Heading and path information from retinal flow in naturalistic environments

JAMES E. CUTTING

Cornell University, Ithaca, New York

PETER M. VISHTON

Amherst College, Amherst, Massachusetts

MICHELANGELO FLÜCKIGER and BERNARD BAUMBERGER

Université de Genève, Geneva, Switzerland

and

JOHN D. GERNDT

Cornell University, Ithaca, New York

In four experiments, we explored the heading and path information available to observers as we simulated their locomotion through a cluttered environment while they fixated an object off to the side. Previously, we presented a theory about the information available and used in such situations. For such a theory to be valid, one must be sure of eye position, but we had been unable to monitor gaze systematically; in Experiment 1, we monitored eye position and found performance best when observers fixated the designated object at the center of the display. In Experiment 2, when we masked portions of the display, we found that performance generally matched the amount of display visible when scaled to retinal sensitivity. In Experiments 3 and 4, we then explored the metric of information about heading (nominal vs. absolute) available and found good nominal information but increasingly poor and biased absolute information as observers looked farther from the aimpoint. Part of the cause for this appears to be that some observers perceive that they have traversed a curved path even when taking a linear one. In all cases, we compared our results with those in the literature.

How do we negotiate cluttered environments during our daily activities? How is it that we can generally do this with relative ease and without injury? What information subserves the determination of our direction of movement, often called heading? For over a decade, we have been developing a theory of wayfinding based on the use of particular sources of information in retinal flow, the complex of motion and displacement information projected to the retina of an individual moving through a rigid environment while fixating an object somewhat off his or her path (Cutting, 1986, 1996; Cutting, Springer, Braren, & Johnson, 1992; Cutting, Vishton, & Braren, 1995; Vishton & Cutting, 1995). Strategically, we have simulated naturalistic environments relatively rich in sources of information about layout—occlusion, relative size, relative density, height in the visual field, in addition to motion perspective.¹

The experiments reported here pursue various aspects of the information available in our naturalistic, pursuit-fixation displays in contexts originally presented elsewhere in the literature. In particular, we consider the measurement of simulated pursuit fixations as they present information to the visual system, the distribution of that information in the central retina during these fixations, and the nature of perceived paths taken during such stimulation.

Pursuit Fixation During Gait

As pedestrians, we look at things around us; rarely do we look in the direction of our heading. Cutting et al. (1995) have suggested that we look near our path at stationary obstacles, in part, for the purposes of avoiding them and of updating information about heading direction; we look at moving obstacles only for the purpose of avoidance because information about heading direction seems poor under conditions of pursuit fixation (but see W. H. Warren & Saunders, 1995; Royden & Hildreth, 1996). In this article, we focus on looking at stationary obstacles. In this situation, the retinal flow field of the moving observer combines the rotational flow of a pursuit eye or head movement and the expanding flow of translational motion. In a cinematic analogy to camera motion, the ro-

We thank Paul Braren, Scott H. Johnson, Nan Karwan, and Daniel Simons for discussions of various topics related to this paper, and G. John Andersen and William Warren for instructive reviews. This research was supported by U.S. National Science Foundation Grant SBR-9212786 and by a John Simon Guggenheim Memorial Fellowship during 1993, both to the first author. Requests for information or reprints should be sent to J. E. Cutting, Department of Psychology, Uris Hall, Cornell University, Ithaca, NY 14853-7601 (e-mail: jec7@cornell.edu).

tational flow field is generated by a pan (the rotation of the camera, typically around the vertical axis) and the expanding flow field by a dolly (typically the linear translation of the camera through space).

Our theory of wayfinding is thus based generally on gaze stabilization and on eye movements, particularly on the pursuit fixations executed during locomotion. Saccades are also considered, but only as necessary when a particular pursuit fixation is completed and another must begin. Feedback from eye muscles during pursuit eye movements may also be available and useful to observers under some circumstances (Royden, Banks, & Crowell, 1992; Royden, Crowell, & Banks, 1994). In our previous research, we were unable to monitor eye movements systematically. Thus, in Experiment 1, we recorded eye positions in order to be sure that we knew where observers were looking.

Information About Heading During Wayfinding and Its Retinal Distribution

Our previous research has also suggested that several sources of local information are used to determine one's heading (Cutting, 1996; Cutting et al., 1992). As our research program has developed, these have changed and become more focused. The current list of effective sources includes the displacement direction of the largest (or nearest) object (DDLO) in the visual field and inward displacement (ID).² DDLO accrues from the fact that, when an observer moves through a cluttered environment, objects closer than a fixated object will generally be displaced on the retina in the direction opposite from one's heading. Thus, if DDLO is to the right, heading direction is likely to be to the left. Cutting (1996) has shown that DDLO predicts responses when the direction of gaze is within $.125^\circ$ to 16° of the heading vector, and beyond. That is, observers' responses follow from the presence of DDLO, whether that information correctly predicts heading direction or not. Thus, DDLO is not a flawless source of information, but its correlation with the true state of affairs increases dramatically as gaze-movement angle increases.

ID occurs when objects move toward the fovea during pursuit fixation. It accrues for objects in certain locations beyond, and in certain locations nearer than, the fixation object. The larger the gaze-movement angle, the larger the spatial regions are within which ID will occur. One's nominal heading is in the same direction as ID for objects farther than fixation, and opposite for objects nearer; thus, a rough depth map of the environment seems needed prior to the use of this information (Vishton & Cutting, 1995). With appropriate depth information, ID is a perfect predictor of heading direction. Cutting et al. (1992; Cutting, 1996) have shown that ID is effective when gaze is 4° or more from the heading vector; it is relatively rare, however, when gaze is less than 4° from one's heading.

Cutting et al. (1992, Experiment 2) and Cutting (1996) have shown that in modestly cluttered environments these two sources of information are uncorrelated and both con-

tribute to performance. They also found that any object undergoing outward deceleration, another source of information, also contributed to performance. Moreover, when none of these three sources was present, observers' performance was near, even below, chance. Thus, the use of these local sources of information in retinal flow forms the basis of not only a theory of correct performance in a wayfinding task, but also a theory of errors.

Given an observer fixating midscreen, how does the distribution of this information in the parafovea and beyond affect performance? Recently, there have been several investigations of the locus of information that is important for wayfinding, extending from the fovea into the periphery (e.g., Crowell & Banks, 1993; W. H. Warren & Kurtz, 1992). These previous studies, however, have employed displays that only mimic the radially expanding flow field of a dolly (or translation), with instructions to their viewers to fixate an unmoving point not structurally part of the nearby environment. To us, this procedure seems incompletely representative of eye movement behavior in natural wayfinding tasks; it occurs only when the moving observer is looking at or near the horizon. Thus, in Experiment 2, we investigated the problem in simulated pursuit-fixation displays. In this way, we explored the sensitivity of the retina to the combined motions more normally projected onto it during natural locomotion. More concretely, on the basis of the results of Experiment 1, which will show that pure simulated pursuit fixations are appropriate and adequate to the wayfinding task, we explored in Experiment 2 how information might be distributed across the central retina.

On the Heading Requirements of Moving Observers

How accurate does an individual need to be in estimating the location of his or her heading? Cutting (1986) formalized this question and computed its requirements on the basis of three phases of an avoidance maneuver and the distances covered during each. Working backward in time, they are (1) The distance covered in negotiating a turn, based in part on the coefficient of friction between foot and turf (or wheel and macadam); (2) the distance covered in adjusting one's footfall so that a turn can begin on an appropriate foot; and (3) the distance covered during reaction time to the visual information in the flow field. The angular requirement at a given velocity is roughly the arctangent of width of the body (moved laterally to avoid the object) divided by the total distance covered during the three phases. By far the most important of these is reaction time, and Cutting et al. (1992) estimated that 3 sec of continuous visual stimulation are necessary for observers to attain 95% performance in avoiding a stationary obstacle. Such an estimate, though long, is not out of line with those assessed in real-world situations (Probst, Krafczyk, Brandt, & Wist, 1984; Road Research Laboratory, 1963). Cutting et al. (1992) and Vishton and Cutting (1995) revised this general approach and showed that wayfinding requirements depend on observer velocity. Thus, one needs to know one's heading

within about $\pm 1.3^\circ$ of gaze if running at 6 m/sec, but only within $\pm 3.7^\circ$ if walking at 2 m/sec.

This approach, and estimates derived from it, have been widely cited in the literature (Beer, 1993; Hildreth, 1992; Perrone & Stone, 1994; Sekuler & Blake, 1994, p. 240; van den Berg & Brenner, 1994a; and W. H. Warren, Morris, & Kalish, 1988). But our calculations do not generally apply to many of the situations in which it is cited. In our approach, aimpoint requirements are assessed in a situation of potential danger: specifically, measuring the span within which performance must be highly accurate *when the heading vector is close to the fixated object*. Traveling at 6 m/sec, a runner must know that the heading vector is within 1.25° of foveal gaze, and *to which side*, if he or she is to initiate an avoidance maneuver. Moreover, we claim that having been moving for some previous period of time, the observer only needs temporally discontinuous, nominal updates. Thus, for example, if a runner is looking instantaneously 45° from the heading, he or she will not make a turn to avoid what he or she is looking at; it is well off to the side, and he or she would likely have already taken steps to avoid possible nearer obstacles moments before. Thus, when looking at an object at 45° to one's path, one probably does not need to know where the heading vector is within a region of $\pm 1.25^\circ$; moreover, the data of Crowell and Banks (1993) show that it is not generally available at such eccentricities.

Aimpoints, Heading Directions, and Perceived Paths

The earlier literature on heading judgments asked observers, at the end of a translational flow sequence, to point in their direction of simulated self-movement (Johnston, White, & Cumming, 1973; Llewellyn, 1971; R. Warren, 1976). Results seemed unimpressive, indicating mean errors of 5° – 10° and more. The error in these early results seemed due, at least in part, to the vicissitudes of memory, the pointing response, and the lack of depth simulated in some of the environments. However, given that almost all previous experiments simulated a linear path of the observer (but see Cutting, 1986, Experiments 10 and 11; W. H. Warren, Mestre, Blackwell, & Morris, 1991), most experimenters (including the first two authors here), seem to have assumed that the observers might also perceive such a path.

Subsequent research on perceived heading circumvented memory and pointing-response difficulties through the use of a single probe at the end of a trial (W. H. Warren et al., 1988), a choice among probes (Royden et al., 1992), a paired-comparison among stimuli (Crowell & Banks, 1993), or the direct manipulation of a computer-controlled analog device superimposed on the display (Cutting et al., 1992, Experiment 6; van den Berg & Brenner, 1994a). In each case, the results indicated considerably better accuracy in aimpoint estimation. Each of these different measures suggested that observers have reasonably good *absolute* information about their heading under the assumption of a linear path. That is, from

these methods, an experimenter can directly plot a probability distribution of responses in space around the aimpoint and measure relative accuracy. From such results, one can infer where the observer thinks the aimpoint is located.

In contrast, as suggested above, most of our research has generally used a *nominal* measure of heading. That is, observers have been given a stationary object to look at throughout the trial (typically a tree) and, at the end of the trial, asked to indicate whether their simulated movement was to its left or right (Cutting et al., 1992; Vishton & Cutting, 1995; see also Cutting, 1986). This methodology has been criticized (W. H. Warren et al., 1988) as not allowing us to infer the exact location of the heading.³ That is, from our previous data one cannot directly plot a probability distribution for the perceived aimpoint; instead, one can only plot a response probability function for the location of the simulated fixation with respect to the aimpoint and then perhaps infer the distribution of responses around the aimpoint from those data. Since our stimuli generally involve pursuit fixations, with the aimpoint continually drifting in position away from the fixation point, the latter inference may not be warranted.

However, consistent with the assumptions of our measurement of wayfinding requirements, we believe that *nominal* information (that for which side of gaze the heading vector lies) is all that is needed for the task at any instant, that nominal information may be all that is normally available in the instantaneous flow field and that absolute knowledge (knowing its exact location) is subject to biases. In Experiment 3, then, we altered our typical methodology to allow observers to indicate their precise heading, and in Experiment 4, we explored this information in stimuli simulating motion through both forests and dot clouds (an environment generally devoid of static depth information), comparing our results with others found in the literature.

GENERAL METHOD

Stimuli

Motion sequences were generated on a Personal Iris Workstation (Model 4D/35GT). The Iris is a UNIX-based, noninterlaced raster-scan system with a resolution of $1,280 \times 1,024$ picture elements (pixels). Sequences were patterned after those used by Cutting et al. (1992) and Vishton and Cutting (1995), mimicking the movement of an observer through a tree-filled environment (except in part of Experiment 4) while the observer is looking at the particular tree off his or her path. All measures reported below are scaled to an observer with an eye height of 1.6 m. A wide range of simulated velocities was used (0.5–2.65 m/sec). There were many trees in this environment, each identical in structure. A small forest was created by translating and replicating this tree at many locations across the ground plane. At each location, the tree was rotated to a new random orientation around its vertical axis. The major branching of tree limbs occurred at 1.5 eye heights (or 2.4 m for an individual with an eye height of 1.6 m), and the top of the highest branch was at 2.7 eye heights (4.32 m).

Each trial simulated forward linear movement of the observer with gaze fixed on a stationary object somewhat off to the side. The angle between the line of gaze and the heading vector, called the *gaze-movement angle*, grew steadily as the trial progressed. The partic-

ular initial and final gaze-movement angles employed will be discussed for each experiment. Both are suggested in Figure 1, but for a much larger gaze-movement angle than used here. In Experiments 1–3 and in part of Experiment 4, a red fixation tree appeared at the center of the screen and stayed there throughout the trial, with the remainder of the environment rotating and expanding rigidly around it. Nonfixation trees were gray, the ground plane brown, and the sky cyan. The trees had no leaves, so the stimulus sequence resembled overland travel through a sparse, wintry scene without snow. As the trial progressed, trees could disappear off the edge of the display because of simulated forward motion of the observer, or because of pursuit fixation of the observer on the focal tree, or both. In one condition of Experiment 4, a cloud of white dots on a black background was substituted for the forest and sky, but the experimental situation was otherwise the same.

Procedure

Fifty-six members of the Cornell University community were tested individually in Experiments 1–4. Each was assumed to have normal or corrected-to-normal vision, and each was naive with respect to the experimental hypotheses at the time of testing. Each sat in a moderately lit room, with the edges of the display screen clearly visible. Viewing was binocular, and the participants were encouraged to look at the fixation object and sit 0.5 m from the screen, creating a resolution of 50 pixels per degree of visual angle and an image size of $25^\circ \times 20^\circ$. The perspective calculations used to generate the stimuli were based on this viewing position and distance. In addition, 91 different naive individuals were tested as a group in Experiment 4, participating as part of a class demonstration. The perspective calculations were appropriate for the middle of the auditorium, with an image size of $20^\circ \times 16^\circ$.

Figure 1. A schematic overview of the geometry of a trial, with the simulated path taken by an observer and the lines of gaze at the beginning and at the end of the trial. Note that the final gaze-movement angle here is 20° , twice as large as any used in this set of studies. See Figure 3 for a suggestion of what the layout looked like, although those in Experiments 1, 3, and 4 had neither a central mask nor an aperture.

In all cases, the viewers were told that they would be watching stimuli that simulated their own movements through an environment, and that the stimulus motion would also mimic their fixation on a central element in the field of view. They were encouraged to keep their eyes at midscreen, but eye position was monitored only in Experiment 1. After the end of the motion sequence on each trial, the last frame remained on the screen until the participant made his or her response.

In Experiments 1 and 2, the participants pressed the right key on the Iris mouse if they thought that they were headed to the right of where they were looking during the trial, and the left mouse key if headed to the left; in Experiment 3, they pressed these keys to indicate whether they were headed to the left or right of a probe. In the laboratory portion of Experiment 4, they used a mouse-controlled cursor to estimate their heading, and in the classroom portion, they estimated heading with respect to a poststimulus array of bars. The observers found the task reasonably natural. No feedback was given. A few practice trials without feedback preceded each test sequence. Laboratory viewers were paid at a rate of \$10/h in Experiment 1 (because of the discomfort in wearing the eye-monitoring equipment) and \$5/h in Experiments 2–4; classroom viewers in Experiment 4 were unpaid.

EXPERIMENT 1

Heading Judgments With and Without Monitored Eye Movements

In our previous research (Cutting et al., 1992; Vishton & Cutting, 1995), we used a simulated pursuit fixation technique in our stimulus sequences, emulating the dolly (translation) and pan (rotation about a vertical axis) of a camera, and holding the position of a fixation object in midscreen. In addition, many trials added small vertical and horizontal oscillatory rotations and translations, which we call *bounce* and *sway*. The combination of these motions generates a display that, when one is fixating an object at the middle of the screen, mimics what is seen during natural gait with a pursuit fixation. In none of our previous research, however, did we actually monitor the eye position of our observers. Instead, we simply instructed them to maintain their gaze at midscreen. Since our theory of wayfinding critically depends on gaze stability and on knowing the position of the eye, and since trial sequences lasted as long as 4 sec (Vishton & Cutting, 1995) or more (Cutting et al., 1992, Experiments 2 and 3), it seems unlikely that *all* viewers followed our instructions *all* the time. The purpose of this experiment, then, was to use an eye-movement recording system to be assured of the viewers' fixation, and then to compare those results with those of an unmonitored situation, replicating the results of our previous studies.

Method

Ten observers participated in two conditions—a fixation-monitored condition and a directed-viewing condition. In the monitored condition, viewers wore a headband-mounted eye-movement recording system (Applied Science Laboratories Eye-Trac Model 210). The continuous image of the display screen was recorded with a Pulnix camera mounted on the forehead, and superimposed on it was the continuous eye position as detected by three sensors for each eye and marked by vertical and horizontal crosshairs superimposed on the Pulnix image. Once the equipment was mounted, viewers sat with their heads confined by a chinrest, minimizing head move-

ments. During the course of testing, the experimenter (P.M.V. or J.E.C.) monitored the position of the eyes, ensuring that they were over the fixation tree on a video display. Effective resolution of the Eye-Trac system is about 1.0° measured horizontally and vertically, but under the conditions of this experiment, any deviations from a held position were scored as inaccurate fixations, and the trial was replaced at the end of the sequence. In the directed-viewing condition, observers were simply instructed to maintain gaze at mid-screen, as they had been in our previous studies.

Motion sequences were 4 sec in duration, generated on line at a median of 115 msec/frame. Since the motion of most trees in the displays was quite slow, motion aliasing problems were not bothersome. Moreover, Vishton and Cutting (1995, Experiment 5) demonstrated that wayfinding performance with such stimuli was unimpaired with frame rates as low as 600 msec/frame. Here, trees with fastest retinal motion moved at rates of only about $1^\circ/\text{sec}$, or about 5.8 pixels/frame, and most motion was much slower. The simulated velocity of the observer was 1.6 m/sec, with a required accuracy of 95% at 4.8° , as estimated by Cutting et al. (1992) and Vishton and Cutting (1995). At the beginning of the trial, the fixation tree was at a distance of 32 m, and the visible horizon clipped at 500 m, less than 0.2° below a true horizon for travel on a flat plane. A total of 101 trees were generated in the environment; a mean of 59 ($SD = 5.2$) were visible at the beginning of a trial, and 54 ($SD = 4$) at the end.

Each observer viewed two different randomly ordered sequences of 40 trials: 2 gaze directions (left or right of the heading vector) $\times$ 5 gaze-movement angles (initial angles of 0.5° , 1° , 2° , 4° , and 8° with corresponding final angles of 0.62° , 1.25° , 2.5° , 5° , and 10°) $\times$ 2 carriage conditions (with and without bounce and sway) $\times$ 2 replications of each token but with different random placements of non-fixed trees. Normally in our studies (and as in Experiments 2 and 3 here), we present many more trials, but the onerousness of wearing the Eye-Trac system limited the number to which we wished to subject our viewers. Maximum simulated eye-rotation rate was $0.5^\circ/\text{sec}$, well within the limits suggested by Royden et al. (1992; Royden et al., 1994) for accurate performance. With the additions of calibrations and rest periods, the experimental session lasted about 40 min. The calibration procedure followed the steps outlined in the Eye-Trac manual. All viewers participated first in the eye-monitored condition. The trials during which eye movements were detected were replaced at the end of the sequence, but the mean of these was less than three trials per observer.

Results and Discussion

As in our previous research and in the studies reported later, there were no effects of the side of gaze or of stimulus replications, so we collapsed across these in subsequent analyses. Also as in our previous studies (Cutting et al., 1992; Vishton & Cutting, 1995), there was no effect of carriage: Overall performance was 92% with and 90% without bounce and sway [$F(1,9) < 1$]. Thus, we collapsed the data further across these conditions as well. And finally, as in our previous research, there was a reliable effect of gaze-movement angle [$F(4,36) = 17.6$, $MS_e = 0.45$, $p < .001$], as is shown in Figure 2. We fit logistics functions to the individual data of each of the 10 observers (see also Vishton & Cutting, 1995) and found that all met the 95% performance criterion at a gaze-movement angle of 4.8° in both conditions.

More importantly for considerations here, there was a nearly reliable difference between performances in the eye-monitored (94%) and directed-viewing conditions

Figure 2. The main results of Experiment 1, a nominal direction task, plotted as a function of the final gaze-movement angles. The data from two conditions are shown—that where eye-movement monitoring equipment was worn and used to ensure that the observer's fixation did not drift from the fixation tree at the center of the screen, and that from a directed-viewing condition, in which instructions were the same, but eye movements were not monitored.

(88%) [$F(1,9) = 4.9$, $MS_e = 0.09$, $p > .054$], and a reliable interaction of viewing condition and gaze-movement angle [$F(4,36) = 2.7$, $p < .046$], as is suggested in Figure 2. Eight of the 10 observers performed better in the eye-movement monitored condition. This result pleased us, because it suggests that uninstructed scanning of the display would seem to inhibit, not facilitate, performance at small gaze-movement angles. It also may be that pursuit fixations inconsistent with natural gaze—that is, for example, looking at an object drifting rightward on the display when it would drift leftward in the real world—may occasionally confound responses. Such an account, if valid, would suggest that eye-movement information plays a role in heading judgments (Royden et al., 1992; Royden et al., 1994).

Overview

Observers' heading-direction judgments were sufficiently accurate to meet the wayfinding task demands under the strict experimental conditions of knowing where the eye is positioned during each trial. This result replicates that of W. H. Warren and Hannon (1990) for an absolute judgment task. What is different here is that the task required only a nominal direction judgment, and the simulated environment consisted of a richer array of sources of information about its layout. Moreover, a disadvantage appeared to accrue when the observer was looking elsewhere rather than at the designated fixation tree, at least in these environments. Next, we pursue the distribution of information across the central retina.

EXPERIMENT 2

Heading Information in Combined Translational and Rotational Flow as It Is Distributed Across the Central Retina

Several hypotheses have dominated the discussion of the relation between aspects of wayfinding and the locus of information in the visual field. The first was the *peripheral dominance hypothesis* proposed by Dichgans and Brandt (1978), in which it was said that the peripheral retina dominates the fovea for spatial orientation (see also Berthoz, Pavard, & Young, 1975; Brandt, Dichgans, & Koenig, 1973). Given that Andersen and Braunstein (1985) found strong vection responses (the feeling of self-motion) with a relatively small display, and given that many studies in many laboratories have used relatively small displays and found adequate wayfinding performance (Crowell & Banks, 1993; Cutting, 1986; Cutting et al., 1992; Royden et al., 1992; van den Berg, 1992; van den Berg & Brenner, 1994a; W. H. Warren & Hannon, 1990; W. H. Warren et al., 1988), the peripheral dominance hypothesis no longer seems tenable.

In their exploration of the roles of central and more peripheral vision for wayfinding, W. H. Warren and Kurtz (1992) superimposed peripheral and central masks of various sizes on a flow field of dots simulating the forward translation of the observer. On the basis of their data, they postulated a *functional sensitivity hypothesis*, where wayfinding and orientation information are picked up on the basis of optical information rather than retinal locus, but the central regions are more sensitive (to radial flow) than the peripheral regions (are to more lamellar flow; see also Stoffregen, 1985, 1986). Arguing that Warren and Kurtz confounded retinal position with information type—lamellar motion typically found orthogonal to movement direction and radial motion found parallel to it—Crowell and Banks (1993) explored this issue using by flow fields containing radial flow (dots moving away from a focal point) and lamellar flow (dots moving generally parallel to one another), both of which were presented to a wide variety of retina positions. On the basis of their data, they proposed a *retinal invariance hypothesis*, where the perception of heading is largely independent of retinal position and can be predicted on the basis of motion detection efficiency at all retinal locations (see also Stoffregen & Riccio, 1990).

Here, we pursue support for a simpler notion, which we call the *retinal sensitivity hypothesis*. That is, we propose that observers' wayfinding responses reflect the degree to which they are sensitive to motion at different parts of the retina, as that motion combines translational and rotational flow. Thus, no flow decomposition or functional specialization is entailed. The problem with *functional sensitivity* and *retinal invariance*, as we see it, is not with either of the hypotheses as stated or researched, but with their generality to natural conditions of pedestrian wayfinding. Both W. H. Warren and Kurtz (1992) and Crowell and Banks (1993) used displays that mimicked

the linear translation of the observer who held his or her fixation at a constant angle with respect to the heading vector. Thus, the pattern of retinal stimulation was only that of the translational flow field, be it the radially expanding or the lamellar portions. Since human beings are mobile-eyed creatures, and since most of our time during pedestrian travel is taken up with a series of pursuit fixations, which combine both translational and rotational flow, we believe that functional sensitivity and retinal invariance are not concepts pertinent to the bulk of eye-movement behavior during human gait. In particular, under normal conditions of pedestrian viewing, neither pure radial nor pure lamellar flow is typically presented to either the fovea or the near periphery. With pursuit-fixation displays, one can assess the sensitivity of various regions of the retina to the complex of motions projected to it during normal pedestrian wayfinding.

Method

Ten observers participated in the two conditions. Half viewed, first, a sequence of stimuli masked in the periphery with the region around a central aperture, followed by sequences with a central mask and unoccluded periphery. The other half participated in the reverse order. Trial durations were 3.67 sec, and sequences were generated at a median of 183 msec/frame. Simulated observer velocity was 1.28 m/sec (a saunter), requiring a 95% accuracy at a gaze-movement angle within 5.7°. The horizon was at a distance of 100 msec, and the fixation tree at 45 m. Initial gaze-movement angles were 0.45°, 0.9°, 1.8°, 3.6°, and 7.2°; respective final gaze-movement angles were 0.5°, 1°, 2°, 4°, and 8°. The most rapid simulated eye-rotation rates were 0.2°/sec, again well within the limit suggested by Royden et al. (1992; Royden et al., 1994) for accurate simulated pursuit fixations. The experiment took about 90 min.

Apertures. In one set of sequences, environments were seen through circular apertures of various radii. The intersection of the horizon and red fixation tree (at an initial distance of 28 eye heights) was at the center of each aperture. The display screen was digitally masked in black beyond a fixed radius on each trial. Radii were 25, 50, 100, 200, 400, and 819 pixels, the latter showing the full screen (but, of course, not leaving a circular image). From a viewing distance of 0.5 m, the apertures were 1°, 2°, 4°, 8°, and 16° in diameter, with the full screen condition 25° × 20°. As a percentage of full screen, the viewing areas were 0.15%, 0.6%, 2.4%, 9.6%, 38.4%, and 100% across the six conditions, and although 100 trees were generated for the environment (as in Experiment 1), the number visible in each condition covaried with the aperture size. The upper panel of Figure 3 shows an example of an 8° aperture. A different random sequence of 240 trials was presented to each observer: 5 gaze-movement angles × 6 apertures × 2 gaze directions × 2 carriage conditions (with and without bounce and sway) × 2 replications.

Central masks. Sequences in this condition were the inverse of Condition 1. Rather than blocking out various amounts of the periphery of each trial, a circular region centered at the middle of the screen was digitally masked in black. To provide observers with a steady object to look at, we placed a white fixation cross (1° × 1°) in the middle of the black shield where the red fixation tree would otherwise have been. Observers were encouraged to fixate the cross throughout the trial. These masks had the same sizes as the apertures—1°, 2°, 4°, 8°, and 16°—with the addition of a no-mask stimulus (a 0° mask). The percentages of the screen area left uncovered were 99.85%, 99.4%, 97.6%, 90.4%, 61.6%, and 100%, respectively, across the six conditions. Again, the number of visible trees covaried with mask size. The lower panel of Figure 3 shows an 8°

Figure 3. A single frame for a sample 8° aperture and a sample 8° mask in sequences of Experiment 2. The display subtended 25° × 20°.

mask. Initial and final gaze-movement angles were as in the aperture condition. Again, simulated observer velocity was 1.22 m/sec, trial durations were 3.67 sec, and a different random sequence of 240 trials with the same general characteristics as in the aperture condition was presented to each observer.

Results and Preliminary Discussion

Apertures. As expected, we found reliable effects of aperture size [$F(5,45) = 14.7$, $MS_e = 15.7$, $p < .0001$] and gaze-movement angle [$F(4,36) = 20.7$, $MS_e = 22.5$, $p < .0001$], and their interaction was significant [$F(20,180) = 2.01$, $MS_e = 1.59$, $p < .005$]. These patterns are shown in the top left panel of Figure 4, collapsed across the two largest (16° aperture and full screen), the two intermediate (4° and 8° apertures), and the two smallest conditions (1° and 2° apertures). As in Experiment 1 here and in the experiments of Cutting et al. (1992) and Vishton and Cutting (1995), there was no effect of carriage ($F < 1.0$).

Again, the individual data in each condition were fit to logistics functions. Table 1 shows that only for the two largest apertures (16° and full screen) did all or nearly all observers meet a 75% performance criterion (used by W. H. Warren et al., 1988), and half or nearly half met a 95%

criterion (used by Cutting et al., 1992; Vishton & Cutting, 1995) at the 5.7° gaze-movement angle. The medians of these individual logistics functions at gaze-movement angles of 0.5°, 1°, 2°, 4°, and 8° for each aperture condition were then determined, and a new group logistics function was fit to them (see also Vishton & Cutting, 1995). The fan of these functions is shown in the top right panel of Figure 4. Such results suggest two things: First, the relative difficulty across all conditions may have depressed performance even in the easiest conditions, and second, a modestly large portion of the visual field (as much as 16° around the fovea) is necessary for observers to perform a wayfinding task under the conditions that we have habitually investigated.

Central masks. We also found reliable effects of central-mask size [$F(5,45) = 9.3$, $MS_e = 8.2$, $p < .0001$] and final gaze-movement angle [$F(5,45) = 65.1$, $MS_e = 44.4$, $p < .0001$], and their interaction was significant [$F(20,180) = 2.04$, $MS_e = 1.16$, $p < .003$]; this can be seen in the bottom left panel of Figure 4. Again, the largest two, the middle two, and the smallest two mask-size conditions were collapsed together. As before, there was no effect of carriage ($F < 1.0$). The individual data were again fit to logistics functions. Table 1 shows that a 75% criterion was met by nearly all observers with masks less than 16°, but that attainment of a 95% criterion was generally possible only with masks smaller than 4°. Again, group median logistics functions were determined from the individual fits for each mask condition, and these are shown in the lower right panel of Figure 4.

Rescaling the Data to Test the Retinal Sensitivity Hypothesis

The data from the 4° and 8° gaze-movement-angle trials were then selected from the six aperture and six central-field masking conditions of Figure 4 and replotted according to the percentage of the visual display that was visible. These data are shown in the left panel of Figure 5. Notice the discrepancy between the aperture data and the central-field mask data as a raw function of the display area. However, if we assume that observers were fixating the center of the display, which Experiment 1 suggested is optimal, we can assume that a general motion detection function is applicable, as shown in the right panel of Figure 5. These

Table 1
The Number of Observers (Out of 10) in Experiment 2 Meeting the Performance Criteria of 75% and 95% at 5.7° Initial Gaze-Movement Angle and at a Simulated Velocity of 1.2 m/sec

Size of Aperture or Mask	Looking Through Apertures		Looking at Central Masks	
	75%	95%	75%	95%
None	—	—	10	9
1°	2	0	10	7
2°	2	1	9	6
4°	3	2	9	2
8°	7	3	9	2
16°	10	5	6	1
Full screen	9	4	—	—

Figure 4. Results of Experiment 2. The upper left panel shows the data for the various apertures; the upper right panel, the group median logistics functions that correspond to them; the lower left panel, those for the various masks; and the lower right panel, the corresponding group median logistics functions. The rectangular outlines of the icons represent the display screen, and the black portion of each is proportional to the area of the screen occluded by the aperture surround or by the central mask. These icons are used again in Figure 5.

data have been taken from Leibowitz, Johnson, and Isabelle (1972) and compared with the data of Johnson and Leibowitz (1979) for static resolution. Both functions show acuity normalized to foveal performance, arbitrarily truncated at an eccentricity of 25° into the periphery.

From these data, one can then rescale the results of the two conditions according to retinal sensitivity, under the assumption that generalizations from threshold to suprathreshold situations are valid. First, since the viewing screen was 25° wide, the relevant area under the motion detection function between 0° and 12.5° . One can use the area under this curve as a reference and normalize it to 1.0. Second, one can then consider viewer performance

at all apertures and masks as a function of the proportion of this area. Thus, apertures always include the left-hand portion of the function and its area; central-field masks always include the right portion of the function and the area underneath. Third, these proportions are then squared to convert from lineal to areal units and are plotted as in the middle panel of Figure 5. The overlap of these results suggests that a simple, single account—a retinal sensitivity hypothesis—can account for the data. It suggests further that there is no need to consider functional sensitivity (W. H. Warren & Kurtz, 1992) or retinal invariance (Crowell & Banks, 1993) accounts of the data under conditions of normal gait and eye-movement behavior.

Figure 5. The left panel shows the performance data for gaze-movement angles of 4° and 8° selected from the six aperture and six central masking conditions in Experiment 2. The right panel shows the data taken from Leibowitz, Johnson, and Isabelle (1972) and from Johnson and Leibowitz (1979) for motion and static resolution detection at various eccentricities, scaled to performance in the fovea. The central panel rescales the data from the left panel as a function of retinal sensitivity, estimated from the motion function in the right panel. The icons correspond to those for particular conditions in Figure 4.

Overview

We have argued that, under normal pedestrian conditions, neither the central visual field nor the periphery is systematically presented with pure radial flow or lamellar flow. Instead, a pedestrian typically executes a pursuit fixation, following an object off his or her path in the middle distance. In such cases, the motions generated by eye rotations are superimposed on those generated by forward movement and create a hybrid motion field, often characterized by opposing motions in the foveal and parafoveal regions. Under such a situation, it appears that, across the six aperture and six central masking conditions, the results are best explained by simple, differential retinal sensitivity to motion.

From this result, it might seem as if we are espousing a neural mechanism that pools information across relatively large regions of the visual field. We are not. Elsewhere, we have documented that local information (the displacement of particular objects in the visual field) rather than global information (various forms of spatial pooling) is a better predictor of wayfinding judgments (Cutting, 1996; Cutting, Flückiger, Baumberger, & Gerndt, 1996). Thus, in this context, we believe that scaled retinal sensitivity reflects the probability of registering the displacements of an informative object within the unmasked field of view. The general lability of the raw data, as suggested in the right panels of Figure 4, support this idea, but as yet we have no data with which to test it.

We next shift gears. Whereas in Experiments 1 and 2, we explored the measurement of fixations and distribution of information during them, in Experiments 3 and 4 we explored the nature of perceived paths taken during these simulated pursuit fixation trials.

EXPERIMENT 3

Nominal or Absolute Information About Heading in Pursuit-Fixation Displays?

How should we best characterize heading perception on the basis of visual information? Shall we say that moving observers know their absolute heading within some degree of accuracy, or only that they know the nominal direction of their heading with respect to where they are looking? To be concrete, in our experimental situation we know the following: At any instant when observers are looking 4° to the right of their aimpoint and going at a velocity of near 2 m/sec, they are about 95% correct in saying that the aimpoint is to their left. Nonetheless, we do not know where exactly they think their aimpoint is located, nor do we know whether or not they perceive themselves on a straight path. If they perceive themselves on a straight path, they may think that the heading vector is systematically located at 4° to the left, but equally they may think that it is only 2° or even 8° to the left. If they perceive themselves on a curved path, that path might curve away from the fixated tree, or even toward it and then behind it. Thus, this experiment was a preliminary exploration of the perception of absolute headings and paths taken; Experiment 4 followed up on it.

Method

Trial sequences simulated the linear translation of the observer across the tree-filled plane at 2.85 m/sec (a jogging pace) for 3.5 sec. Required accuracy at this velocity would be $\pm 2.6^\circ$ according to Cutting et al. (1992) and Vishton and Cutting (1995). Unlike in Experiments 1 and 2, no oscillatory rotational or translational additions of bounce and sway of the observer were simulated. Due to graphics optimizations, sequences were generated at a median of 65 msec/

frame, considerably faster than in Experiments 1 and 2. The ground plane was covered with a mean of 29.5 ($SD = 1.6$) trees visible at the beginning of the trial and 25 ($SD = 2.1$) at the end. Again, fixation was to be maintained on the red tree at the middle of the screen at an initial distance of 40 m, with the visible horizon set at 110 m, or about 0.8° below a true horizon. Eye position was not monitored.

The initial gaze-movement angles were 0.5° , 1° , 2° , or 4° , and the respective final angles were 0.62° , 1.25° , 2.5° , or 5° . Either throughout or at the end of each trial, a red probe appeared at the visible horizon 0.5° , 1° , 2° , or 4° to the left or right of the aimpoint. At the end of the trial, the observers used the left and right mouse keys to indicate whether the probe was to the left or right of the true heading. Mean simulated eye/head rotation rate for the 4° initial gaze-movement trials was less than $0.3^\circ/\text{sec}$, again well within the limits suggested by Royden et al. (1992; Royden et al., 1994) for accurate aimpoint estimation with such stimulus sequences.

Sixteen observers participated. Each viewed two sequences, one with the probe continuously present during the course of the trial and one with the probe appearing on the screen after all motion had terminated, as in Warren et al. (1988, Experiment 1). Thus, each participant looked at two different randomly ordered sequences of 128 trials: 2 gaze directions (left and right of the aimpoint) $\times$ 4 gaze-movement angles $\times$ 2 probe directions (left and right of the aimpoint) $\times$ 4 probe-movement angles $\times$ 2 replications of each sequence type with different randomly placed nonfixation trees. Half of the subjects first viewed the sequence with a probe continuously present on each trial and then the sequence with a probe appearing at the end of each trial; half viewed the sequences in reverse order. The experiment lasted about 45 min.

Results and Preliminary Discussion

As in the study of Warren et al. (1988, Experiment 1), there was no effect of probe presentation, whether of those continuously present during trial sequences (73% correct performance) or those presented only at the end (71%; $F < 1$). Nor were there any interactions involving probe condition. Thus, in our additional analyses we collapse

across probe types. There was a reliable effect of probe-movement angle [$F(3,45) = 55.5$, $MS_e = 3.98$, $p < .0001$], shown in the left panel of Figure 6. These results are compatible with those of Warren and Hannon (1990, Experiment 2), which are also shown, who did not systematically vary gaze-movement angles but always kept them within a range somewhat smaller than that used here.

There was also a reliable effect of gaze-movement angle [$F(3,45) = 15.6$, $p < .0001$], shown in the middle panel of Figure 6, with performance *decreasing* with increasing angle. Note that in the context of judgments around probes this effect is in the reverse direction from judgments around a fixation object. In particular, these data indicate that performance in estimating the aimpoint location deteriorates the farther one looks away from one's heading, as Crowell and Banks (1993) found for a much larger range of eccentricities. However, nominal judgments about the direction of one's heading increase in accuracy the farther one looks away from the heading vector. The import of this result for us is a suggestion that perhaps one should not characterize heading information as absolute and decreasing with gaze-movement angle, but rather as nominal and increasing with gaze-movement angle. We will discuss this idea in more detail later.

As a partial replication of our previous work with a nominal-direction judgment task we selected those trials in which the probe was nearest the fixation tree. Such a situation occurred on four types of trials when probe and gaze were closest—when pairs of end gaze and probe positions were $.62^\circ$ and $.5^\circ$, 1.25° and 1° , 2.5° and 2° , and 5° and 4° , respectively, to the same side of the movement vector. That is, for example, when gaze ended 0.62° to the left of heading and the probe was 0.5° to the left of heading, a judgment that the heading was to the left of

Figure 6. The main results of Experiment 3, a probe task. The left panel compares our results with those of W. H. Warren and Hannon (1990). The middle panel shows the decline in performance with increases in gaze-movement angle, which replicates most closely the nominal direction task of Experiment 1 and our previous work (Cutting, Springer, Braren, & Johnson, 1992; Vishton & Cutting, 1995).

the probe would be essentially the same as a judgment that it was to the left of the fixation tree, equivalent to our nominal task. These data are also shown in the right panel of Figure 6, and are similar to those of Figure 2 for Experiment 1.

Beyond this, however, there were several reliable interactions that we did not initially understand. After reinspection of the data, it became clear that we should recast the analysis of variance with two old and one new factor: gaze angle $\times$ probe angle $\times$ side, where the last factor concerned whether the heading was to the same side as the gaze with respect to the probe (*same*) or not (*different*). Results differed if, for example, both the simulated gaze and heading were to the left of the probe rather than gaze to the left and heading right, or vice versa. In particular, when gaze and heading were on the same side, performance was considerably higher (82%) than when they were on different sides [62%; $F(1,15) = 8.5$, $MS_e = 10.5$, $p < .015$]. Recast in this manner, there was also a gaze-movement angle $\times$ side interaction [$F(3,45) = 13.2$, $MS_e = 1.57$, $p < .0001$]. That is, as the gaze-movement angle increased, there was an increase in the discrepancy between same-side and different-side performance. These data indicate that the probability distribution of the observers' responses is not centered on the true heading vector; instead, there is an underestimation of its location as displaced from fixation.

To demonstrate the latter effect in greater detail, we converted the individual data in each condition to a psychophysical function, used probit analysis to fit it, and

then determined the location of the perceived aimpoint at each gaze-movement angle for each individual. The results are shown in the left panel of Figure 7. In these derived data, there was a main effect of gaze-movement angle [$F(3,45) = 7.86$, $MS_e = 10.95$, $p < .0001$], indicating that as the true gaze-movement angle increased, so did its perceived counterpart. Nonetheless, also suggested in the figure, the effect of underestimation of perceived heading was even more robust [$F(3,45) = 14.07$, $MS_e = 19.6$, $p < .0001$]; perceived heading was a bit less than half its true value. The fact that the perceived aimpoint hovers near gaze position might also drive other biases in judgments of aimpoint (see, e.g., Cutting et al., 1992, Experiments 4–6; Johnston et al., 1973; Llewellyn, 1971; W. H. Warren et al., 1991).

Overview

The left panel of Figure 6 shows that we replicated the probe findings of W. H. Warren and Hannon (1990), as well as our own (Cutting et al., 1992, Experiments 4 and 5), in that observers were as accurate in the laboratory as they needed to be to perform the task in the real world. Nevertheless, also in our data and perhaps in those of others, there is an additional trend denoting underestimation of the perceived aimpoint, as shown in the left panel of Figure 7. As shown in the right panel of Figure 6, nominal information for the direction of heading is excellent and increases with the magnitude of the gaze-movement angle, but as shown in the central panel of Figure 6, absolute information becomes increasingly in-

Figure 7. The left panel shows the heading estimates derived from the results of Experiment 3, compared with those of Royden, Banks, and Crowell (1992, Experiment 3; Royden, Crowell, & Banks, 1994, Experiment 7). The Royden et al. data correspond to the estimated means of Subjects M.S.B. and T.R.C. in the two conditions when the true headings were $\pm 4^\circ$ from midscreen and fixation. The right panel shows the results of Experiment 4. In general, nominal heading judgments are highly accurate for forest stimuli and inaccurate for dot clouds; absolute heading information is relatively poor for both.

distinct. Moreover, and more importantly, its mean error also increases, biased in the direction of fixation. Because of this bias, absolute heading information becomes more infirm with increase in gaze-movement angle. Since nominal heading judgments are unaffected by such bias, we suggest that this is another reason to consider that nominal heading direction judgments, rather than absolute heading judgments, reflect necessities of everyday pedestrian locomotion.

A Comparison With the Data of Royden et al. (1992)

Our results also appear to contrast with those of Royden et al. (1992, Experiment 3; see also Royden et al., 1994, Experiment 7), which are also shown in the left panel of Figure 7. Here we have replotted their results as a function of final gaze-movement angle, our independent variable; Royden et al., of course, plotted them as a function of simulated rotation rate. Although the parameters of rotation rates and fixation distances were different in the two studies, our observers generally estimated their heading to be between the line of gaze and the true heading, whereas those in Royden et al. often perceived their heading to be to the side of the gaze direction opposite the true heading. The reason for their result seems fairly straightforward; their observers "thought they were moving along a curvilinear path in the direction of the simulated eye movement" (Royden et al., 1992, p. 585). Observers in our study, in contrast, did not mention this fact nor were their data generally consistent with it; only 2 of 16 observers showed perceived aimpoints consistently in a direction opposite the true heading vector. The geometries of linear and curvilinear paths are suggested schematically in Figure 8, and if these were perceived, their difference explains the difference in results.

More particularly, if moving observers perceive that they are on a straight path, the gaze-movement angle increases during the course of the trial, and the visual display simulates a combined dolly and pan. If, as in Figure 8, that object is to the right, the pan is also to the right. If, on the other hand, observers perceive that they are on a curvilinear path, the most likely impression is that the fixation object is on or near that curved path (Motorcycle Safety Foundation, 1992).⁴ In this situation, the gaze-movement angle decreases during the course of the trial, and the visual display would be perceived to emulate a curvilinear dolly and a pan to the left. Such a pan is in the reverse direction of the linear case. Royden (1994) has shown that there is optically very little difference between these two cases, at least for trial durations of 1.25 sec.

At least three methodological differences might contribute to the differences in viewers' perceptions, and to the difference between the results shown in the left panel of Figure 7. First, simulated eye rotation rates were much slower here (less than $0.3^\circ/\text{sec}$) than in the studies of Royden et al. (1992; Royden et al., 1994; 0° – $5^\circ/\text{sec}$). More-

Figure 8. The geometry of pursuit fixations in situations of possible perceived straight and curvilinear paths. The fixated object is placed here on the circular path, in part because the Motorcycle Safety Foundation (1992) suggests that drivers fixating objects off their instantaneous linear path are likely to swerve and collide with that object. This difference might be taken to account for the difference between the results shown in Figure 7 for Experiment 3 and the studies of Royden, Banks, and Crowell (1992). Some of the results of Experiment 4 suggest that a hybrid of the two paths may be perceived in the case of the forest stimuli.

over, Royden et al. showed a striking difference for stimuli with simulated rotations above and below $1^\circ/\text{sec}$.

Second, the trial durations of Royden et al. (1992) were quite brief—1.25 sec. In contrast, the sequence durations in our study were 3.5 sec. In our work, we have generally used sequences longer than 3 sec in duration because it is only with such stimuli that most observers can achieve the 95% correct nominal-direction performance at a given criterion gaze-movement angle (Cutting et al., 1992; Vishton & Cutting, 1995). Perhaps the brevity of their trials encouraged the perception of curvilinearity; and perhaps the length of ours discouraged it.

Third, and perhaps most important, the stimulus sequences of Royden et al. (1992; Royden et al., 1994) simulated a situation in which the observer was moving through a dot cloud—that is, no information, other than motion and density (farther dots are clustered closer together than nearer dots), about the depth and layout of the environment was available to the observer. In contrast, our stimuli mimicked the movement through a sparse forest, where motion and density information are joined by occlusion, relative size, and height in the visual field (Cutting & Vishton, 1995). Such additional information is known to improve performance on heading estimation tasks (Cutting et al., 1992, Experiment 2; van den Berg & Brenner, 1994a, 1994b; Vishton, Nijhawan, & Cutting, 1994).

The bulk of Experiment 4, then, was designed to test the effects of stimulus-element type (dot clouds vs. forests) while the simulated eye rotation rates and durations were held constant. Experiment 4 also included a small foray into the effects of duration. In general, our hope was to replicate both the results of Royden et al. (1992, Experiment 3) and those of our Experiment 3, in an effort to understand the causes of the differences between them.

EXPERIMENT 4

Heading Information While Moving Through Dot Clouds and Forests

Method

Data were gathered in two situations, one in a laboratory setting (like that of Experiments 1–3) with 20 individuals run individually, and the other in a large classroom, before a group of 91 students as a short demonstration. For both, stimulus sequences simulated movement through a dot cloud (as in Royden et al., 1992; Royden et al., 1994) or through a forest (as in Experiment 3 and in Cutting et al., 1992, and Vishton & Cutting, 1995). No attempt was made to equate the dot-cloud and forest stimuli in any way; the numbers of dots and trees were chosen simply to match their general use in the literature (see Vishton & Cutting, 1995, Experiments 4 and 5, for further discussion). All sequences were generated at a median of 66 msec/frame. We will consider the sequences in the laboratory setting first.

Laboratory setting. A first set of sequences was patterned after a particular condition in Royden et al. (1994, Experiment 7; but cf. Royden et al., 1992, Experiment 3). Each trial mimicked forward motion (2.5 m/sec) while one is looking 5° off to the side at a white fixation cross, which always remained at midscreen. A mean of 220 dots ($SD = 17$) were present at the beginning of each trial. Dots were laid out randomly at different x - and y -coordinates, and linearly spaced along the z -axis at depths between 0 and 37 m. Each dot was white and subtended 15' of arc against a black background. The initial distance of the fixation cross was set at 5.6, 8.1, and 15.6 m, and at 10 km. All trials were of 1.25-sec duration, yielding mean simulated eye rotation rates 5°, 2°, 1°, and 0°/sec, respectively. Initial gaze-movement angles were always $\pm 5^\circ$ from the center of the screen; final angles were $\pm 11.25^\circ$, 7.5° , 6.25° , and 5° , respectively. Different, randomly ordered, sequences of 48 trials were presented to each subject: 2 fixation directions (left and right of the heading vector) $\times$ 4 simulated eye-rotation rates $\times$ 6 replications of each trial type with different random positions of dots.

A second sequence of 48 stimuli with the same general parameters was generated to simulate travel through sparse forests, as in Experiments 1–3, with a mean of 25 trees ($SD = 4.5$) visible at the beginning of the trial. Given that 25 trees spread between 0 and 37.3 m would have created very dense, generally nonnegotiable environments, we scaled up the simulated observer velocity and distances; that is, simulated movement was at 10 m/sec (beyond humanly sustainable footspeed) through a forest with trees spread out between 0 and 149.2 m. In this rescaled situation, the near-fixation trees were at similarly increased distances—22.4, 32.4, and 62.4 m. The farthest fixation tree, which was at 250 m, lay well beyond the far edge of the rest of the forest (and was occasionally occluded during much of the trial). Since dot clouds offer no significant eye height or distance information, such rescaling does not change the comparability of the two conditions. To increase contrast among the various trees in the forest in comparison with the displays in Experiments 1–3, each nonfixation tree was colored a different shade of gray.

Viewers were run individually. They all watched the dot-cloud sequences first, then the forest sequences. Such a procedure was followed because we generally knew what to expect from the forest stimuli from previous studies, but were unsure of the dot-cloud stimuli. Since Royden et al.'s (1992; Royden et al., 1994) viewers had never seen forest stimuli, we thought it best to begin with these here. In both cases, at the end of each trial, the final frame of the sequences remained on the screen, and the viewer moved the mouse to position a cursor bar (which straddled the horizon) to the left or right of the fixation cross or tree, indicating the position of his or her aimpoint at the end of the trial. Because the cursor moved with the mouse and always occluded trees and dots in its region, it did not appear to be part of the environment, and in particular it did not ap-

pear to be located at the horizon in the forest stimuli. The experiment lasted about 20 min.

Classroom setting. Here, a first sequence consisted of dot-cloud stimuli patterned after those in Royden et al. (1992, Experiment 3). Three types of trials were presented, each mimicking forward motion at 0.5 m/sec while one was looking 5° off to the side at a fixation cross that was structurally part of the field of about 220 visible dots scattered between 0 and 37 m. In one trial type, the fixation cross was presented at an initial distance of 1.62 m and motion lasted for 1.25 sec, generating a simulated eye rotation of 2.5°/sec; in a second trial type, the fixation distance was increased to 16.2 m, but trial duration remained at 1.25 sec, generating a simulated eye rotation of .25°/sec; and in a third, the fixation distance was 16.2 m, but trial duration was increased to 4 sec, generating a mean eye rotation of .3°/sec. Again, initial gaze-movement angles were always 5°; final gaze-movement angles were about 8.1°, 5.3°, and 5.6°, respectively. There were 18 randomly ordered stimuli in all: 3 trial types $\times$ 2 gaze directions $\times$ 3 replications of each with differently positioned random dots. These were presented as a 4-min video sequence and back-projected in a large auditorium.

Three types of forest stimuli were generated in a similar manner with identical values of simulated observer velocity, gaze-movement angles, simulated eye-rotation rates, and durations. Thus, here and unlike in the laboratory sequences, there was no rescaling of depth and velocity. Each sequence had a red fixation tree and a mean of seven other trees laid out between 0 and 37 m. Viewers watched 18 trials of the same types and with the same durations as those for the dot-cloud stimuli.

Again, for both stimulus types, the last frame of the trial sequence remained on the screen. Here we superimposed 19 vertical bars, straddling the horizon and occluding all behind them. These bars were numbered left to right from -9 to 0 to 9 . The middle bar was at midscreen, and the others were spaced at 1° increments to the left and right. After each trial, the viewers wrote the number of the bar closest to their perceived aimpoint on an answer sheet, using the numbers -10 to 0 to 10 , with the greatest numbers indicating aimpoints off the screen, to the left and right, respectively. The viewers also indicated their location in the auditorium, so that analyses could be undertaken concerning their angle and distance with respect to the screen. Each test sequence was preceded by three practice trials, one of each type.

Results and Preliminary Discussion

Laboratory setting. The results generally replicated both those of Experiment 3 and of Royden et al. (1992; Royden et al., 1994). The major effects are shown in the right panel of Figure 7 as the functions with larger circles or squares. Notice that, as was suggested in the left panel, there was a striking difference between the dot-cloud and forest stimuli [$F(1,19) = 15.4$, $MS_e = 29,886$, $p < .001$]. The mean perceived location of the heading for the forest stimuli remained between the line of gaze and the heading vector; the headings for the dot cloud stimuli, however, deviated more markedly and, except in the case of the 0°/sec stimulus, all mean headings were perceived to be to the side opposite from the heading vector. Only 5 of the 20 participants, when viewing the forest stimuli, consistently indicated their heading to be opposite the true heading vector, whereas 19 of 20 did so for the dot-cloud stimuli. Individual functions are shown in Figure 9, plotted in the same manner as in Royden et al. (1992; Royden et al., 1994).

There was also a reliable main effect of simulated eye rotation rate [$F(3,57) = 30.9$, $MS_e = 18,697$, $p < .0001$],

as suggested by Royden et al. (1992; Royden et al., 1994), and most importantly there was an interaction between stimulus type and rotation rate [$F(3,57) = 14.6$, $MS_e = 4,336$, $p < .0001$]. Thus the patterns of results for the two stimulus conditions—dot cloud versus forest sequences—were substantially different.

Classroom setting. Let us consider, first, only the trials with short durations (1.25 sec), summarized as the smaller circles and squares in the left panel of Figure 7. There was a prominent effect of simulated environment—dot cloud versus forest stimuli [$F(1,90) = 99.7$, $MS_e = 1,320$, $p < .001$]. For the trials simulating paths through a forest, aimpoint estimates remained close to the line of gaze, as before, whereas those for the dot clouds were farther from the heading vector and to the opposite side. Only 26 of the 91 observers consistently placed their heading to the wrong side of fixation for the forest stimuli, whereas 85 of 91 did so for the dot cloud stimuli [$\chi^2(1) = 31.3$, $p < .0001$]. Again, there was a pronounced effect of simulated rotation rate [$F(1,90) = 80.9$, $MS_e = 801$, $p < .001$] and an interaction between stimulus type and rotation rate [$F(1,90) = 91.8$, $MS_e = 881$, $p < .001$].

In addition, among the stimuli differing only in duration (not shown in Figure 7), there was also a reliable effect of simulated environment [$F(1,90) = 66.9$, $MS_e = 145.3$, $p < .001$]. That is, there was a mean error of $+0.9^\circ$ from fixation toward the heading vector for forest stimuli and one of -0.1° away from the heading vector for dot-cloud stimuli. For the 4-sec forest sequences only 7 of 91 observers consistently misplaced their nominal heading direction, whereas 58 of 91 did so for the dot-cloud sequences [$\chi^2(1) = 44.6$, $p < .0001$]. Moreover, the decrease in nominal errors for the 4-sec forest stimuli over the 1.25-sec forest stimuli was reliable [$\chi^2(1) = 10.9$, $p < .001$], as it was for the dot-cloud stimuli [$\chi^2(1) =$

5.1, $p < .05$]. Finally, there was no effect on the result of observer position in the auditorium. This last finding replicates various results of Gibson (1947) and Goldstein (1987), and suggests that no recalibration of the gaze-movement angle or simulated-rotation results is necessary as a function of each observer's position.

On the Nature of One's Perceived Path

To try to understand the nature of the perceived paths, we interviewed the 2 subjects whose forest data were most aberrant from the mean, shown in Figure 9. That is, these subjects, as they did with the dot-cloud sequences, always indicated their heading through the forests to be on the opposite side of their gaze from that of the true heading vector. Upon questioning, these viewers indicated that their path was a hybrid of the two possibilities shown in Figure 8. That is, for the arrangement shown in that figure, they indicated that in the nearground their path was somewhat to the left of the fixated tree but curved in space around behind it, so that in the background it was to the right of the tree. In Experiments 1 and 2, and in all our previous studies, we asked subjects to judge their trajectory with respect to the fixation tree; here, although it occluded everything else on the screen at its location as it moved, the cursor bar was centered at the level of the horizon. Thus, for observers who perceived it to be at the horizon, it would have been entirely consistent to indicate that they were going to the left of the tree, and beyond it to be headed to the right of it when the path stretched to the horizon. It appears, then, that some viewers in the forest situations think that the motion simulated during a trial (at least those as short as 1.25 sec) is generated along a curved, not a straight, path. In the dot clouds, it appears that essentially all viewers think they are on a curved path.

Figure 9. Individual data from the laboratory setting of Experiment 4, plotted in the same manner as in Royden, Banks, and Crowell (1992; Royden, Crowell, & Banks, 1994).

It may seem that not knowing whether one is on a straight or curved path, and not knowing exactly where one is going at all points stretched out to the horizon leaves the pedestrian in a rather precarious situation. Our counter to this is severalfold. First, longer display durations clearly ameliorate this effect. Second, and consistent with directed perception (Cutting, 1986, 1991a, 1991b), we believe that perceivers use multiple sources of information for any given, relatively complex, task. Thus, not only two or more local sources of visual information are available (Cutting, 1996; Cutting et al., 1992), but also perhaps eye-movement information (Royden et al., 1992; Royden et al., 1994) and vestibular information (Berthoz, Israël, Georges-François, Grasso, & Tsuzuku, 1995), and certainly kinesthetic information from the direction of one's feet. Third, the monitoring of one's progress through the environment is ongoing. One can remember headings and trajectories over a period of time and update them as needed. Thus, gazes away from the heading vector, at least for brief periods of time, do not mean that the heading vector is continuously available; instead, they may mean that no update is needed because one can remember the heading direction for that period of time. Only in situations of landing an airplane, of driving a car at highway speeds, of downhill skiing, or of high-speed running is continuous monitoring of heading needed.

Overview

As is shown in the panels of Figure 7, there is a difference between the pattern of perceived heading for travel through forests and that for dot clouds. Simulated travel through forests generates the impression that the perceived aimpoint is to the appropriate side of gaze and nominally correct. Simulated travel through dot clouds, on the other hand, consistently yields the impression that heading is to the side opposite the true heading vector. The major variable accounting for the difference in results between our studies and those of Royden et al. (1992; Royden et al., 1994) and others, then, appears to be the nature of the simulated environments. The data of W. H. Warren, Li, Ehrlich, Crowell, and Banks (1996) corroborate this idea.

CONCLUSIONS

The results of these four experiments allow us to draw several conclusions. First, on the basis of Experiment 1, we note that observers can easily follow instructions and fixate on the tree at the center of the screen while performing a wayfinding task. Indeed, at least in these environments, their performance is somewhat better if they do so. Such a result can serve as a partial justification for the framework of our previous research (Cutting, 1986; Cutting et al., 1992; Vishton & Cutting, 1995).

Second, on the basis of Experiment 2, we found that a reasonably large portion of the visual field appears to be necessary for wayfinding tasks with pursuit-fixation displays such as ours. Within the limitations of a display de-

vice of $20^\circ \times 25^\circ$, apertures centered at the fovea must be at least 16° in diameter, and central field masks must be no larger than 2° , for 95% correct performance to be attained at the simulated velocities. When the pattern of data is scaled to retinal sensitivity as defined by the motion-detection function, the differences between apertures and central-field mask conditions disappear. Within the framework of our previous results, we believe that these results reflect the probability that one or more sources of information will occur and be registered by the visual system.

Third, on the basis of Experiments 3 and 4, we found that observers could consistently indicate on which side of their gaze the heading vector lay, but they could not actually indicate its precise location, biasing it consistently toward fixation. Thus, as gaze-movement angle increases, there are increasingly good nominal, but increasingly poor and biased absolute, judgments of the heading vector. In addition, we suggest that the extreme aimpoint inaccuracies and curvilinear trajectories found by Royden et al. (1992; Royden et al., 1994) in their simulated pursuit-fixation displays may be due less to the necessity of extraretinal information for wayfinding than to their choice of stimulus environments.

REFERENCES

- ANDERSEN, G. J., & BRAUNSTEIN, M. (1985). Induced self-motion in central vision. *Journal of Experimental Psychology: Human Perception & Performance*, **11**, 122-132.
- BEER, J. M. A. (1993). Perceiving scene layout through an aperture during visual simulated self-motion. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 1066-1081.
- BERTHOZ, A., ISRAËL, I., GEORGES-FRANÇOIS, P., GRASSO, R., & TSUZUKU, T. (1995). Spatial memory of body linear displacement: What is being stored? *Science*, **269**, 95-98.
- BERTHOZ, A., PAVARD, B., & YOUNG, L. R. (1975). Perception of linear horizontal self-motion induced by peripheral vision (linearvection). *Experimental Brain Research*, **23**, 471-489.
- BRANDT, T., DICHGANS, J., & KOENIG, E. (1973). Differential effects of central and peripheral vision on egocentric and exocentric motion perception. *Experimental Brain Research*, **16**, 476-491.
- CROWELL, J. A., & BANKS, M. S. (1993). Perceiving heading with different retinal regions and types of optic flow. *Perception & Psychophysics*, **53**, 325-337.
- CUTTING, J. E. (1986). *Perception with an eye for motion*. Cambridge, MA: MIT Press.
- CUTTING, J. E. (1991a). Four ways to reject directed perception. *Ecological Psychology*, **3**, 25-34.
- CUTTING, J. E. (1991b). Why our stimuli look as they do. In G. Lockhead & J. R. Pomerantz (Eds.), *Perception of structure: Essays in honor of Wendell R. Garner* (pp. 41-52). Washington, DC: American Psychological Association.
- CUTTING, J. E. (1993). Perceptual artifacts and phenomena: Gibson's role in the 20th century. In S. Masin (Ed.), *Foundations of perceptual theory* (pp. 231-260). Amsterdam: Elsevier.
- CUTTING, J. E. (1996). Wayfinding from multiple sources of local information in retinal flow. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 1299-1313.
- CUTTING, J. E., FLÜCKIGER, M., BAUMBERGER, B., & GERNDT, J. (1996). Heading information from retinal flow in naturalistic environments [Abstract]. *Investigative Ophthalmology & Visual Science*, **37**, S454.
- CUTTING, J. E., SPRINGER, K., BRAREN, P. A., & JOHNSON, S. H. (1992). Wayfinding on foot from information in retinal, not optical, flow. *Journal of Experimental Psychology: General*, **121**, 41-72.

- CUTTING, J. E., & VISHTON, P. M. (1995). Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth. In W. Epstein & S. Rogers (Eds.), *Perception of space and motion* (pp. 69-117). San Diego: Academic Press.
- CUTTING, J. E., VISHTON, P. M., & BRAREN, P. A. (1995). How we avoid collisions with stationary and moving objects. *Psychological Review*, **102**, 627-651.
- DICHGANS, J., & BRANDT, T. (1978). Visual-vestibular interaction: Effects of self-motion perception and postural control. In R. Held, H. Leibowitz, & H.-L. Teuber (Eds.), *Handbook of sensory physiology: Vol. 8, Perception* (pp. 755-804). New York: Springer-Verlag.
- GIBSON, J. J. (Ed.) (1947). *Motion picture testing and research* (Army Air Forces Aviation Psychology Research Reports, No. 7). Washington, DC: U. S. Government Printing Office.
- GOLDSTEIN, E. B. (1987). Spatial layout, orientation relative to the observer, and perceived projection in pictures viewed at an angle. *Journal of Experimental Psychology: Human Perception & Performance*, **13**, 78-87.
- HILDRETH, E. C. (1992). Recovering heading for visually-guided navigation. *Vision Research*, **32**, 1177-1192.
- JOHNSON, C. A., & LEIBOWITZ, H. W. (1979). Practice effects for visual resolution in the periphery. *Perception & Psychophysics*, **25**, 439-442.
- JOHNSTON, I. R., WHITE, G. R., & CUMMING, R. W. (1973). The role of optical expansion patterns in locomotor control. *American Journal of Psychology*, **86**, 311-324.
- KIM, N.-G., GROWNEY, R., & TURVEY, M. T. (1996). Optical flow not retinal flow is the basis of wayfinding by foot. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 1279-1288.
- LEIBOWITZ, H. W., JOHNSON, C., & ISABELLE, E. (1972). Peripheral motion detection and refractive error. *Science*, **177**, 1207-1208.
- LEWELLYN, K. R. (1971). Visual guidance of locomotion. *Journal of Experimental Psychology*, **91**, 245-261.
- MOTORCYCLE SAFETY FOUNDATION (1992). *Evaluating, coaching, and range management instructor's guide*. Irvine, CA: Author.
- PERRONE, J., & STONE, L. (1994). A model of self-motion estimation within primate visual cortex. *Vision Research*, **34**, 1917-1938.
- PROBST, T., KRAFCZYK, S., BRANDT, S., & WIST, E. R. (1984). Interaction between perceived self-motion and object-motion impairs vehicle guidance. *Science*, **225**, 536-538.
- ROAD RESEARCH LABORATORY (1963). *Research on road safety*. London: Her Majesty's Stationary Office.
- ROYDEN, C. S. (1994). Analysis of misperceived observer motion during simulated eye rotations. *Vision Research*, **34**, 3215-3222.
- ROYDEN, C. S., BANKS, M. S., & CROWELL, J. A. (1992). The perception of heading during eye movements. *Nature*, **360**, 583-585.
- ROYDEN, C. S., CROWELL, J. A., & BANKS, M. S. (1994). Estimating heading during eye movements. *Vision Research*, **34**, 3197-3214.
- ROYDEN, C. S., & HILDRETH, E. C. (1996). Human heading judgments in the presence of moving objects. *Perception & Psychophysics*, **58**, 836-856.
- SEKULER, R., & BLAKE, R. (1994). *Perception* (3rd ed.). New York: McGraw-Hill.
- STOFFREGEN, T. A. (1985). Flow structure versus retinal location in the optical control of stance. *Journal of Experimental Psychology: Human Perception & Performance*, **11**, 554-565.
- STOFFREGEN, T. A. (1986). The role of optical velocity in the control of stance. *Perception & Psychophysics*, **39**, 355-360.
- STOFFREGEN, T. A., & RICCIO, G. E. (1990). Responses to optical looming in the retinal center and in the periphery. *Ecological Psychology*, **2**, 251-274.
- VAN DEN BERG, A. V. (1992). Robustness of perception of heading from optic flow. *Vision Research*, **32**, 1285-1296.
- VAN DEN BERG, A. V., & BRENNER, E. (1994a). Humans combine the optic flow with static depth cues for robust perception of heading. *Vision Research*, **34**, 2153-2167.
- VAN DEN BERG, A. V., & BRENNER, E. (1994b). Why two eyes are better than one for judgments of heading. *Nature*, **371**, 700-702.
- VISHTON, P. M., & CUTTING, J. E. (1995). Wayfinding, displacements, and mental maps: Velocity fields are not typically used to determine one's aimpoint. *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 978-995.
- VISHTON, P. M., NUHAWAN, R., & CUTTING, J. E. (1994). Moving observers utilize static depth cues in determining their direction of motion [Abstract]. *Investigative Ophthalmology & Visual Science*, **35**, 2000.
- WARREN, R. (1976). The perception of egomotion. *Journal of Experimental Psychology: Human Perception & Performance*, **2**, 448-456.
- WARREN, W. H., & HANNON, D. J. (1990). Eye movements and optical flow. *Journal of the Optical Society of America A*, **7**, 160-169.
- WARREN, W. H., & KURTZ, K. J. (1992). The role of central and peripheral vision in perceiving the direction of self-motion. *Perception & Psychophysics*, **51**, 443-454.
- WARREN, W. H., LI, L. Y., EHRLICH, S. M., CROWELL, J. A., & BANKS, M. S. (1996). Perception of heading during eye movements uses both optic flow and eye position information [Abstract]. *Investigative Ophthalmology & Visual Science*, **37**, S454.
- WARREN, W. H., MESTRE, D. R., BLACKWELL, A. W., & MORRIS, M. W. (1991). Perception of circular heading optical flow. *Journal of Experimental Psychology: Human Perception & Performance*, **17**, 28-43.
- WARREN, W. H., MORRIS, M. W., & KALISH, M. (1988). Perception of translational heading from optical flow. *Journal of Experimental Psychology: Human Perception & Performance*, **14**, 644-660.
- WARREN, W. H., & SAUNDERS, J. A. (1995). Perceiving heading in the presence of moving objects. *Perception*, **24**, 315-331.

NOTES

1. Some might call these "full-cue" environments, a term that is apt and brief. Nonetheless, we generally do not like the term "cue" because of its connotations of probabilism in perception (see Cutting, 1986, 1993), and instead typically substitute the term "source of information." Given the lack of economy of our term, we will call these "full-cue" settings by the term "naturalistic environments." We well recognize that, in some absolute sense, our simulated environments are barely closer to reality than those used by any others; yet because we employ multiple sources of information in our displays, whereas others generally have not, and because we think that it is important to do so (Cutting, 1991a, 1991b), we think that by use of the term we are suggesting the direction toward which we should all move.

2. Cutting et al. (1992) referred to these two sources of information as *differential motion parallax* and *inward motion*. Vishton and Cutting (1995), however, discovered that it is not motion, but displacement, in the visual field that is the carrier of this wayfinding information. Thus, they changed the names to *differential parallax displacements* and *inward displacements*. Cutting (1996) then reanalyzed the data of Kim, Growney, and Turvey (1996) and found that the displacement of the largest (or nearest) object in the visual field was a better predictor than differential parallax displacements, but that inward displacements and outward deceleration were still needed to explain the data.

3. Cutting et al. (1992, Experiment 2) measured reaction times and suggested indirect evidence for information about the absolute location of the aimpoint available in the relative velocities of near trees in the environment, but they provided no direct test of the use of this information. Here we suggest that there may be no direct information. We believe that most aimpoint estimation tasks are really little different than scaling tasks, where the observer scales eccentricity judgments of the heading vector against by velocity of motion he or she sees.

4. The Motorcycle Safety Foundation (1992) suggested that "Riders steer in the direction they are looking. Most riders have experienced situations in which they were unable to avoid hitting an object or defect in the roadway because their gaze was fixed on the hazard rather than on the clear path of travel. This phenomenon, sometimes called target fixation, is an example of the need for proper visual direction control" (p. XVI-4).