Limits on integrating motion information across saccades

IBRAHIM DAHLSTROM-HAKKI and ALEXANDER POLLATSEK
University of Massachusetts, Amherst, Massachusetts

In two experiments, we investigated whether people could detect changes in the rotary motion of a cube. A rendering of a cube rotating at a constant angular velocity was presented on a video monitor and, at a key point in the trial, a cross was presented to one side of the cube as a cue for a saccade. On some trials, a change in the rotation occurred either about 100 msec before the saccade or during the saccade; on other trials, there was no change. The change consisted of moving the cube to a new position in the "rotation sequence," after which it continued to rotate at the same angular velocity as before. There was also a control on all trials to ensure that change detection was not due to the detection of low-level motion. Although detection of the change was well above chance when it occurred during the fixation, it was at chance when it occurred during the saccade, except in the case of one participant (who was in both experiments). This chance performance also occurred in Experiment 2 for (1) a slower rotation speed and (2) an axis of rotation that made the rotation planar. The participant who had above chance performance (and as good as that when the change occurred during a fixation) reported using a "strategy" that did not track the path of the cube. It thus appears that there is no natural way in which the visual system tracks this rotary motion, and that detection of change requires some sort of recoding. This finding raises the question of whether good performance in other, apparently similar, motion-detection tasks is a result of similar recoding.

The question of what happens to visual information when the eyes move from one fixation to another is a central one in visual perception. That is, the eyes typically make saccadic movements three to five times a second in viewing conditions that do not involve tracking a moving object, but because our introspections are of a continuing visual world (and we are usually unaware of saccades), some information about the visual world is maintained from fixation to fixation; otherwise we could not function. Our introspections also suggest that detailed visual information is maintained from one fixation and somehow integrated with the visual information from the next fixation. However, many studies in the last 40 years have shown that (1) much less information is retained from a fixation than one might think, and (2) it is not clear whether much, if any, of this information is sensory (i.e., preserves a snapshot-like image of what had been seen on the last fixation).

The question of how much information is maintained from fixation to fixation has usually been studied by one of two paradigms, both involving display changes during a saccadic eye movement between two fixations (Rayner,

1998). The first is a change-detection paradigm, in which the participant is asked to make a judgment about whether or not a change occurred during the saccade. The second is a "preview-benefit" paradigm, in which the participant is not asked to make a conscious decision about whether there has been a change; instead, the processing of the stimulus following the saccade is assessed, and the key question is how this processing is affected by information from the same spatial location seen on the previous fixation. Although the present experiments are change-detection experiments, we will review relevant findings from both these paradigms.

Beginning with research as early as the 1950s, experiments investigating the ability of participants to correctly detect intrasaccadic change indicated time and again that humans are surprisingly bad at detecting these changes when they occur in parts of a scene that they are not specifically attending to. Ditchburn (1955), and later Wallach and Lewis (1966), ran experiments designed to investigate the mind's ability to maintain a coherent, stable image of the world in the face of constantly changing visual input. Participants in these experiments were asked to detect changes in their environment and were surprisingly poor at detecting reasonably large displacements of stimuli if those displacements occurred during an eye movement. Subsequent experiments by Mack (1970), Rayner (1975), Bridgeman, Hendry, and Stark (1975), and others showed poor detection of saccade-contingent change (of both form and location) for a number of different types of stimuli, ranging from written text to simple targets. In each of these experiments, participants were presented with a visual stimulus that changed during a saccade and, in each

This research was supported by Grants HD17246 and HD26765 from the NIH, KDI Grant 9980045 from the NSF, and by a grant from the General Electric Fund. This work was part of the first author's master's thesis, and he thanks the members of his committee (Alexander Pollatsek, Keith Rayner, and Donald Fisher) for their valuable input, and his wonderful wife, Añira Dahlstrom-Hakki, for her love and support. Correspondence concerning this article should be sent to A. Pollatsek, Department of Psychology, University of Massachusetts, Amherst, MA 01003 (e-mail: pollatsek@psych.umass.edu).

case, they were rarely able to report detecting a change in the stimulus if that change occurred during a saccadic eye movement.

Although some sort of information is retained, because change detection is possible, it is not at all clear in what form the information is stored. Even though it probably seems obvious to most laypeople that one is continuously maintaining a detailed representation of the visual world, this actually seems unlikely, and a closer look at what it would entail to realize such a system makes it clear that maintaining such a detailed representation would be impractical.¹ First, maintaining a detailed representation would require the brain to dedicate a large amount of resources to a task of little potential benefit. Because people make three to five fixations a second, the amount of visual information one receives in only a minute is immense—even in a static environment. Thus, it seems fairly inefficient to store all that information in the brain; most of it will be of little use, because most of it can be reacquired by resampling the visual environment. Second, one would need to integrate all this detailed visual information from all these fixations, which would be quite difficult, if not impossible. Among other things, saccades rarely land where intended, so somehow one has to determine how the information from a new saccade aligns with the image from the preceding fixation. Moreover, these images would not be identical, because the visual acuity of the region from which new information is to be added to an existing image is impoverished, making a match between the two extremely difficult. In addition, the system would have to be able to handle changes in lighting or perspective (or both) between different fixations and integrate that information in a few milliseconds.

Thus there is a general consensus that there is a more abstract “visual” representation of the world that is maintained across fixations (Jonides, Irwin, & Yantis, 1983; Pollatsek & Rayner, 1992; Rayner & Pollatsek, 1983). However, it is far from clear what this representation is. Some clues come from the work on preview benefit. In this paradigm, a visual object is presented in the parafovea, the object location is then fixated, and the fixated object is named. The key trials are those on which the parafoveal object is different from the fixated object.

The original studies were in reading (Rayner, 1975) and demonstrated that an orthographically similar word seen in the parafovea facilitated the time to name the fixated target word. However, a change of case in the letters between fixations made no difference, again indicating that the representation was not “visual” (McConkie & Zola, 1979; Rayner, McConkie, & Zola, 1980). Moreover, work using homophones (see, e.g., Pollatsek, Lesch, Morris, & Rayner, 1992) has indicated that the phonological overlap of the stimuli is also important, so that the representation need not be purely orthographic. Other work, however, has indicated that a semantically related word in the parafovea did not speed naming when compared to the time needed to name an unrelated word (Rayner, Balota, & Pollatsek, 1986; see also Altarriba, Kambe, Pollatsek, & Rayner, 2001). It is also worth noting, with respect to

change detection, that people are usually unaware of these letter changes across saccades unless the number of letters increases or decreases.

The results in experiments that used line drawings of pictures came up with a similar pattern. The original study by Pollatsek, Rayner, and Collins (1984) showed that the representation was not literal, because a mirror image of the object or a size change of the object during the saccade had minimal effects on naming time. However, there was evidence that something more than a semantic representation of the object was maintained, because (1) there was some facilitation when the parafoveal object had a similar general shape as the foveal object (e.g., a baseball bat as a preview for a banana) and (2) there was no evidence of semantic priming (e.g., a preview of a banana did not facilitate naming of an apple).

Irwin (1992) suggested that this representation is basically the same one that people think of when they think of (active) “visual short-term memory.” This metaphor suggests that attention to the object may be necessary to maintain the information (and detect changes in it), or that attention at least modulates the maintenance of this information. In fact, there are studies that indicate a substantial role for attention. In all the parafoveal preview studies mentioned above, the parafoveal preview has been found to be attended to prior to fixating. In addition, in a display change experiment with naturalistic scenes, Hollingworth and Henderson (2002) showed that if an object is fixated immediately after the change has occurred, participants are more likely to detect the change than if it is not; presumably the object was attended to and then fixated. Hollingworth and Henderson also showed that the change is more likely to be noticed if it is fixated immediately before the change occurs, which also suggests an attentional component to the detection of change.²

To summarize, the above research from both the change-detection and preview literature suggests that little of the visual detail of what is seen on one fixation is maintained across the saccade to the next fixation, and that, instead, the representation is a rough, schematic one, perhaps largely excitation of visual object detectors that provide information reasonably consistent with the visual information from the parafovea. However, there are some studies looking at saccade-contingent change with moving stimuli that appear to argue for a more veridical representation of the object from the prior fixation.

The first of these experiments were studies in biological motion (Verfaillie, 1997; Verfaillie & De Graef, 2000; Verfaillie, De Troy, & Van Rensbergen, 1994). Participants in these experiments saw a set of moving dots on a computer screen that were positioned in such a way as to mimic the locations of major joints on the human body and that moved so as to give the illusion of a walking man. (Participants had no problem seeing this simple collection of moving dots as a walking man.) In many of the experiments, the location of the “walking man” was unchanged, as if he were walking in place on a treadmill. Most relevant to the current experiments were those that introduced a change in the walking sequence across the saccade. That

is, if one thinks of the walking sequence as a repeating sequence of "still frames," the change was made so that the "frame number" was moved either backward or forward from where it should have been during the saccade, but then after this single change, the walking sequence continued normally from the changed position. (On catch trials, there was no interruption of the normal sequence.) Although performance was not perfect, participants were much better at this task than on many of the other change-detection tasks previously mentioned. Similarly, people were very good at detecting a change in the orientation of the walking dot man (i.e., which direction he was facing) across saccades. They appeared to be somewhat less good at detecting changes in location of the walking dot man when he was walking across the screen (in the prior experiments, he was walking on the treadmill).

In almost all their experiments, Verfaillie and colleagues compared the change-detection performance across a saccade with detection within a fixation. To control for the possibility of people detecting some large motion change during the saccade, there was a large jump in the position sequence for a frame or two on all trials in which the change was made across the saccade, and then either the walking man would be back where he should be on the control trials or starting in a new position on the change trials. Detection of change in the walking sequence and change in the orientation of the figure were not as good across saccades as within a fixation, but performance was still quite high; detection of change in location was notably poorer across saccades than within a fixation, though still well above chance. (Note that there was no such control in the within-fixation condition, which could explain why performance was somewhat better in the within-fixation condition than in the between-fixation condition.)

One question, of course, is whether there is something special about biological motion. To test this, Pollatsek and Rayner (2002) used a fairly different stimulus in a similar experimental paradigm: a line that rotated at a constant angular velocity about a fixed point in the center of the display. The change consisted of a "jump" in the location of the line to a position ahead or behind that of the current location, followed by continued rotation at the same velocity—fairly analogous to the change in the walking dot man. In Experiment 2, when a control was instituted both in the within-saccade condition and in the within-fixation condition to ensure that the jump in the sequence was not detected by low-level detection of motion, detection of the change in the rotation of the line was well above chance, and as good across fixations as it was within a fixation (as measured by d'). (In this case, the control was a static display of the line in all its positions for two frames.)

Another series of experiments by Gysen and colleagues (Gysen, De Graef, & Verfaillie, 2002; see also Gysen, Verfaillie, & De Graef, 2002a, 2002b) also demonstrated fairly good detection of change of a geometric object (a frustum of a cone). Gysen and colleagues directly compared participants' intrasaccadic change-detection performance for static and translating stimuli. They found

some differences in the pattern of how well intrasaccadic change was detected between these two types of stimuli. They also found that detection of displacement change was usually better in the translating condition than in the static condition. Most relevant to the current experiments, they found equal (and above chance) performance in detecting the orientation change for static and rotating stimuli. The speed of rotation in these experiments was somewhat slower than for the walking man and the rotating line experiments. The rotation speed in the Gysen et al. (2002) experiments was about 110°/sec and the orientation change was a rotation of 45°, whereas the rotation speed of the line in the Pollatsek and Rayner (2002) experiment was 167°/sec and the orientation changes were 50° and less.³ (The motion in Verfaillie et al., 1994, is not a simple rotation. However, its speed seems roughly comparable to the others, because the time to complete one cycle is 1.6 sec, compared with 2.16 sec in Pollatsek & Rayner, 2002.)

The findings above with stimuli that are moving in some sort of pattern (either repetitively going through a walking sequence or a rotation sequence) suggest that perhaps some sort of more detailed representation is preserved across fixations. That is, it appears that, in order to be able to perceive a "bad continuation" of a pattern of motion, one has to have a fairly detailed representation of what the motion was before the saccade. However, as in all the change-detection experiments, one doesn't have much control over the level of coding. For example, if there is a repetitive sequence, what may be encoded is a temporal sequence, such as the times when the right foot of the dot man hits the ground or when the rotating line is vertical, and then what is noticed is that the temporal sequence is disrupted across the saccade. We will return to this point later. However, we want to emphasize that such a recoding need not be a conscious strategy.

The present experiments were designed to explore the reasons behind good detection of a change in moving stimuli by probing the limits of the phenomenon. In particular, we wanted to understand whether good detection of change in rotary motion would still occur with a stimulus that was a representation of a 3-D object rotating in depth: a computer rendering of a rotating 3-D cube. Although the cube is still a fairly simple figure, it differs from the types of stimuli that have been used thus far in some significant ways. The walking man stimulus is a 2-D representation of apparent biological motion. The rotating line is a very simple 2-D shape providing a minimal amount of visual information. The fully rendered cube, however, does still represent a potential real life shape that one is likely to observe in the environment. In addition, the lighting and the texture effects of the rendered cube provided a richer, and possibly harder to process, wealth of visual information than did the rotating line or the walking man dots. Although the rotating frustum of a cone provided a rendering of a 3-D stimulus that is comparable to the cube used in these experiments, we believe that the good performance with regard to detection of

change in the path of the frustum's rotation may not have been a result of perceptual processes. (We will return to this point in the General Discussion section.)

EXPERIMENT 1

As indicated above, the aim of this experiment was to investigate participants' ability to detect saccade-contingent changes in the rotating cube paradigm outlined above and to compare that performance with their ability to detect the same changes when they occurred within a fixation. Because people were surprisingly good at detecting such changes in the gait of a walking man (Verfaillie et al., 1994) and the rotation of a simple line (Pollatsek & Rayner, 2002), we thought it plausible that people would be able to detect fairly analogous changes in a rotating cube.

Method

Participants. Eight participants were recruited from the University of Massachusetts community to participate in this experiment. They all had normal or corrected-to-normal vision, and were offered either pay or course credit for their participation.

Apparatus. Each participant viewed the stimuli on a 19-in. NEC Trinitron monitor connected to a 2.4 GHz Pentium 4 PC. The refresh rate of the monitor was 160 Hz. Participants were required to wear a lightweight helmet that is part of the EyeLink II SR eyetracking system. The eyetracking system samples at the rate of 500 Hz and provides eye movement data both for data analysis and to the display PC for the purposes of conducting display-contingent changes. No head restraint was used because the EyeLink II system is able to compensate for head movements. Eye movement data were collected from the right eye.

Stimuli and Design. The overall screen resolution was set at 800×600 pixels, and each stimulus was presented at the center of that display. The stimulus was a blue rotating cube that was fully rendered and set against a black background. The rotation of the cube can be described as a succession of "frames," analogous to movie frames, with a representation of the cube at a particular orientation being presented on each frame. The duration of a frame was 12 msec, which was two refresh periods for the video display. The orientation of the cube changed 3° from frame to frame, so that the cube rotated 3° per 12 msec or 250° per sec. During a trial, the cube rotated about an axis that ran through its topmost and bottommost corners. The cube appeared in one of two different orientations (see Figure 1) an equal number of times, randomly distributed across all trials. Orientation A allowed the cube to rotate about the vertical axis with its topmost and bottommost corners on the axis of rotation. Orientation B was achieved by rotating the cube described in the first condition by 30° in the z -plane.

A change in the rotary motion of the stimulus occurred on three quarters of the trials. On these trials, the cube underwent a one-time jump of $\pm 12^\circ$, 24° , or 48° , and continued its rotation from its new location at the same rate of rotation it had before the jump. A control for low-level motion detection was introduced during the 12-msec interval directly preceding the jump, in trials in which a change occurred, and at a corresponding point in time, in trials in which no change occurred. Note that although frames were incremented every other refresh period during the cube's normal motion, the change was initiated at the beginning of the first possible refresh period.

On each trial, one of two possible controls for low-level detection of rapid motion was used. The 135° forward control consisted of a 135° jump ahead of where it should have been in the cube's rotation (or 138° ahead of the previous frame) that was followed by a return to its original trajectory in the case of the no-change condition

(which would be a 132° jump back), or to the new location in the case of the change conditions. (For a 24° forward change, this would be 30° ahead of where it was before the jump.) The 135° return control was identical to the forward control condition in the no-change condition. However, in the change condition, it consisted of a jump ahead in the cube's rotation of a magnitude so that the return jump in the following frame was 132° . For example, in the 135° return control condition, for a 24° forward change, the control jump ahead would be 162° (note that this includes an extra 3° to compensate for the frame that is taken up by the control). As indicated above, the purpose of both types of controls was to eliminate the use of low-level motion detection to determine whether there was a change (especially when the change was within a fixation).

A trial consisted of the following sequence of events. First, a fixation cross appeared in the center of the screen for a calibration check, and a drift correction was applied, if needed. Then the experimenter initiated the trial, and a cross and a rotating cube simultaneously appeared on the screen. The cross was 6° to the left of the center of the screen and the cube was centered on the screen. (The cube subtended approximately 7° – 8° of visual angle, its exact size depending on its orientation.) Participants were instructed to make a saccade to the cross location. After a preset, varying amount of time (but after we ensured that the left-hand cross was indeed fixated) the cross then jumped 12° to the right. Participants were instructed to make a saccade to the new cross location as soon as possible. On half the trials, the change occurred during the induced saccade, whereas on the other half of the trials, the change occurred 100 msec prior to the appearance of the second cross and hence at least 200 msec before the saccade. The saccade-contingent changes were initiated 4 msec after the eyetracker had determined that the saccade had crossed a boundary that was $.65^\circ$ to the right of the initial fixation cross location. At the end of each trial, participants indicated using a response box whether there was a change in the motion of the stimulus during the trial. They were also instructed to try to ignore the motion caused by our motion control.

Subsequent to running the experiment, the eye movement record was examined in order to ensure that the changes in the between-fixation condition had occurred during the saccade. The "jump" stimulus was displayed for two refresh periods (each refresh period was 6 msec) as soon as the onset of the saccade was detected by the system, and then the continuation of the rotating cube (either in its previous path in the no-change condition or starting in a new path in the change condition) started in the next refresh period. Thus, the display change was completed $34 \text{ msec} \pm 3 \text{ msec}$ after the saccade began. An examination of the eye movement record indicated that the saccade durations were $54 \text{ msec} \pm 9 \text{ msec}$ and thus that the completion of the display change (when it occurred) was always well within the saccade duration.

To summarize, each participant saw the cube in (1) two orientation conditions; in each, it rotated about the axis running through its topmost and bottommost corners, (2) two jump directions, either forward or backward, (3) two control conditions, and (4) four angles of rotation, including the no-change condition. (The forward and backward versions of the no-change condition were, of course, identical.) Each of these appeared in either the within-fixation condition or the between-fixation condition five times, for a total of 320 experimental trials. A break was introduced in the middle of the experiment to give participants a chance to rest. There were also 10 practice trials at the beginning of the experiment in order to familiarize participants with the task, the equipment, and the stimulus. These practice trials were preceded by a demonstration in which participants were shown what the cube would look like, and what change and control conditions would look like.

Results

Contrary to our expectations, the data from Experiment 1 indicated that participants were close to chance at

Figure 1. Experiment 1: Rotation axes for Orientation A (left) and Orientation B (right).

detecting the presence of a change in the between-fixation condition. The false alarm rate was 5.3%, and the hit rates were 4.7%, 5.6%, and 11.6% in the 12°, 24°, and 48° jump conditions, respectively [$F(3,21) = 1.19, p > .30$]. Even the pairwise comparison of the hit rate in the 48° jump condition to the false alarm rate was not close to significant [$t(7) = 1.02, p > .30$]. In signal-detection terms, the average d' 's for the 12°, 24°, and 48° conditions were $-.06$, $-.05$, and $.17$, respectively ($ts < 1$).

In contrast, detection of the change was reasonably good in the within-fixation condition. The false alarm rate was 45.3%, and the hit rates were 63.8%, 72.2%, and 89.4% in the 12°, 24°, and 48° jump conditions, respectively [$F(3,21) = 22.99, p < .001$]. The average d' 's for the 12°, 24°, and 48° conditions were $.58$, $.84$, and 1.53 [$t(7) = 3.98, p < .01$; $t(7) = 3.84, p < .01$; and $t(7) = 7.76, p < .001$, respectively].

The other variables had, at most, a minor effect on the participants' abilities to detect change in either the between-fixation or within-fixation conditions. We had thought that Orientation A (with the rotation axis going through corners of the cube) would be more natural and that rotation changes would thus be easier to detect. However, there was virtually no difference between the two orientation conditions in both the between- and within-fixation conditions. For the between-fixation condition, with Orientation A, the false alarm rate was 5.6% and detection rates were 3.8%, 5.6%, and 11.3% in the 12°, 24°, and 48° jump conditions, whereas with Orientation B, the false alarm rate was 5.0% and detection rates were 5.6%, 5.6%, and 11.9% in the 12°, 24°, and 48° jump conditions. Similarly, for the within-fixation condition, with Orientation A, the false alarm rate was 42.5% and detection rates were 64.4%, 73.1%, and 90.6% in the 12°, 24°, and 48° jump conditions, whereas with Orientation B, the false alarm rate was 48.1% and detection rates were 63.1%, 71.2%, and 88.1% in the 12°, 24°, and 48° jump conditions.

In both cases, neither the main effect of orientation nor the interaction of orientation with size of jump (counting no-change as 0° change) was close to significant (F 's < 1). The hit rate for forward jumps was a bit higher than for backward jumps in the within-fixation condition, but there was no significant effect associated with the direction of jump in either the within- or between-fixation conditions. In the between-fixation condition, the hit rates were 6.3%, 5.0%, and 10.6% for the forward jumps, and 3.1%, 6.3%, and 12.5% for the backward jumps, in the 12°, 24°, and 48° jump conditions, respectively ($F < 1$) for the main effect of direction and [$F(2,14) = 1.137, p > .30$], for the direction $\times$ jump size interaction. In the within-fixation condition, the hit rates were 70.6%, 80.0%, and 91.9% for the forward jumps, and 56.9%, 64.4%, and 86.9% for the backward jumps, in the 12°, 24°, and 48° jump conditions, respectively [$F(1,7) = 1.951, p > .20$] for the main effect of jump size and ($F < 1$) for the interaction. There was also no significant difference between the two types of manipulations to control for perception of rapid movement (F 's < 1.1).

The data thus indicated that, averaged over participants, performance was close to chance in the between-fixation condition and well above chance in the within-fixation condition. However, because performance was somewhat above chance in the largest change condition in the between-fixation condition, we wanted to determine whether this was a small effect for most of the participants, or whether there were any participants who were performing substantially above chance. In fact, as indicated in Figure 2, 1 participant (Participant 2) was able to perform significantly above chance in both the within-fixation condition [$F(3,128) = 6.93, p < .001$] and the between-fixation condition [$F(3,128) = 9.49, p < .001$], and the difference in his performance between the within- and between-fixation conditions was marginal [$F(1,256) = 2.80, p > .05$]. Moreover, as illustrated in Figure 3, once

Figure 2. Change detection in the within-fixation and between-fixation conditions for Participant 2.

Participant 2 was removed from the analysis, the hint that detection of change was greater than chance in the between-fixation condition disappeared completely. The false alarm rate for the remaining participants was 2.9%, and the hit rates were 2.1%, 1.4%, and 3.2% in the 12°, 24°, and 48° jump conditions, respectively ($F < 1$).

Given the high level of performance Participant 2 was able to achieve, even in the between-fixation condition, we conducted a follow-up interview to ascertain how he was able to perform so well. He articulated a strategy that

allowed him to detect the change, apparently without any need to maintain visual information about the geometry of the cube across the fixations. Because the cube in Experiment 1 was rendered and had lighting and texture effects, Participant 2 was able to detect a peak in the reflection of light off of one of the cube's faces at uniform intervals. This allowed him to construct a "rhythm" of light reflection peaks at uniform intervals. He was able to maintain this rhythm as a mental beat, which helped him determine whether or not the cube had undergone a change by com-

Figure 3. Change detection in the within-fixation and between-fixation conditions, excluding Participant 2.

paring the rhythm maintained mentally with the pattern of light reflection peaks the cube exhibited after the saccade. This enabled him to detect change in the between-fixation condition just about as well as in the within-fixation condition (see Figure 2). Subsequently, Participant 2 was asked to participate in Experiment 2, to determine whether or not he would be able to obtain similar results with the cube rotating about an axis that might not allow use of similar light cues.

EXPERIMENT 2

Contrary to the results of the “walking man” and rotating stick experiments, detection of change for the stimulus in Experiment 1 proved to be essentially impossible for all but 1 of the participants when the change occurred during a saccade. In Experiment 2, we investigated whether there was some detail of the stimulus condition in Experiment 1 that hindered change detection, or whether detection of change in this paradigm, unlike performance on earlier paradigms involving moving stimuli, is close to impossible without the aid of a conscious strategy specifically designed to detect that change. Specifically, Experiment 2 tested whether either (1) a slower rotation speed or (2) a “simpler” axis of rotation might improve the detection of change in rotation across fixations. Our hypothesis about speed was largely driven by intuition. It seemed that a slower rotation was, in some sense, easier to follow. The simpler axis of rotation that we investigated was the one that appears to be perpendicular to the screen. This means that the actual rotation on the screen is a 2-D object rotating around an axis parallel to the line of sight. As a result, this makes it much closer to the rotating line stimulus employed by Pollatsek and Rayner (2002). A second motivation for trying a planar rotation of the cube was to test Participant 2’s strategy, because the 2-D rotation did not provide the same light reflection patterns as did Experiment 1.

Method

Participants. Eight participants were recruited from the University of Massachusetts community to participate in this experiment. Participant 2 from Experiment 1 was invited to participate in Experiment 2, and was designated as Participant 2 in this experiment as well. All participants had normal or corrected-to-normal vision, and were offered either payment or course credit for their participation.

Apparatus. The apparatus used was identical to that used in Experiment 1.

Stimuli and Design. Experiment 2 used stimuli that were very similar to those used in Experiment 1. The overall screen resolution was set at 800×600 pixels, and the stimulus was presented at the center of that display. The stimulus in each case was a blue rotating cube that was fully rendered and set against a black background. In Experiment 2, the cube rotated about two different axes (see Figure 4): the y -axis rotation referred to as Orientation A in Experiment 1, and a z -axis rotation (i.e., a rotation in one’s visual plane in an orientation allowing the cube to rotate with one of its corners in the center of the display). Each of these rotations appeared in two different speed conditions, either at the “fast” rate of 3° every 12 msec, or at the “slow” rate of 1.5° every 12 msec.

A change in the rotation path of the stimuli in Experiment 2 occurred on half of the trials. As in Experiment 1, the change consisted of a jump in the sequence of rotation of the cube, but unlike in Experiment 1, there was only one (large) jump condition, the 48° forward jump condition of Experiment 1. In addition, since no significant difference was found between the two different types of control used in Experiment 1, we only employed the 135° forward control in Experiment 2.

Because of hardware constraints, our computer system was unable to hold enough frames in memory to display rotations across both the y - and z -axes without reloading the program. As a result, the axis of rotation was blocked and counterbalanced, with half the participants seeing 160 y -axis rotations followed by 160 z -axis rotations and the other half seeing the opposite order. In addition, speed of rotation was blocked and counterbalanced within the blocking of axis of rotation.

On half the trials, the change occurred during an induced saccade, whereas on the other half of the trials, the change occurred 100 msec prior to the induced saccade. The saccade was induced in the same manner as in Experiment 1, and as in Experiment 1, saccade-contingent changes were introduced only after the eyetracker had determined that the saccade had indeed begun. At the end of each trial, participants indicated using a response box whether or not they

Figure 4. Experiment 2: The y -axis of rotation (left) and the z -axis of rotation (right).

were able to detect a change in the stimulus during the trial. The within-fixation and between-fixation conditions were randomly presented within a block, as were the jump and no-jump conditions. At the end of each block of the experiment (i.e., after every 80 trials), participants were given a break. Each block was identical to the others except that either the axis of rotation, the speed of the rotation, or both changed.

To summarize, (1) each participant saw the cube in two rotation conditions, one about the y -axis, and one about the z -axis, (2) the cube underwent either a 0° or 48° jump, (3) half the trials rotated at a "slow" rate of 1.5° every 12 msec, whereas the other half rotated at the "fast" rate of 3° every 12 msec, and (4) each change occurred in either the within-fixation condition or the between-fixation condition. Each of these types of stimuli was presented 20 times. The 320 experimental trials were in four blocks of 80 trials each, one block for each of the four conditions defined by rotation axis and speed. There were also 10 practice trials at the beginning of the experiment in order to familiarize participants with the task, the equipment, and the stimulus. These practice trials were preceded by a demonstration in which participants were shown what the cube would look like, and what change and control conditions would look like.

Results

As in Experiment 1, participants were able to reliably detect the presence of a change in the motion of the cube in the within-fixation condition (18.9% false alarm rate, 86.9% hit rate) [$F(1,7) = 104.5, p < .001$]. However, in the between-fixation condition, they were only slightly more likely to say that a change had occurred in the no-change condition than when a change had indeed occurred (4.7% false alarm rate, 11.7% hit rate) ($F < 1$). The mean d' s across participants were 2.35 [$t(7) = 7.77, p < .001$] and 0.21 ($t < 1$) in the within- and between-fixation conditions, respectively.

Although detection was slightly better in the y -axis condition than in the z -axis condition, the effect was far from significant in either the between-fixation or within-fixation conditions. For the between-fixation condition, the false alarm and hit rates in the y -axis condition were 4.7% and 14.4% ($t < 1$) and the false alarm and hit rates in the z -axis condition were 4.7% and 9.1% ($t < 1$). The d' s were .31 and .08, respectively, and the difference between them was obviously not significant ($t < 1$). For the within-fixation condition, the false alarm and hit rates in the y -axis condition were 15.0% and 86.6% [$t(7) = 7.70, p < .001$] and the false alarm and hit rates in the z -axis condition were 22.8% and 87.2% [$t(7) = 6.92, p < .001$]. The d' s were 2.45 and 2.23, and again the difference between them was not significant [$t(7) = 1.28, p > .20$].

Slowing down the rotation speed did appear to improve performance a bit; however, the effect was reliable only in the within-fixation condition. For the between-fixation condition, the false alarm and hit rates in the fast condition (the speed employed in Experiment 1) were 6.9% and 10.0% ($t < 1$), and in the slow speed condition, they were 2.5% and 13.4% [$t(7) = 1.09, p > .30$], and the difference between the two d' s in the two conditions, .06 and .38, was not significant [$t(7) = 1.30, p > .20$]. For the within-fixation condition, the false alarm and hit rates in the fast speed condition were 26.6% and 83.1% [$t(7) = 5.22, p < .01$], the false alarm and hit rates in the slow speed condition were 11.3% and 90.6% [$t(7) = 8.756, p < .001$], and the difference between the d' in the two conditions, 1.89 and 2.87, was significant [$t(7) = 3.13, p < .05$].

Given Participant 2's performance in Experiment 1, we decided to look at his performance more closely in Experiment 2. Among other things, we wanted to determine whether he was responsible for the slightly above chance performance in the between-fixation condition, as he was in Experiment 1. As in Experiment 1, Participant 2 performed very well and about the same in the between-fixation and within-fixation conditions. In the between-fixation condition, his d' was 1.57 [hit rate = 73.7%, false alarm rate = 17.5%; $t(152) = 9.00, p < .001$], and in the within-fixation condition, his d' was 1.51 [hit rate = 76.2%, false alarm rate = 21.3%; $t(152) = 9.92, p < .001$]. When Participant 2 was removed from our analysis, the detection rate in the between-fixation condition went to virtually zero, with a mean ($d' = .01$) ($t \approx 0$) (mean hit rate = 2.9%, mean false alarm rate = 2.9%).

We were also interested in whether Participant 2 was able to perform equally well across the different speed and axis conditions. He was able to detect the presence of a change very well at the slow speed, but less well at the fast speed. At the slow speed, his d' s in the within- and between-fixation conditions were 2.68 [hit rate = 85.0%, false alarm rate = 5.0%; $t(38) = 8.34, p < .001$] and 3.29 [hit rate = 95.0%, false alarm rate = 5.0%; $t(38) = 12.73, p < .001$], respectively, in the vertical (y -axis) rotation condition, and were 3.29 [hit rate = 95.0%, false alarm rate = 5.0%; $t(38) = 12.73, p < .001$] and 2.32 [hit rate = 85.0%, false alarm rate = 10.0%; $t(38) = 7.01, p < .001$] in the planar (z -axis) condition. At the fast speed, his d' s in the within- and between-fixation conditions went down to 0.97 [hit rate = 55.0%, false alarm rate = 20.0%; $t(38) = 2.39, p < .05$] and 2.12 [hit rate = 90.0%, false alarm rate = 20.0%; $t(38) = 6.10, p < .001$] in the vertical rotation condition, and were only .40 [hit rate = 70.0%, false alarm rate = 55.0%; $t < 1$] and -.29 [hit rate = 25.0%, false alarm rate = 35.0%; $t < 1$] in the planar rotation condition.

In a follow-up interview, Participant 2 reported using the same strategy he reported using in Experiment 1 for the y -axis rotation conditions: establishing a mental "rhythm" corresponding to light reflection peaks, and comparing that rhythm to the post-saccadic pattern of light reflection peaks. However, he attempted to use a slightly different strategy in the z -axis conditions. The strategy consisted of the establishment of a reference point along the cube's perimeter of rotation. As each corner of the cube crossed that reference point, he was able to establish a beat and maintain that beat as a mental rhythm. He then compared this mental rhythm with the beat exhibited by the cube's rotation following the saccade to determine whether or not a change had occurred. Participant 2 reported that he had more difficulty with this strategy than with the strategy he employed with the y -axis rotations, and found it very

difficult to apply it to the “fast” planar rotation condition, the condition in which he performed poorly. However, he found it difficult to express why this was the case.

DISCUSSION

In the introduction, we discussed research that suggested that moving stimuli may be coded differently than non-moving stimuli so that the trajectory of a path of motion might be integrated across a saccade, whereas many of the visual details of static objects are not maintained across saccades. Quite obviously, the results obtained in the present experiments were fairly different from those obtained in earlier research on moving stimuli, because even large changes in the motion of the cube were not detected at all by 7 of the 8 participants in each of the experiments. Moreover, the 1 participant who did achieve good levels of performance in most conditions apparently did so by using a conscious strategy in which the rotary motion of the object was coded into a rhythm (of the recurrence of a periodic event) that had little to do with maintaining the visual representation of the form of the rotating object. This indicates that, for stimuli of this complexity, there is little information about the rotation that is preserved over the saccade that is integrated with information from the next fixation.

We have no firm hypotheses for why performance in the rotating cube task is so different from that in the rotating line task (Pollatsek & Rayner, 2002), the dot representation of a walking man task (Verfaillie, 1997; Verfaillie & De Graef, 2000; Verfaillie et al., 1994), or the rotating frustum of a cone task (Gysen et al., 2002). One possibility is that there is something crucially different in the experimental setup (most plausibly, either the type of video display or the type of “control” for motion). The most likely difference in procedure between the present experiment and the other studies is the type of control (i.e., the control for low-level motion detection during a fixation may have been less effective in the present experiments than in the ones cited above). First, the control in Experiment 2 of Pollatsek and Rayner employed a bright display of all the lines, which was quite disruptive (and produced quite high false alarm rates). Second, because the display rate was faster in the current experimental setup than in that of Verfaillie and colleagues, the move to the control stimulus and back to the rotating cube was also faster, and may have been less disruptive. However, although some differences between the motion-control manipulation in the present experiment and the other experiments might explain why performance was better in the within-fixation condition than in the between-fixation condition, it doesn’t explain why performance was at chance for 7 of the 8 participants in the between-fixation condition. If anything, a less effective control condition should have made the change in rotation easier to detect. Thus, we think that this difference in procedure is unlikely to explain the difference between the experiments.⁴

A second possibility is that the rotation of the cube is harder to perceive—whether within or between fixations—than the other types of motion studied. However, within-fixation performance in Experiment 2 was actually quite a bit better than that in Experiment 2 of Pollatsek and Rayner (2002). The d' s in Experiment 2 in the y -axis and x -axis rotation conditions were 2.43 and 2.23 (averaged over rotation speed), whereas the d' in the largest rotation condition in Experiment 2 of Pollatsek and Rayner (2002) was 1.10. (In Experiment 1 of Pollatsek & Rayner, 2002, a much higher d' was observed in the within-fixation condition, but there was no control for low-level motion detection.)

This indicates that there is likely a deeper reason for the difference between our findings and the earlier studies investigating detection of changes in motion across saccades. One possibility is that the rotating cube paradigm provides a richer wealth of visual information and is, in some sense, “harder to digest.” Certainly, the rotating line is a much simpler stimulus than the rotating cube. However, whether the dot walking man or a rotating frustum of a cone is simpler than the rotating cube is less clear. A second, perhaps related, possibility is that rotation in depth is less related to the information obtained from retinal motion detectors than is 2-D motion, and that 3-D motion is therefore in a different visual system module than the one for 2-D motion. The case may be made that it is the 2-D motion module that subserves detection of changes in the experiments using the rotating stick or the walking man. This hypothesis, however, runs into some problems because of the z -axis rotation conditions in Experiment 2. That is, in z -axis rotation conditions, the low-level motion that would be picked up by motion detectors should have been as closely related to the physical motion of the cube as for the other types of figures that produced good detection of change across saccades in previous experiments. However, this planar motion may have been much more complex and harder to deal with for the cube than for either of the other stimuli. (See the first hypothesis.)

This raises the possibility that the difficulty in detecting change in the cube rotation task is an attentional one (especially in the planar rotation condition). That is, there may be 2-D motion cues (such as the rotation of one of the edges of the 2-D screen representation of the cube) that, if attended to, would allow well above chance detection of change in rotation across fixations. However, the “natural strategy” might be to attend to the 3-D representation, and changes in this representation might be close to impossible to detect.

This discussion raises the question of whether the record of physical motion of this type is ever stored across saccades by the visual system and compared to the motion on the new fixation as an automatic feature of the visual system. The above two paragraphs suggest that this is a possibility, but that somehow, the computation succeeds for rotating sticks, walking men, and rotating cone frusta, but fails for cubes. There is another possibility, however, suggested by the performance of Participant 2. That is, it

could be that the computations necessary to detect these types of changes in the sequence of visual motion are never computed by the visual system. One could argue that, evolutionarily, there is no need to make such computations (which are likely to be difficult and perhaps error-prone) when saccades are brief and fixations are long. That is, any such changes in motion trajectory are extremely unlikely to occur only during a saccade, and when they occur during a fixation, they are likely to be detectable by much simpler computations. If so, then the chief difference between the present paradigm and the other paradigms may be that coming up with a strategy such as that adopted by Participant 2, of noticing periodicity, is much easier in the other situations. As indicated earlier, this periodicity coding in the other paradigms may even be unconscious and thus not even a “strategy.” In the case of the rotating line, it may be that noticing the line in a canonical place in its rotation (such as when it was either vertical or horizontal) easily allowed people to compute a “rhythm” for the rotating line. The rotating frustra were oriented so that the vertical axis of the cone was always horizontal in the rotation. This meant that the small circle created by cutting off the top of the cone formed a circle that was likely to be quite salient, perhaps either at the point when it reappeared after being occluded or when it was directly facing the viewer. It is somewhat less transparent what such a canonical place in the walking man’s gait would be—perhaps when one or the other foot touched the ground—but it doesn’t seem that far-fetched that there was one for each participant. However, whether conscious or unconscious, such coding may require focused attention and thus not be that different from detection of change in static properties. We are currently attempting to determine the role of attention in these tasks.

Although further research is needed to determine whether the detection of intrasaccadic change in a moving stimulus is possible without the aid of attention or a conscious strategy, our research shows that the assumption that there is naturally good detection of change for all moving stimuli is false. We think the question is still open as to the type of coding that allows good performance in some of these motion-detection tasks.

REFERENCES

- ALTARRIBA, J., KAMBE, G., POLLATSEK, A., & RAYNER, K. (2001). Semantic codes are not used in integrating information across eye fixations in reading: Evidence from fluent Spanish–English bilinguals. *Perception & Psychophysics*, **63**, 875–890.
- BRIDGEMAN, B., HENDRY, D., & STARK, L. (1975). Failure to detect displacement of the visual world during saccadic eye movements. *Vision Research*, **15**, 719–722.
- DITCHBURN, R. W. (1955). Eye-movements in relation to retinal action. *Optica Acta*, **1**, 171–176.
- FELDMAN, J. A. (1985). Four frames suffice: A provisional model of vision and space. *Behavioral & Brain Sciences*, **8**, 265–313.
- GYSEN, V., DE GRAEF, P., & VERFAILLIE, K. (2002). Detection of intrasaccadic displacements and depth rotations of moving objects. *Vision Research*, **42**, 379–391.
- GYSEN, V., VERFAILLIE, K., & DE GRAEF, P. (2002a). The effect of stimulus blanking on the detection of intrasaccadic displacements of translating objects. *Vision Research*, **42**, 2021–2030.
- GYSEN, V., VERFAILLIE, K., & DE GRAEF, P. (2002b). Transsaccadic perception of translating objects: Effects of landmark objects and visual field position. *Vision Research*, **42**, 1785–1796.
- HOLLINGWORTH, A., & HENDERSON, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception & Performance*, **28**, 113–136.
- IRWIN, D. E. (1992). Memory for position and identity across eye movements. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 307–317.
- JONIDES, J., IRWIN, D. E., & YANTIS, S. (1983). Failure to integrate information from successive fixations. *Science*, **222**, 188.
- MACK, A. (1970). An investigation of the relationship between eye and retinal image movement in the perception of movement. *Perception & Psychophysics*, **8**, 291–298.
- MCCONKIE, G. W., & RAYNER, K. (1976). Identifying the span of the effective stimulus in reading: Literature review and theories of reading. In H. Singer & R. B. Ruddell (Eds.), *Theoretical models and processes of reading* (pp. 137–162). Newark, DE: International Reading Association.
- MCCONKIE, G. W., & ZOLA, D. (1979). Is visual information integrated across successive fixations in reading? *Perception & Psychophysics*, **25**, 221–224.
- POLLATSEK, A., LESCH, M., MORRIS, R. K., & RAYNER, K. (1992). Phonological codes are used in integrating information across saccades in word identification and reading. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 148–162.
- POLLATSEK, A., & RAYNER, K. (1992). What is integrated across fixations? In K. Rayner (Ed.), *Eye movements and visual cognition: Scene perception and reading*. New York: Springer, 166–191.
- POLLATSEK, A., & RAYNER, K. (2002). Simple rotary motion is integrated across fixations. *Perception & Psychophysics*, **64**, 1120–1129.
- POLLATSEK, A., RAYNER, K., & COLLINS, W. E. (1984). Integrating pictorial information across eye movements. *Journal of Experimental Psychology: General*, **113**, 426–442.
- RAYNER, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, **7**, 65–81.
- RAYNER, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, **124**, 372–422.
- RAYNER, K., BALOTA, D. A., & POLLATSEK, A. (1986). Against parafoveal semantic preprocessing during eye fixations in reading. *Canadian Journal of Psychology*, **40**, 473–483.
- RAYNER, K., MCCONKIE, G. W., & ZOLA, D. (1980). Integrating information across eye movements. *Cognitive Psychology*, **12**, 206–226.
- RAYNER, K., & POLLATSEK, A. (1983). Is visual information integrated across saccades? *Perception & Psychophysics*, **34**, 39–48.
- VERFAILLIE, K. (1997). Transsaccadic memory for the egocentric and allocentric position of a biological-motion walker. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 739–760.
- VERFAILLIE, K., & DE GRAEF, P. (2000). Transsaccadic memory for position and orientation of saccade source and target. *Journal of Experimental Psychology: Human Perception & Performance*, **26**, 1243–1259.
- VERFAILLIE, K., DE TROY, A., & VAN RENSBERGEN, J. (1994). Transsaccadic integration of biological motion. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 649–670.
- WALLACH, H., & LEWIS, C. (1966). The effect of abnormal displacement of the retinal image during eye movements. *Perception & Psychophysics*, **1**, 25–29.

NOTES

1. In fact, some early models proposed the existence of such a “visual buffer” that maintained essentially a detailed “pixel by pixel” representation of information between fixations (see, e.g., Feldman, 1985; McConkie & Rayner, 1976).

2. We think attention is likely also needed for preview benefit that results from integration of information across saccades. That is, even though readers are not conscious of the change, they are undoubtedly attending to the word to the right of fixation prior to landing on it.

3. The Method section of Gysen et al. (2002) suggests a much slower rotation speed (1.9°/sec), but this is because of a confusing description, and the rotation speed cited here is the correct one (K. Verfaillie, personal communication).

4. However, in current work in progress in the lab, detection performance in the rotating stick paradigm, using the same setup, refresh rate, and type of control as in the present experiments, was quite similar to that reported in Pollatsek and Rayner (2002).

(Manuscript received October 26, 2004;
revision accepted for publication March 8, 2005.)