improvement in performance [$F(1,6) = 106, p < .001$].

The results of Experiment 1 show that the best heading control is obtained with 3-D objects (trees) even in the absence of terrain texture. Thus, we conclude that the vertical extension of the 3-D objects above the ground plane enhanced heading control. Furthermore, we propose that the mechanism for this improvement is an enhanced

Figure 3. (A) Root mean square (RMS) error in the constant-crosswind task for three terrain types presented with and without trees. Error bars represent ± 1 standard error of the mean for 7 observers. (B) RMS error in the variable-crosswind task for three terrain types presented with and without trees. Error bars represent ± 1 standard error of the mean for 7 observers.

motion parallax cue mediated by the vertical extension of the 3-D objects.

Note that in Experiment 1 the addition of trees shifted the Fourier energy in the scenes to higher spatial frequencies, and that shift might have improved performance independently of the vertical contours associated with the trees. We assessed this possibility in two control experiments, each conducted with 8 observers. In the first experiment, we increased the texture density of the noise terrain, which shifted the Fourier energy to higher spatial frequencies. (It also lowered contrast from 0.63 to 0.43.) Two densities of noise terrain were used: 512 texture elements/m² and 2,048 texture elements/m², each presented with trees (density was either 16 or 100 trees/km²) and without trees. The result showed that increasing tree density significantly reduced RMS error [$F(2,12) = 30.9, p < .001$], but increasing terrain density did not ($F = 3.3, p = .12$; mean RMS error was 14.6° for the low-density terrain and 14.7° for the high-density terrain).

In the second control experiment, we compared 1-m-wide trees to 1- and 4-m-wide poles, all presented on a gray untextured terrain (density was 16 trees or poles/km²). Removing the textured foliage from the top of the trees to create poles shifted the Fourier energy to lower spatial frequencies, as did increasing pole width. No significant differences among conditions were found ($F_s < 0.02$, $p_s > .9$; mean RMS error ranged from 16.2° to 16.5°). Thus, shifting the Fourier energy in the scenes to higher or lower frequencies did not affect heading control. These results suggest that it is the vertical contours of the trees that improved heading control in Experiment 1.

EXPERIMENT 2

In Experiment 1, we found that the presence of 3-D objects (trees) was more important than terrain texture for the control of heading. Improved performance with 3-D objects may have been due to their vertical extension above the ground plane, which provided enhanced motion parallax information. That is, the locally sharp edges of the vertically extended objects may have provided salient cues to motion parallax. Moreover, in Experiment 1 the crosswind created a lateral disturbance that required the observer to steer to the left or to the right to minimize heading error, and the differential retinal image motion (motion parallax) would have been primarily horizontal. The best carrier of horizontal motion parallax may have been the vertical tree contours, along or across which the motion information was spatially integrated. This hypothesis is similar to that put forward by Longuet-Higgins and Prazdny (1980), who stated that the vanishing point on the retina, which corresponds to the focus of expansion in the optic array, can be located by using the motion parallax at a number of separate retinal positions.

To determine whether or not the visual cues for motion parallax and heading control are enhanced by increasing the extension and/or density of 3-D objects, we investigated the effects of tree height and tree density on heading control for the variable-crosswind task only.

Method

There was a total of 21 test databases, each constructed using the gray untextured terrain that contained no static perspective depth cues for relative depth, no flow field, and no motion parallax. Twenty of the databases were obtained by combining each of four tree densities (1, 4, 16, and 64 trees/km²) with each of five tree heights (1.13, 2.25, 4.5, 9, and 18 m). The trees provided object perspective, increased motion parallax closer to eye level, dynamic occlusion and disclosure (which increased with tree density), and multiple reference objects. One database, which included the target but had no trees or texture, served as a zero-density, zero-height control. This database thus corresponded to the gray untextured/no-tree condition of Experiment 1. Note that a database constructed using trees of zero height is effectively a textured terrain. Such a terrain would be analogous to the two types of textured terrain examined in Experiment 1, and so it was not retested here.

Results

Figure 4 shows the RMS error for the various combinations of tree height and tree density. It can be seen that

increasing either tree density or tree height produced a decrease in RMS error, with tree density having the greater effect.

These data were statistically analyzed in several steps. First, an ANOVA was used to test the effect of tree height by collapsing across the tree density factor with the no-tree control group included. Here, the main effect of tree height was significant [$F(5,30) = 47.39$, $p < .001$], with less error produced by increasing height. Next, the effect of tree density was tested by collapsing across the tree height factor with the no-tree control group included. Here, the main effect of tree density was significant [$F(4,24) = 31.77$, $p < .001$], with less error produced by increasing density.

Next, to test for a tree height $\times$ tree density interaction, an ANOVA of the data from the 20 combinations of tree height and tree density was computed with the no-tree control group excluded. This analysis revealed that the interaction between tree height and tree density was significant [$F(12,72) = 2.5$, $p < .01$]. To probe this interaction, an ANOVA was used to test for simple main effects. This analysis showed that the effect of tree density was significant at each level of tree height, with the F ratios ranging from 8.6 to 17.3 and all $p_s < .01$. Post hoc tests were again performed using Tukey's HSD statistic. For the tree height of 1.13 m, the following differences were reliable (all $p_s < .01$): a density of 1 tree/km² versus 4, 16, and 64 trees/km²; a density of 4 trees/km² versus 16 and 64 trees/km²; and a density of 16 trees/km² versus 64 trees/km². For the tree heights of 2.25 and 4.5 m, the following differences were reliable (all $p_s < .01$): a density of 1 tree/km² versus 4, 16, and 64 trees/km²; and a density of 4 trees/km² versus 16 and 64 trees/km². For the tree heights of 9 and 18 m, all densities were reliably different from one another (all $p_s < .01$). Moreover, the effect of tree height was significant at each level of tree density,

Figure 4. Root mean square (RMS) error in the variable-crosswind task for five tree heights and five tree densities; terrain was untextured. Zero height/zero density represents the absence of trees. Error bars represent ± 1 standard error of the mean for 7 observers.

with the F ratios ranging from 3.23 to 10.0 and all $ps < .05$. Post hoc tests were again performed using Tukey's HSD test. For the densities of 1, 4, and 16 trees/km², the difference between a height of 1.13 and one of 2.25 m was reliable (all $ps < .05$).

The results of Experiment 2 reveal that increasing tree density produced more precise heading control, whereas increasing tree height exerted a much smaller effect. We infer that the primary carrier of horizontal motion parallax information was the vertical tree contours across which the motion information was spatially pooled, which is consistent with the suggestion by Longuet-Higgins and Prazdny (1980).

DISCUSSION

The principal results of the present study are that the vertical extension of 3-D object contours increases heading control when relatively complex imagery is viewed. This increased control was evident even in the absence of terrain texture. Moreover, increasing the density of the vertical contours improved heading control, and this effect did not appear to saturate up to a density of 100 trees/km². These results were obtained from an active heading control task under conditions in which a rotational component in the retinal flow was introduced by simulated crosswinds that acted on an aircraft as it was steered toward a distant target.

The present results extend the literature on heading perception (see, e.g., Li & Warren, 2000; Royden et al., 1994; see the review by Warren, 1998) into the domain of active heading control (see also, e.g., Frey & Owen, 1999; Li & Warren, 2002; Rushton et al., 1999), which is important given that heading perception and heading control may be mediated by anatomically and functionally distinct cortical systems (see, e.g., Aglioti et al., 1995; Bridgeman et al., 2000; Goodale et al., 1991; see the review by Milner & Goodale, 1995). Consistent with Li and Warren (2002), we found that, provided that motion parallax information is available, retinal flow is sufficient for the control of self-motion even when extraretinal information is absent. These results are also consistent with those obtained in the passive heading perception literature.

The improvement in heading control with the addition of vertically extended contours to scenes of untextured or textured terrain was due neither to a change in the global spatial frequency content of the scene nor to the shape or width of the vertically extended contours. Rather, the improvement in heading control was tied to the density and, to a lesser extent, to the length of the vertical contours in the scene. This, in turn, leads us to propose that the density and length (above the ground plane) of vertical contours improved heading control by providing enhanced visual cues to motion parallax. This proposal is consistent with that of Longuet-Higgins and Prazdny (1980), who showed that motion parallax could be used to recover heading direction by a differential velocity process that eliminates the rotational component of the retinal flow and produces

a radial pattern from which heading can be determined. These authors noted that the vanishing point on the retina, which corresponds to the focus of expansion in the optic array, can be located by using the motion parallax at a number of separate retinal positions, an idea that is consistent with the density effect found in the present study.

In recent investigations, Li and Warren (2000, 2002) suggested that vertically extended objects serve primarily as reference objects by which observers update their heading over time. However, in the present study we found that heading control was improved as the density of vertical contours was increased up to a value of 100 trees/km². We think it is unlikely that observers were updating their heading by comparing it to such a large number of objects, either simultaneously or sequentially. It is more likely that the additional improvement in heading control with higher tree densities was due to a spatial pooling of motion parallax information (although some of the trees may have served as reference objects in the way envisaged by Li & Warren, 2000). This conclusion is in contrast to the findings of Li and Warren (2004), who reported that increasing dot density did not improve the ability of observers to judge the path of their egomotion and argued that dense motion parallax is not essential for accurate path perception. However, in their study dot density was varied only from 0.7 to 2 dots/m². These densities may have been too low to affect path perception.

It is surprising that increasing the length (i.e., height) of vertical contours produced only a small effect on heading control, because Andersen and Weymouth (1923) and Blake, Camisa, and Antoinetti (1976) have shown that sensitivity to binocular disparity—the cue for stereoscopic depth perception—increases with the length of vertical test lines. Because binocular disparity results from the horizontal separation of the eyes, it is defined in the horizontal dimension, and thus the studies cited above show that an inherently horizontal cue is best carried by extended vertical contours. Also, an analogous phenomenon has been reported for vernier acuity, which involves the visual discrimination of spatial offset. In this case, judgments of horizontal offset were improved by vertically extended test lines and judgments of vertical offset were improved by horizontally extended test lines (see, e.g., Levi, Klein, & Carney, 2000; Wang & Levi, 1994). From these studies on stereoscopic depth and vernier acuity, one may infer that the visual information was spatially integrated along the length of orthogonally extended contours. However, on the basis of the results of Experiment 2, we conclude that such length summation effects are not significant when motion parallax information is used for heading control.

Finally, our results have implications for differential models of heading based on spatial derivatives of the flow field (see, e.g., Koenderink & van Doorn, 1981; Longuet-Higgins & Prazdny, 1980; Waxman & Ullman, 1985). As was pointed out by Li and Warren (2000; see also Warren, 1998), those models rely on the existence of locally smooth flow fields, which would not exist when trees or

poles are presented on a gray untextured terrain. Thus, the fact that in our study precise active heading control was obtained under such conditions is inconsistent with those models.

In conclusion, the precision of active heading control is greater when 3-D objects, as opposed to textured terrain, are the primary component in flight-simulator imagery. The vertically extended contours of 3-D objects enhance the visual processing of motion parallax information. This enhancement, in turn, facilitates active heading control in the absence of task-related extraretinal information.

REFERENCES

- AGLIOTI, S., DESOUSA, J. F. X., & GOODALE, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, **5**, 679-685.
- ANDERSEN, E. E., & WEYMOUTH, F. W. (1923). Visual perception and the retinal mosaic: I. Retinal mean local sign—an explanation of the fineness of binocular perception of distance. *American Journal of Physiology*, **64**, 561-594.
- BARDY, B. G., WARREN, W. H., JR., & KAY, B. A. (1996). Motion parallax is used to control postural sway during walking. *Experimental Brain Research*, **111**, 271-282.
- BARDY, B. G., WARREN, W. H., JR., & KAY, B. A. (1999). The role of central and peripheral vision in postural control during walking. *Perception & Psychophysics*, **61**, 1356-1368.
- BEST, L. G., WIGHT, D. R., & PEPPLER, P. W. (1999). M2DART: A real image rear-projection display. In D. G. Hopper (Ed.), *Cockpit displays VI: Displays for defense applications* (Proceedings of SPIE, Vol. 3690, pp. 348-355). Bellingham, WA: International Society for Optical Engineering.
- BLAKE, R., CAMISA, J. M., & ANTONETTI, D. N. (1976). Binocular depth discrimination depends on orientation. *Perception & Psychophysics*, **20**, 113-118.
- BRIDGEMAN, B., GEMMER, A., FORSMAN, T., & HUEMER, V. (2000). Processing spatial information in the sensorimotor branch of the visual system. *Vision Research*, **40**, 3539-3552.
- FREY, B. F., & OWEN, D. H. (1999). The utility of motion parallax information for the perception and control of heading. *Journal of Experimental Psychology: Human Perception & Performance*, **25**, 445-460.
- GIBSON, J. J. (1950). *The perception of the visual world*. Boston: Houghton Mifflin.
- GOODALE, M. A., MILNER, A. D., JAKOBSON, L. S., & CAREY, D. P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, **349**, 154-156.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1981). Exterspecific component of the motion parallax field. *Journal of the Optical Society of America*, **71**, 953-957.
- LEVI, D. M., KLEIN, S. A., & CARNEY, T. (2000). Unmasking the mechanisms for Vernier acuity: Evidence for a template model for Vernier acuity. *Vision Research*, **40**, 951-972.
- LI, L., & WARREN, W. H., JR. (2000). Perception of heading during rotation: Sufficiency of dense motion parallax and reference objects. *Vision Research*, **40**, 3873-3894.
- LI, L., & WARREN, W. H., JR. (2002). Retinal flow is sufficient for steering during observer rotation. *Psychological Science*, **13**, 485-491.
- LI, L., & WARREN, W. H., JR. (2004). Path perception during rotation: Influence of instructions, depth range, and dot density. *Vision Research*, **44**, 1879-1889.
- LONGUET-HIGGINS, H. C., & PRAZDNY, K. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society of London: Series B*, **208**, 385-397.
- MILNER, A. D., & GOODALE, M. A. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- RIEGER, J. H., & LAWTON, D. T. (1985). Processing differential image motion. *Journal of the Optical Society of America A*, **2**, 354-360.
- ROSCOE, S. N., EISELE, J. E., & BERGMAN, C. A. (1980). Information and control requirements. In S. N. Roscoe (Ed.), *Aviation psychology* (pp. 33-38). Ames: Iowa State University Press.
- ROYDEN, C. S., CROWELL, J. A., & BANKS, M. S. (1994). Estimating heading during eye movements. *Vision Research*, **34**, 3197-3214.
- RUSHTON, S. K., HARRIS, J. M., & WANN, J. P. (1999). Steering, optic flow, and the respective importance of depth and retinal motion distribution. *Perception*, **28**, 255-266.
- VANDUFFEL, W., FIZE, D., PEUSKENS, H., DENYS, K., SUNAERT, S., TODD, J. T., & ORBAN, G. A. (2002). Extracting 3D from motion: Differences in human and monkey intraparietal cortex. *Science*, **298**, 413-415.
- WANG, H., & LEVI, D. M. (1994). Spatial integration in position acuity. *Vision Research*, **34**, 2859-2877.
- WARREN, W. H., JR. (1998). The state of flow. In T. Watanabe (Ed.), *High-level motion processing: Computational, neurobiological, and psychological perspectives* (pp. 323-341). Cambridge, MA: MIT Press.
- WAXMAN, A. M., & ULLMAN, S. (1985). Surface structure and 3-D motion from image flow: A kinematic analysis. *International Journal of Robotics Research*, **4**, 72-94.

(Manuscript received May 18, 2004;
revision accepted for publication July 13, 2005.)