

Trade-offs Between Gaze and Working Memory Use

Jason A. Droll and Mary M. Hayhoe
University of Rochester

Eye movements during natural tasks suggest that observers do not use working memory to capacity but instead use eye movements to acquire relevant information immediately before needed. Results here however, show that this strategy is sensitive to memory load and to observers' expectations about what information will be relevant. Depending upon the predictability of what object features would be needed in a brick sorting task, subjects spontaneously modulated the order in which they sampled and stored visual information using working memory more when the task was predictable and reverting to a just-in-time strategy when the task was unpredictable and the memory load was higher. This self organization was evidenced by subjects' sequence of eye movements and also their sorting decisions following missed feature changes. These results reveal that attentional selection, fixations, and use of working memory reflect a dynamic optimization with respect to a set of constraints, such as task predictability and memory load. They also reveal that change blindness depends critically on the local task context, by virtue of its influence on the information selected for storage in working memory.

Keywords: eye movements, working memory, change blindness, executive control

In the course of visually guided behavior, the brain must make multiple decisions about what visual information from the environment to sample. Visual scenes typically abound with complex information, but the brain is fundamentally limited in how much information can be encoded and stored. Gaze is directed to only one area at a time, visual attention may modulate the information selected during each of these fixations, and the capacity of visual working memory sets strict limits on the ability to store the attended information across successive fixations. Thus, the brain must constantly decide where to look, what to attend, and what to remember. The complexity of coordinating each of these processes is made especially clear when considering that they occur in the context of extended behavior, when different visual information is required at different stages of a task. The purpose of the present experiment was to see if we could control how observers choose to select visual information for encoding and storage in working memory by manipulating the degree to which they could predict what information would be relevant for their task.

Where to Look?

While making a peanut butter and jelly sandwich or making tea, fixations of the eye are predominantly directed to areas in the

scene that are relevant to the immediate demands of the task (Hayhoe, Shrivastava, Mruczek, & Pelz, 2003; Land, Mennie, & Rusted, 1999; Land, 2004; Land & Hayhoe, 2001). Although this observation is intuitively unsurprising, the implications are quite profound. Strategically directing gaze to areas that are relevant for a behavioral goal suggests both that observers have the ability to control their sampling of the visual environment and that this sampling reflects the operation of some mechanism that controls how the task is structured. The structure of the task includes not only high-level decisions, such as whether to spread the peanut butter first or the jelly, but also the ordering of component microtasks, such as searching for the peanut butter jar, attending to the shape or color of the jar, guiding grasp, and recalling the location of the bread. How might these microtasks be related to fixation behavior? One possibility is that, as a consequence of limitations on visual working memory, fixations may be used to acquire information on a need-to-know basis just in time for the demands of the task. This approach is consistent with the idea that observers may use the world as a form of external memory, using shifts in gaze to acquire information as opposed to storing information internally (Ballard, Hayhoe, Pook, & Rao, 1997; O'Regan, 1992; O'Regan & Noë, 2001). This is evidenced by the observation that gaze is often redirected to the same location in the scene, presumably because the information needed was not encoded from an earlier fixation or previously encoded information was not stored in working memory.

What to Attend?

One method of assessing what visual information is being attended and encoded is to assume that observers are attending to the information targeted by the direction of gaze. However, it is difficult to draw strong inferences about what is being attended or other underlying cognitive operations solely from observations of fixation behavior. Although a cognitive operation may reliably lead to a particular fixation, the fixation itself does not uniquely

Jason A. Droll and Mary M. Hayhoe, Department of Brain and Cognitive Sciences & Center for Visual Science, University of Rochester

Jason A. Droll is now at the Department of Psychology, University of California, Santa Barbara. Mary M. Hayhoe is now at the Department of Psychology, University of Texas at Austin.

This work was supported by National Institutes of Health/Public Health Service Research Grants EY07125, EY05729, and RR09283. We thank Keith Parkins for programming the virtual environment and Constantin Rothkopf for development of the Fixation Finder software.

Correspondence concerning this article should be addressed to Jason A. Droll, Department of Psychology, University of California, Santa Barbara, Santa Barbara, CA 93106. E-mail: droll@psych.ucsb.edu

specify the ongoing cognitive operation. The challenge in ascribing cognitive processes to individual fixations is particularly acute when considering the distinct role of task demands. Although the direction of gaze is often interpreted as the locus of visual attention, the details of exactly what operations are occurring during each fixation must be more complicated than a “spatial spotlight” (Ballard et al., 1997; Cave & Bichot, 1999; Roelfsema, Lamme, & Spekreijse, 2000). How might these task-specific visual computations be used in ordinary behavior? Figure 1 illustrates the hierarchical organization of the operations when making a peanut butter and jelly sandwich. The cognitive goal of making the sandwich may require many subtasks, such as spreading the peanut butter, that in turn involve microtasks, such as grasping the peanut butter jar. To successfully complete this microtask, the observer fixates the peanut butter jar. During this fixation to the jar, what information is acquired? How specific is this information, and how does it relate to the demands of the microtask or to the larger cognitive goal?¹ If the fixation on the jar was guided on the basis of the brown color, what other irrelevant information, such as jar shape or texture, is visually processed during the fixation?

There is increasing evidence to suggest that different visual goals require different visual computations, appropriate for the current task (Ballard, Hayhoe, & Pelz, 1995; Droll, Hayhoe, Triesch, & Sullivan, 2005; Hayhoe, Bensinger, & Ballard, 1998; Triesch, Ballard, Hayhoe, & Sullivan, 2003; Ullman, 1984; Wallis & Bulthoff, 2000). Thus, fixations to objects at different stages of a task may result in different internal representations, depending on the immediate demands of the observer. For example, Ballard et al. (1995) had observers copy a model pattern of blocks on a computer screen by picking up blocks with the mouse and moving them to make a copy. In the course of copying a single block, subjects commonly fixated individual blocks in the model pattern twice, once before picking up a matching block and once before placement. Given the requirements of the task, a reasonable hy-

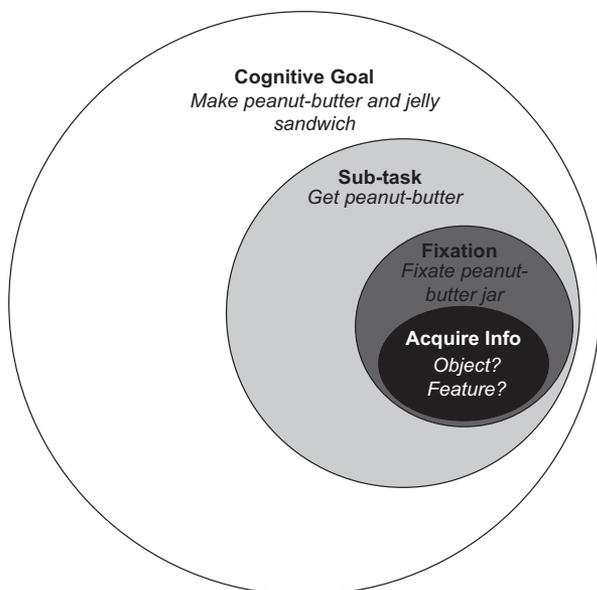


Figure 1. Cognitive goals requiring extended behavior may be organized as a hierarchy of operations.

pothesis is that block color is acquired during the first fixation and the next fixation on the block is to acquire its location.

The suggestion here is that observers may acquire different information from the point of fixation depending on their immediate goal or task. This idea is consistent with recent work in inattention blindness, demonstrating that the visual information acquired at the point of fixation is strongly modulated by task demands (Mack & Rock, 1996). Observers are often unaware of highly salient but task-irrelevant stimuli, even when it is present in the scene for several seconds and is at the point of fixation. However, observers are most sensitive to these unexpected stimuli when stimuli features (e.g., color) match those used in the current task set (Most, Simons, Scholl, Jimenez, Clifford, & Chabris, 2001; Simons & Chabris, 1999). This suggests that the attentional control setting defined by the task is a strong predictor of what information observers acquire from the scene. The view that visual representations are governed by task set differs from more traditional theories that posit general purpose representations, or object files, in which the role of attention is to bind object features and that attention is conceptualized as a mental spotlight yoked to direction of gaze (Kahneman, Treisman, & Gibbs, 1992; Noles, Scholl, & Mitroff, 2005; Treisman & Gelade, 1980). Psychophysical demonstrations of task-specific representations are also consistent with evidence on how neural mechanisms represent visual information. Neural activity is greater for behaviorally relevant information throughout the hierarchy of visual areas from the low-level primary visual cortex (Li, Piech, & Gilbert, 2004) to higher level areas involved in visuomotor transformations, such as the lateral intraparietal area (Gottlieb, Kusunoki, & Goldberg, 1998) or the dorsolateral prefrontal cortex (Asaad, Rainer, & Miller, 1998; Boettiger & D’Esposito, 2005; Miller & Cohen, 2001). However, although there are many experiments that examine task effects for individual operations, such as the planning of saccades, visual search, feature- or object-based attention, little is known about how these operations might be composed into extended tasks or how conclusions drawn from standard testing paradigms extend to an understanding of visual processes during natural behavior.

What is Stored?

Natural behavior is remarkably different from what might be predicted from many standard vision experiments, most especially with respect to the usage of working memory. Although estimates of working memory capacity suggest storage of three to seven objects (Cowan, 2000; Luck & Vogel, 1997; Vogel, Woodman, & Luck, 2001), performance during natural tasks often suggests that observers structure their behavior in a manner that minimizes the use of working memory. For example, as discussed above, relevant objects are often refixated many times throughout the course of a task (Ballard et al., 1995). These frequent refixations suggest that the subject chose not to store that information from previous fixations, despite its being well within the estimated capacity of visual working memory.

¹ Other researchers have suggested alternate ways in which to characterize the planning required in extended tasks (Land et al., 1999; Schwartz, 1995). Our present intention is simply to acknowledge that a hierarchy of goals and tasks is necessary for composing natural behavior.

However, inferring the contents of working memory through observation of eye movements is a somewhat indirect measure of what information was selected for storage. A more direct method of testing working memory is to measure subjects' sensitivity to visual changes. Insensitivity to visual changes, when these changes are masked by a visual transient, is known as change blindness (Simons, 2000). Failure to quickly find the change is often interpreted as evidence for sparse internal visual representations (O'Regan, 1992; O'Regan & Noë, 2001; Rensink, 2000) and is generally consistent with the just-in-time account of how vision is typically used (Ballard et al., 1997). However, it is not clear how demonstrations of change blindness reflect operations performed under ordinary circumstances, because subjects are often encouraged to adopt a relatively arbitrary decision criterion of where to direct gaze or of what information to remember. Relatively few change detection experiments monitor gaze, and thus they are unable to determine if observers even attended to the changed object. Thus, rates of change blindness are hard to evaluate because they are often not compared against evidence that observers fixated, attended, or stored the changed information in the first place (Henderson & Hollingworth, 1999; Hollingworth & Henderson, 2002). In fact, observers are much more likely to detect visual changes for objects that they have fixated both before and after the change (Henderson & Hollingworth, 1999). However, the reasons changes might be missed in these cases are not clear. One possibility is that despite directing gaze to the object, subjects did not encode or store the features that underwent the change, perhaps because the relevance of this information in relation to the task was not known.

In contrast to more standard change-detection paradigms, in which the subject's only task is to search for a change between alternating images, measuring sensitivity to visual changes in the midst of an ongoing task may more accurately assess how visual memory is used in natural behavior (Droll et al., 2005; Hayhoe et al., 1998; Shinoda, Hayhoe, & Shrivastava, 2001; Simons & Levin, 1998; Triesch et al., 2003; Wallis & Bulthoff, 2000). The results from these experiments suggest that the demands of the task determine what visual information is selected for storage.

Trade-Offs Between Gaze and Working Memory Use

Simply describing the information required in a task is unlikely to sufficiently characterize the visual representations used in extended behavior. This is because a task requiring extended behavior can be organized in a variety of ways. For example, when fixating an object, the observer may either decide to store this information in working memory or not to store this information and simply refixate the object the moment the information becomes necessary. The decision to use one or the other strategy may depend on the cost of each. These costs may reflect the burden placed on each resource, either in metabolic cost for storing information (Haxby, Petit, Ungerleider, & Courtney, 2000), redirecting attention (Serences, Schwarzbach, Courtney, Golay, & Yantis, 2004; Yantis et al., 2002), or perceptual or attentional load (Lavie, Hirst, de Fockert, & Viding, 2004). Thus, in the course of performing a task, observers may need to balance the use of each strategy to satisfy a set of such constraints (Gray, Schoelles, & Sims, 2005). Other related work has investigated this trade-off in relation to eye movements. For example, in Ballard et al.'s (1995)

block copying task, subjects often fixated a block twice while they were copying it, once before picking up a brick of that color and then again just before putting it down. In one condition, the experimenters increased the cost of the eye movements by placing the regions in the display farther away from each other. Subjects had to make a larger head movement to fixate the model blocks, and this may have added to the cost, either in energy or time, for refixating the blocks. Consequently, subjects reduced the frequency of fixations on the model bricks and were more likely to remember the fixated information throughout the next few stages of the task. This suggests that as the cost of a change in gaze increased, observers offset this cost by increasing their reliance on working memory. Throughout the course of a task, observers may find a balance point between using changes in gaze to acquire information in the scene or storing information internally. Thus, opposite to the manipulation performed in Ballard et al., if a task increased the demand on working memory, subjects might offset this cost by more frequently using changes in gaze to acquire the necessary visual information. One of the purposes of the present experiment was to perform such a manipulation; that is, to increase the working memory load in a task to see if observers would switch to a strategy that relied more heavily on the just-in-time gaze strategy.

In the block copying task, as described above, observers used a just-in-time fixation strategy in preference to using working memory. However, in a recent experiment, Droll et al. (2005) found instead that subjects used working memory in preference to a just-in-time strategy. In this experiment, subjects selected and then sorted bricks in a virtual environment. Subjects performed trials in which they knew which brick feature would be relevant for both the pick-up and put-down decisions (e.g., select and sort on the basis of width). During this task, subjects only rarely refixated the brick in hand after they had picked it up, but instead looked directly to the location where they were putting it down. Because subjects did not refixate the brick to get the information just in time, this suggests that they instead used information stored in working memory for the sorting decision. Why might subjects sometimes use memory and sometimes use a just-in-time fixation? One possibility was suggested by another condition in the experiment. In this condition, a second feature was used for put-down. Thus, the working memory load was increased from one to two features. This resulted in a small increase in the number of fixations and, by inference, a reduction in the use of working memory. One interpretation is that the modest increase in memory load required when a second feature was used for put-down caused subjects to increase the use of the just-in-time strategy.

In the present experiment, we explored the hypothesis of increased reliance on the just-in-time strategy by adding a further increase to the working memory load, in an attempt to better understand when subjects use working memory as opposed to overt fixations. The critical manipulation was to randomize which brick feature would be relevant for the put-down decision. Because subjects could not anticipate which feature would be relevant for sorting until later in the task, they either had to remember *all* brick features (four, in this case), or, they could opt for a just-in-time strategy and refixate the brick to get the relevant information just before put-down. We found that the modest increase in memory load from two to four features had a dramatic influence on eye movement patterns and the storage of visual information in work-

ing memory. Subjects were much more likely to revert to a just-in-time strategy when the memory load required by the task was higher. Because the selection to use each strategy was spontaneous, this suggests that such trade-offs are an intrinsic component of natural visually guided tasks. Thus, the use of eye movements, attentional selection, and working memory reflect a dynamic optimization with respect to a set of constraints, such as metabolic energy, time, or memory load.

Method

Equipment

Subjects viewed a virtual scene through a Virtual Research (Aptos, CA) V8 head mounted display. The visual display was generated by a Silicon Graphics computer at a rate of 60 Hz and was rendered in stereo on two LCD screens in the headset with 640 by 480 pixel resolution. The viewpoint updating had 1–2 frame latency with respect to the position and orientation of the head, tracked at 120 Hz with a 6-ms internal latency (Fastrack; Polhemus, Colchester, VT). Two devices monitored the movements of the eyes. An Applied Science Laboratory (ASL) 501 video-based eye tracker monitored the position of the left eye with 60 Hz temporal resolution and approximately 1° in accuracy. A limbus eye monitor (ASL 210) monitored the velocity of the right eye with 1000 Hz resolution and was used to detect saccades or blinks in order to trigger display changes. Force feedback from interaction with objects in the virtual environment was provided with two haptic stimulation devices that allowed subjects to experience realistic forces when grasping objects. Two Phantom-3 devices from SensAble Technologies (Woburn, MA) were used in opposition—one for the thumb and one for the index finger. The usable work-space volume was 55 cm × 55 cm × 40 cm. For more details on the equipment, see Droll et al. (2005).

Sorting Task

Figure 2 shows the scene visible to the observer within the helmet. Trials began with an array of five virtual bricks on a table (Figure 2A). Each brick was defined by four features: color, width, height and texture. Each feature was in one of two different states: (a) brick color was either red or blue; (b) width was either wide or thin (7.7 or 6.0 cm); (c) height was either tall or short (10.0 or 8.0 cm); and (d) texture was either thick or thin diagonal stripes (0.5 cycles/cm or 0.25 cycles/cm). At least one instance of each feature state appeared within the array of bricks. The small red spheres in the display represent thumb and index finger position.

The example shown in Figure 2 is taken from a trial in which color was relevant for both pick-up and put-down. Above the bricks, a pick-up cue indicated which feature state the subject should select (Figure 2A; videos of subject performance can be viewed online at <http://www.psych.ucsb.edu/%7Edroll/BricksParadigm/index.html>). The yellow arrows indicate that the subject should pick up a red brick. The feature value relevant for pick-up was varied pseudo-randomly from trial to trial, with the constraint that each value of the feature occurred five times in a block of ten trials. Subjects were instructed to select any brick with the relevant feature value regardless of the value of the other three features. After lifting up the brick, the pick-up cue disappeared

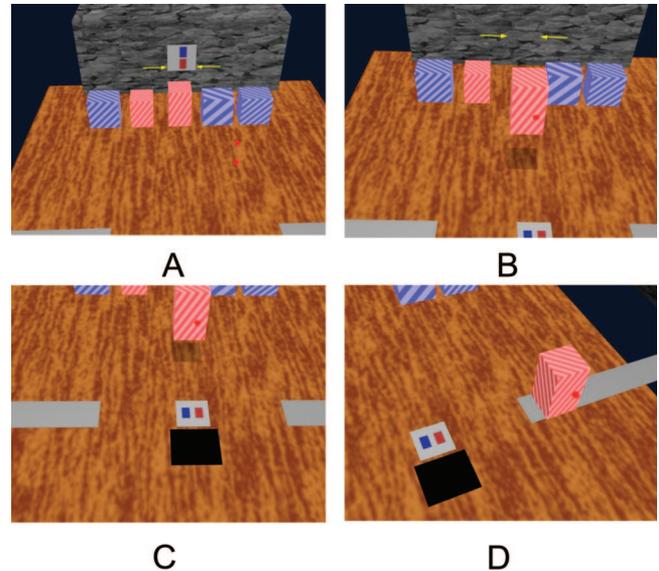


Figure 2. Scene during a *One Feature* trial when brick color was task relevant. Fingertips are represented as small red spheres. In a single trial, a subject (A) selects a brick based on the pick-up cue, (B) lifts the brick, (C) brings it toward himself, (D) decides on which conveyor belt the brick belongs based on a put-down cue.

from the scene (Figure 2B), and the brick was carried toward the subject (Figure 2C). Upon disappearance of the pick-up cue, the put-down cue was displayed between two conveyor belts (Figure 2B–D). The put-down cue showed the two values of the task-relevant feature, and their spatial arrangement indicated the appropriate conveyor belt for put-down. For example, in this trial, because the brick in hand is red, the cue in Figure 2D instructs the subject to place the brick on the right belt. (See Figure 3 for cues used for other features.) Similar to the pick-up cue, the put-down cue (e.g., “red on right/blue on left” or “red on left/blue on right”) was determined pseudorandomly across trials. After the brick was placed on the conveyor belt (Figure 2D), it was carried out of the scene and the put-down cue was removed, indicating the end of a trial. The next trial began immediately, presenting five new bricks and a new pick-up cue. Throughout the entire experiment, no time constraint was imposed on the participants. In an effort to capture movements representative of ordinary behavior, subjects were encouraged to perform the task at a pace they considered comfortable and natural.

Figure 3 shows the two trial types and experimental conditions. In the *One Feature* trials, only one of the four features was relevant for both pick-up and put-down (e.g., pick-up by color, put-down by color). In the *Two Feature* trials, one feature was used for pick-up, and a second feature was used for put-down (e.g., pick-up by color, put-down by width). In the *Predictable* condition (Figure 3A), subjects performed a block of 80 *One Feature* trials followed by a block of 80 *Two Feature* trials. These trials were predictable in the sense that the same features were always relevant at each stage of the sorting task (e.g., always pick-up by color, always put-down by either color or width). Because of the blocked trial design, subjects knew which brick feature was relevant for both pick-up and put-down, regardless of trial type. Seventy-two subjects participated in the *Predictable* condition and were counter-

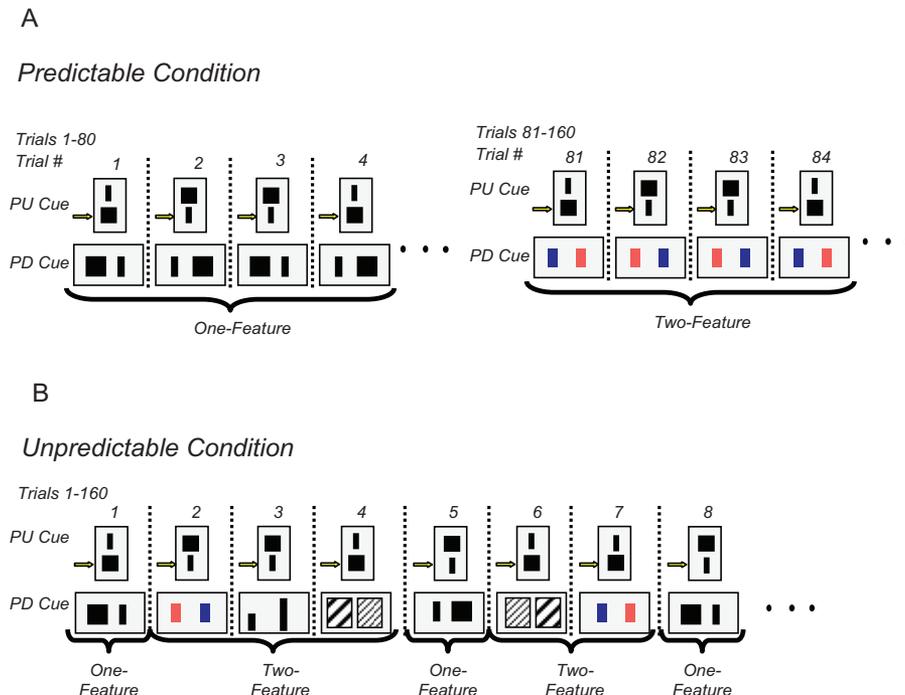


Figure 3. Sample pick-up and put-down cues for each trial type in (A) *Predictable* and (B) *Unpredictable* conditions. In the *Predictable* condition, the same one or two features were relevant during all trials. In the *Unpredictable* condition, subjects always used the same feature for pick-up, but could not predict which brick feature would be relevant for put-down. *One Feature* trials used the same feature for both pick-up and put-down; *Two Feature* trials used separate features for pick-up and put-down. PU = pick-up; PD = put-down.

balanced across the four brick features used for pick-up and put-down. Data from the *Predictable* condition have been reported previously (Droll et al., 2005) and are presented again here for the purpose of comparison.

The new manipulation in this experiment was the *Unpredictable* condition (Figure 3B). In the *Unpredictable* condition, subjects performed a block of 160 trials in which one brick feature was always used for pick-up, but the feature used for put-down was randomly assigned. These trials were unpredictable in the sense that subjects could not accurately predict which brick feature would be relevant for put-down on each trial (e.g., always pick-up by color, but put-down may be by color, width, height, or texture). A separate group of 32 subjects participated in the *Unpredictable* condition and were counterbalanced across the four brick features used for pick-up. Each of the eight put-down cues was randomly assigned in each block of 16 trials.

Change Detection Task

Subjects performed a change detection task concurrently with the sorting task. Before the beginning of the experiment, subjects were told that any of the features of the brick they were carrying might change, regardless of the feature’s relevance in the sorting task. If subjects detected a feature change they were instructed to place it into a virtual trash bin, a black hole located between the conveyor belts (Figure 2C–D). The change detection task was designed to probe the current information in working memory used in sorting. In order to minimize the influence of this secondary

change detection task, changes occurred in fewer than 10% of trials.

Before data collection, subjects performed 16 practice trials to familiarize themselves with the task. To ensure that subjects were familiar with the trash bin and comfortable with the instructions, yet naïve at detecting actual changes, at some time between the 10th and 16th practice trial the experimenter interrupted the subjects and instructed them to “Imagine that the brick you are carrying suddenly changed in either color, width, height, or texture. Where would you place the brick?” Subjects always placed the brick in the trash bin, indicating that they understood the change-detection task. Following recalibration of the eye tracker, subjects were reminded a third time to place any brick that changed in any of the four features into the trash bin.

Within each block of 80 trials, up to two changes could occur for each of the four brick features. The order of the feature changes was randomized. Feature changes occurred only during a saccade and only when the brick in hand was within a change zone extending from the back edge of the conveyor belts to the half-way point between the conveyor belts and the brick array. Thus, changes occurred between pick-up and put-down, and any transient signal was masked by the saccade. Ninety-three percent of visual changes in this system are generated in response to a saccade or a blink and changes are displayed with a mean latency of 50 ms following the initiation of a saccade. For more detail on saccade contingent updating in virtual reality (see Triesch, Sullivan, Hayhoe, & Ballard, 2002).

Because the number of trials in which a change to a brick feature was variable (as a consequence of saccade-triggered changes), there was a small number of subjects who were not exposed to a change of a particular feature. Thus, although there were 72 and 32 subjects in the *Predictable* and *Unpredictable* conditions, respectively, it was not possible to use all subjects in each condition when comparing rates of change detection. The reported degrees of freedom reflect the number of subjects used in each analysis.

Eye Movement Analysis

Fixations were determined using in-house Fixation Finder software, which implements three algorithms incorporating eye velocity and position: (a) a velocity-based algorithm, (b) an adaptive velocity-based algorithm that adapts the velocity threshold depending on an estimate of the noise level present in the signal for each subject, and (c) a Hidden Markov Model. The initial threshold for the eye movement velocity was set to 50 deg/s. This high threshold was used due to the noise present in the tracking signal. All recorded fixations needed to meet the additional criteria of having angular velocity less than 50 deg/s for at least 50 ms. Successive fixations occurring less than 50 ms apart and with a displacement of less than 1.5° were consolidated. The Fixation Finder then provided a confidence value associated with each fixation, depending on the agreement between the algorithms. This automated scoring of fixations was judged to be comparable to that of manual scoring. In house Matlab (Mathworks) functions were used to analyze eye movements during the experiment, including identifying what object each fixation fell on and the duration of the fixation.

Due to the dynamic nature of the virtual environment, an additional criterion was used to categorize fixations directed to the brick in hand. In a small number of trials, subjects would carry the brick into the line of gaze during a fixation. Consequently, although subjects did not initially target the brick in hand, the brick was nevertheless the centrally fixated object throughout most of the fixation. Thus, we categorized fixations to the brick in which at least two thirds of the frames within the fixation were directed to the brick, or a minimum of 15 frames, sampled at 60 Hz. This criterion was not applied to the data previously published in Droll et al. (2005), and thus there are minor quantitative differences presented here. The qualitative pattern of results is identical.

Eye position was monitored for all 104 subjects. Automated fixation analysis is sensitive to noise in the tracking signal so it was necessary to screen subjects based on frequency of track loss. Because of the difficulty in maintaining an accurate track within the virtual reality helmet, only a subset of subjects had adequate eye position data throughout the experiment to merit analysis with the automated fixation finder. Analysis of the video records resulted in the selection of 73 subjects for whom eye position was judged adequate for automated analysis (43 and 30 in the *Predictable* and *Unpredictable* experimental conditions, respectively). Based on the video records, the results reported for these 73 subjects are representative of all 104 subjects.

Experimental Sessions

A total of 104 subjects participated in the experiment for \$10 per hour: 72 in the *Predictable* condition (Droll et al., 2005); and 32 in the *Unpredictable* condition. Experimental sessions typically lasted about an hour. All experiments were tested in accordance to regulations required by the University of Rochester Human Subjects Review Board. Subjects were recruited through posters around campus and ranged in age from 18 to 39 years.

Results

Certain aspects of performance in the *Predictable* condition have been reported previously (Droll et al., 2005). We report them here again for purpose of comparison to the *Unpredictable* condition in order to highlight the different pattern of results found in the present experiment. Videos of subject performance can be viewed online (<http://www.psych.ucsb.edu/%7Edroll/BricksParadigm/index.html>).

Performance During Trials Without a Feature Change

Subjects performed the sorting task with near perfect accuracy in each condition. Improper brick selection or brick placement during put-down occurred in less than 1% of trials. The speed of their behavior appeared to be neither particularly rushed nor slowed. The general pattern of behavior in a single normal trial, with no feature change, began with a fixation to the pick-up cue and the selection of a brick from the array. As subjects brought the brick toward the conveyor belts, they fixated the put-down cue before deciding which belt to place the brick on. We were particularly interested in performance at this stage in the task (see Figure 4). If subjects looked directly to the appropriate belt immediately after the put-down cue (Figure 4A, left) and began to move the brick toward it, then it is likely that the decision for brick placement was made on the basis of their working memory representation of the brick feature. If, however, the subject refixated the brick in hand before looking at the belt to guide placement (Figure 4A, right), it is likely that the decision was made on the basis of feature information acquired during the refixation, using a just-in-time strategy. To investigate which strategy subjects were using, we examined the frequency with which subjects refixated the brick in hand before the brick was carried to the belt it would ultimately be placed on.

In the *Predictable* condition, shown in Figure 4A, following a fixation to the put-down cue, subjects overwhelmingly looked directly to the belt to guide put-down rather than refixating the brick in hand before sorting, as previously reported in Droll et al. (2005). This suggests that the sorting decision was based on information stored in working memory, and presumably acquired at time of pick-up. Subjects may have stored the task-relevant features in working memory because they could predict which features would be relevant in every trial because of the blocked trial design. In comparison, in the *Unpredictable* condition, subjects were much more likely to refixate the brick in hand before fixating the conveyor belt to guide placement (Figure 4B). In this case subjects could not predict which of the four features would be used for put-down and would have to store four brick features if

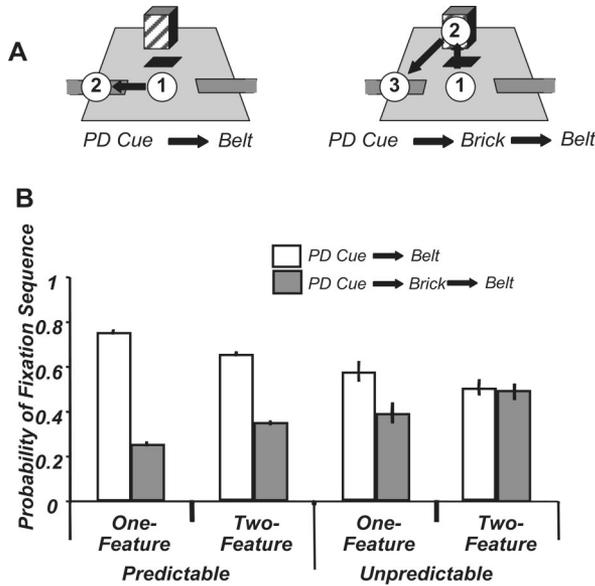


Figure 4. Two fixation sequences during put-down. (A) After fixating the put-down cue, subjects either fixated the belt (left icon) or fixated the brick before fixating the belt (right icon) before sorting. (B) Probability of each fixation sequence for each trial type when the put-down feature was either *Predictable* or *Unpredictable*. Subjects were more likely to refixate the brick when they couldn't predict which feature was used for put-down. The probability of refixating the brick was also greater when a second feature was used for put-down. PU = pick-up; PD = put-down.

they used a strategy that relied on working memory. Instead, subjects seemed to use a strategy that lessened their reliance on working memory. Consistent with this interpretation, the addition of a single extra feature in the *Two Feature* trials increased the probability of a refixation even in the *Predictable* condition (25 vs. 34% for one versus two features, paired *t* test, $p < .01$, $df = 42$). For more detail on performance in the *Predictable* condition, see Droll et al., (2005).

Greater frequency of brick fixations in the *Unpredictable* condition compared with the *Predictable* condition is especially interesting for the *One Feature* trials. In both instances subjects sorted the brick using the same feature they had used for pick-up, so this information must have been encoded in both cases. However, the frequency of refixation increased from 25% in the *Predictable* condition to 40% in the *Unpredictable* condition (unpaired *t* test, $p < .01$, $df = 71$). This suggests that when subjects are not certain if the pick-up feature will continue to be relevant for put-down, they may be less likely to store this information in working memory. More frequent brick refixations in the *Unpredictable Two Feature* than in the *Predictable Two Feature* trials (48% vs. 34%, unpaired *t* test, $p < .01$, $df = 71$) suggests that when memory load increases to four features of the same object, subjects offset this increase in load by using a refixation strategy, acquiring the information just in time for put-down once the demands of the sorting task are known.

Performance During Change Trials

Change detection. Within each block of 80 trials, up to two changes could occur for each of the four brick features. Because

the triggering of these changes also required a saccade or a blink during the time the brick was being carried, there was slight variation in the number of visual changes that occurred for each subject. Feature changes were relatively infrequent to ensure that subjects would not prioritize the secondary task of detecting changes. As changes overwhelmingly occurred soon after pick-up but before a sorting decision had been committed, noticed changes were almost always detected after the put-down cue had been inspected.

In the *Predictable* condition shown in Figure 5A–B, subjects detected changes to relevant features nearly twice as often as irrelevant features in both *One* and *Two Feature* trials, as reported previously in Droll et al. 2005; $F(1, 66) = 18.37$, $p < .001$, for *One Feature* trials; $F(2, 124) = 7.00$, $p = .01$ for *Two Feature* trials (Figure 5B). This pattern of results suggests that subjects selectively acquire and store individual object features relevant to the task rather than information on the entire object. (For more detail, see Droll et al., 2005). During *Unpredictable One Feature* trials, changes to the relevant feature were also detected more frequently (57%) than changes to irrelevant features (37%) (paired *t* test, $p < .05$, $df = 24$, Figure 5C). This suggests that once a feature is used for pick-up, subjects have some residual representation of this feature in working memory. Once the put-down cue prompts them to reevaluate the state of this feature, subjects may be able to detect the change by using this residual representation. However, observations of subjects' fixation sequence, as discussed above, suggest that this storage did not necessarily occur on all trials. In the *Unpredictable Two Feature* trials, rates of change detection followed no clear effect of task relevance (Figure 5D). Changes were detected with similar frequency regardless of whether the feature was relevant to pick-up (39%), relevant to put-down (38%), or irrelevant (37%), $F(2, 62) = 0.01$, $p = .98$. Detecting changes irrespective of task relevance is presumably a consequence of subjects' inability to accurately predict which object features to store in working memory. Because subjects could not predict which feature would be relevant for put-down prior to fixating the put-down cue, the relevance of the put-down feature was usually not known until after the feature change. Before the feature change, during pick-up, all features are potentially relevant (or irrelevant) to put-down, so there is no basis for preferential encoding of the remaining features.

It is interesting to note that rates of change detection for the pick-up relevant feature are higher in *Unpredictable One Feature* trials than in *Two Feature* trials (paired *t* test, $p < .01$, $df = 24$). One possibility is that, despite the highest rates of refixation in the *Unpredictable Two Feature* trials, these refixations were not for the purpose of reevaluating the changed pick-up feature, but instead for the purpose of evaluating the put-down feature. Rates of change detection were slightly higher overall in the *Unpredictable* than the *Predictable* condition (41% vs 35%), although this difference failed to reach significance (*t* test, $p = .22$, $df = 100$). However, it is also interesting to note that rate of change detection for the pick-up relevant feature is higher in the *Unpredictable* condition than in the *Predictable* condition (57% vs. 37%), although this difference only approached significance (*t* test, $p = .07$). What might account for this difference? In the *Predictable* condition, subjects are certain that the pick-up feature will be used again for put-down. As a consequence, subjects only infrequently refixate the brick during their sorting decision, and fixations to the

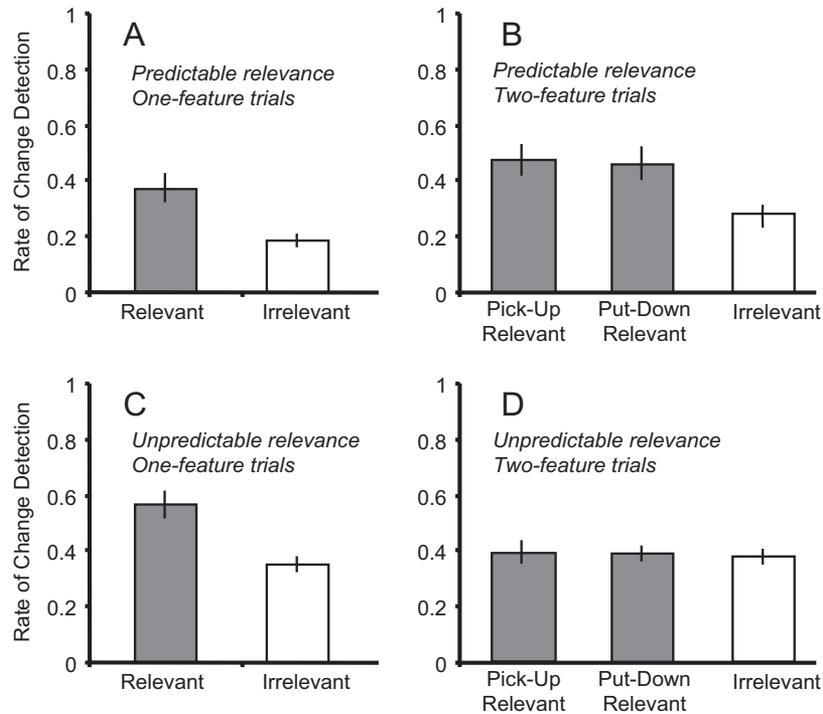


Figure 5. Rates of change detection with respect to task relevance for both the (A, B) *Predictable* and (C, D) *Unpredictable* experiments during (A, C) *One-Feature* and (B, D) *Two-Feature* trials.

brick after pick-up typically occurred during placement of the brick on the belt. It is possible that this task of brick placement consumes attentional resources and inhibits inspection of the new, postchange feature and detection of the change. This interpretation is consistent with subjects' sorting behavior following missed changes. In contrast, in the *Unpredictable* condition, following pick-up, subjects were uncertain if this same feature would be used for put-down and were more likely to refixate the brick. This reinspection operation, combined with storage of the prechange pick-up feature, may have resulted in the higher rates of change detection for *One Feature* trials in the *Unpredictable* compared with the *Predictable* condition.

Sorting behavior on missed trials. In trials when subjects failed to detect the feature change, sorting behavior provides a useful insight into the nature of the representation of the changed feature. In the *Predictable* condition, bricks with changes relevant to put-down were most often sorted on the basis of their old, prechange feature (see Figure 6), as reported previously in Droll et al. (2005). For example, if a wide brick changed to a thin brick while being carried, the brick was most often sorted as if it were still wide. This suggests that in the *Predictable* condition (both *One* and *Two Feature* trials), subjects were using their working memory for the put-down decision and not the feature value that was present within the scene at the time of put-down. This failure to reacquire or to update the changed feature occurred even though subjects later fixated the changed brick for an average of 750 ms as they guided the bricks onto the conveyor belt following the feature change. During *Unpredictable One Feature* trials, subjects were also more likely to sort the changed brick by the old feature. This suggests storage of the old, prechange feature, perhaps be-

cause it had been used as the basis for selection during pick-up. However, subjects were more likely to sort changed bricks by the new feature than the old feature during *Unpredictable One Feature* trials (5 of 16) than they were during these same trials in the *Predictable* condition (6 of 78; binomial, $p < .05$). Although there was only a small number of trials, this suggests that when subjects are uncertain if the pick-up relevant feature will continue to be relevant for put-down, this information may be less likely to be stored in working memory. This result parallels the increased frequency of refixations in the *One Feature* trials in each condition (Figure 4). Note that in the *Predictable* trials, subjects had a lower frequency of refixations, consistent with use of working memory.

Sorting behavior was markedly different in the *Unpredictable Two Feature* trials. During these trials, changed bricks were most often sorted by the new, postchange feature. Sorting by the new feature suggests that subjects acquired this feature value just in time for sorting, following the change. This interpretation is consistent with the higher frequency of refixating the brick in hand during normal trials, after fixating the put-down cue. Preference to sort changed bricks by the new feature is strong evidence for just-in-time acquisition of visual information, using the world as a form of external memory (O'Regan, 1992). Again, these trials are the same as those that were performed by subjects in the *Predictable* condition, the only difference being subjects' inability to accurately predict which brick feature would be relevant for put-down. Subjects could have adopted a strategy of storing all four brick features on each trial (c.f. Luck & Vogel, 1997) and used this stored information in working memory as the basis of their sorting decision. Instead, subjects seemed to have adopted a strategy of using

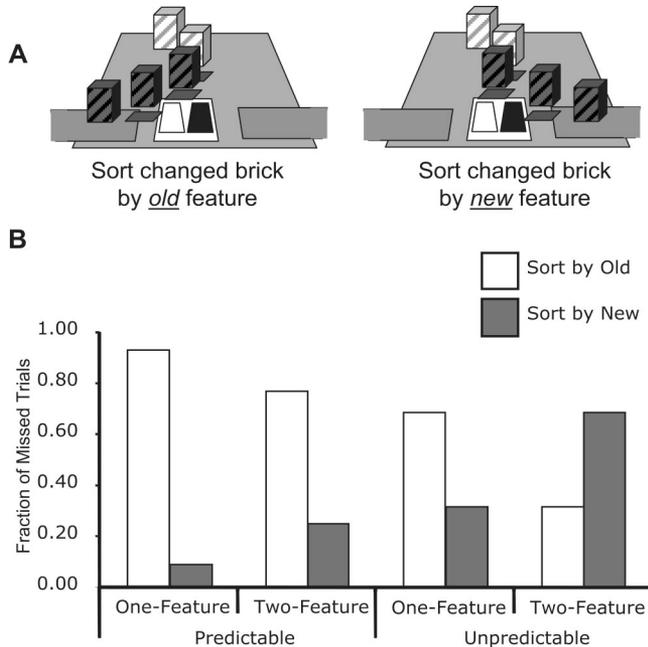


Figure 6. (A) Following a missed change, subjects either sorted the brick by the old, prechange feature (left) or the new, postchange feature (right). (B) Frequency of each sorting decision on trials in which the subject missed a change to the put-down relevant feature. When the changed feature was either predicted to be relevant for put-down or had also been used for pick-up, subjects sorted changed bricks by the old feature. However, when the change was to a feature that the subject had not used for pick-up and they could not predict its relevance for put-down, the brick was sorted by the new feature.

gaze to acquire the relevant brick feature just in time for sorting (Figure 4). Because this feature acquisition typically occurred after the feature change, missed changes were more frequently sorted by the new feature (Figure 6)².

Note also that sorting by the new feature in the *Predictable One Feature* task occurred in only 8% of missed change trials (6 of 78), whereas subjects sorted on the basis of the new feature in 24% of trials in the *Predictable Two Feature* task (16 of 66; binomial test, $p < .05$). This means that the sorting decision in the *Two Feature* task was more likely to be based on information acquired after the change, just before the time of put-down. This is consistent with normal trials in which subjects were more likely to refixate the brick in hand before guiding it to the belt in the *Predictable Two Feature* than in the *One Feature* trials.

Discussion

Subjects performed a sorting task in which they were required to select a brick for pick-up on the basis of one feature and then to sort the selected brick by the same or a second feature. The memory load of the sorting task was modestly increased by manipulating both the number of features used in each trial (one or two) and the subjects' ability to predict which of the four brick features would be used for sorting. *Predictable One Feature* trials required the storage of only one feature; *Unpredictable Two Feature* trials required the storage of four brick features. As the

memory load increased, subjects were more likely to shift from a strategy using visual working memory to a strategy in which visual information was acquired just in time through a change in gaze. This change in strategy was evidenced both in subjects' fixation behavior and also their sorting decision for missed changes.

Task Specific Representations

Implicit in the interpretation of our results is the idea that different features are encoded differentially. Subjects' fixation behavior, rates of change detection, and sorting performance are all consistent with the idea that only partial information about the brick may be encoded during pick-up. Thus, the way the brick is represented varies from trial to trial and with respect to the different task conditions. How do task demands influence the way in which visual information is represented in the brain? The present results support a growing body of evidence that suggests that visual information in the scene is represented differently when it relates to the task or when relevant features are attended (Corbetta & Shulman, 2002; Li et al., 2004; Roelfsema, Lamme, & Spekreijse, 2000). For example, cortical areas whose neurons are selective for a particular feature have enhanced activity when this feature is attended to (Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991; Desimone & Duncan, 1995; McAdams & Maunsell, 2000). When monkeys are rewarded for performing a particular action in response to a particular visual stimulus (e.g., look left when shown stimulus 'A'), activity in the prefrontal cortex is modulated by the learned contingency between the stimulus and the rewarded action (Asaad, Rainer, & Miller, 1998; Miller & Cohen, 2001). This body of work suggests that visual representations are not passively acquired for general-purpose use but are instead actively coded with respect to their behavioral relevance (Corbetta & Shulman, 2002; Curtis & D'Esposito, 2003; Fox, Snyder, Vincent, Corbetta, & Van Essen, 2005). Such specific encoding may underlie the change detection results in which task relevant changes were noticed more frequently than irrelevant changes. This selective encoding is consistent with inattentional blindness, in which observers are seemingly unaware of highly salient but irrelevant visual information (Mack & Rock, 1996; Marois, Yi, & Chun, 2004; Most, Scholl, Clifford, & Simons, 2005; Most et al., 2001; Simons & Chabris, 1999).

The present experimental results suggest that representations of objects may be incomplete, with not all of the features represented. It has traditionally been thought that one of the roles of visual attention is to bind object features into a single representation, or object file, that is then stored in working memory (Kahneman & Treisman, 1984; Kahneman et al., 1992; Luck & Vogel, 1997; Noles et al., 2005). However, characterizations of object files typically come from tasks that may not be representative of how

² Although there is general agreement between subjects' fixation sequence and their sorting behavior on missed trials, the correspondence is not perfect. One reason for differences is that data on the subjects' fixation sequence and their sorting behavior are taken from different trials. Interpreting refixations during change trials is problematic (it is not clear if those refixations are a cause or a consequence of detecting the change), and thus only normal trials, with no feature change, were analyzed for fixation sequence. Sorting behavior on missed trials, of course, was only analyzed for trials in which a change occurred.

vision is typically used, for example, identifying letters of the alphabet within moving frames (Kahneman et al., 1992). For the purpose of many memory tasks, integrated object files may indeed be required and, thus, formed when needed. This may not, however, be an invariant feature of working memory and there is growing evidence of task-specific representations (Gordon & Irwin, 2000; Wheeler & Treisman, 2002). Unselective feature binding would not predict the differential effects of the various conditions on fixation patterns, change detection performance, and sorting behavior found in the present experiment. Subjects often refixated the brick in hand when more features were needed for the task, sorted bricks on the basis of either old or new features in a manner that paralleled the frequency of refixations, and missed changes to this object of central interest, especially when the feature change was not relevant to their sorting decision.

The pattern of our results suggests that not all fixated information should be presumed to be acquired or stored in working memory as if in a passive buffer (Hollingworth & Henderson, 2002) nor should the updating of object representations be considered obligatory (Treisman & Gelade, 1980). However, neither is working memory so sparse that observers exclusively rely on using the environment as a form of external memory (O'Regan, 1992). When subjects could predict the relevance of each brick feature, the relevant features were often stored in memory, as revealed by subjects' fixation sequence during normal trials and their sorting decision following a missed change (*Predictable* condition). These results are consistent with the observation that the presence of a predictive cue will increase the likelihood of storing the cued stimulus in working memory (Schmidt, Vogel, Woodman, & Luck, 2002). In the present experiment, the predictive cue is not an explicit signal, but rather an internal expectation of what visual information is likely to be relevant to the task, learned across trials.

It should also be noted that in our interpretation of the results, we assume that selectivity occurs at the time of encoding of the perceptual features. It is also possible that all features are encoded nonselectively, but that the irrelevant features are not maintained in working memory. This possibility is consistent with evidence suggesting that processes of perceptual binding during encoding may be distinct from the binding mechanisms involved in memory storage (Allen, Baddeley, & Hitch, 2006; Mitchell, Johnson, Raye, & D'Esposito, 2000). In the present paradigm, we are unable to distinguish between these possibilities. Both accounts are consistent with the observed pattern of change detection and refixations. However, the predominance of the sort-by-old strategy during put-down reveals a remarkable insensitivity to the stimulus features and is consistent with the inattentional blindness literature. This result suggests that the limitation is indeed due to encoding rather than memory.

Another possible cause of subjects' insensitivity to encoding changed features is that the changes that occurred in this experiment occurred between two feature states that were generally familiar in the context of the scene (e.g., wide or thin, red or blue). As a consequence, insensitivity to these changes may have reflected proactive interference stemming from previous trials (Monsell, 1978). A related idea is that irrelevant changes were missed because, despite the strong low-level difference between episodic physical states, these changes did not sufficiently violate observers' expectations of what features to expect in the scene and thus did not attract attention. Such immunity to low-level changes may

perhaps allow objects to be represented with high-level continuity, despite large changes in physical detail (Gordon & Irwin, 2000).

Trade-Offs Between Gaze and Working Memory Use

The most striking aspect of our results is the ease with which we were able to control the subjects' use of either using gaze or working memory by increasing the memory load from one to four features. Such sensitivity suggests that such trade-offs are an intrinsic aspect of natural behavior. In Ballard et al.'s (1995) block copying task, the authors postulated that frequent fixations to the model area reflected a strategy that minimized memory load. In support of this, they found a reduced frequency of fixations on the model bricks when subjects had to make a larger head movement to fixate the model blocks, which may have added to the cost, either in energy or time, for refixating the blocks. In the present experiment, subjects made a similar trade-off. Fixations to the brick were more frequent when the task required a high memory load, but subjects preferred to use working memory when they knew in advance which one or two features were needed. Both these results suggest that the balance between eye movements and working memory reflects some kind of optimization or trade-off with respect to a set of constraints on the part of the observer. In the Ballard et al. (1995) experiment, this trade-off was influenced by the time or energy demands of head movements. In the present experiment, the trade-off appears to be influenced by the predictability of the necessary information, which changed the memory load.

The mechanisms by which observers settle on a particular strategy or balance between working memory use and just-in-time fixations is not clear. Note that subjects were instructed to perform the task at a pace they felt was comfortable and natural. The variability in performance between subjects is quite small, suggesting that this trade-off reflects a pervasive and stable aspect of natural behavior. Subjects also did not appear to be aware that they were making these trade-offs, or explicitly making decisions about their strategy. It is also unclear why subjects appear to opt for reduced memory strategies, well below traditional measures of working memory capacity, in the context of such natural tasks. One possibility is that there may be other demands on attention and working memory in natural tasks that are not obvious to the experimenter, such as controlling the grasp or remembering the location of the conveyor belts in order to program the eye and hand movements for put-down. Prioritizing each of the many microtasks required for each trial may have also been influenced by the expected value or execution cost of performing each operation (Fu & Gray, 2006; Gray et al., 2005). Had we provided extra incentive, in the form of reward, or different task instruction, subjects' may have reprioritized perceptual encoding or memory use, resulting in a different pattern of behavior. Another possibility is that organizing multiple subtasks requires the use of an executive control mechanism, demanding attentional resources that could otherwise be allocated toward working memory (Hester & Garavan, 2005). Another attentional demand may have been the covert acquisition of visual information in the periphery without the use of eye movements. Our interpretation of the results has implicitly assumed that perceptual encoding requires the use of overt fixations, but this may not necessarily be the case. Indeed, because sorting by the new, postchange, feature was more frequent than the rate of

refixations in the *Unpredictable Two Feature* trials, subjects may have occasionally acquired the new feature from peripheral information. However, it is not clear how subjects may have used peripheral information in each condition or trial type.

The trade-offs between fixations and working memory use observed in the present experiment serve as a caution against overgeneralizing the performance from many standard experimental paradigms. Traditional visual paradigms are often designed with the intention of identifying the properties of particular processes, such as working memory (Vogel et al., 2001). However, such experiments may reflect the limits on performance in the context of the particular experiment rather than the usage of these operations during natural tasks. In the present experiment, subjects were encouraged to perform in a manner in which they felt most comfortable. Thus, the decision of whether to store information in memory or to use gaze may reflect the operation of sets of constraints, and the trade-offs between these constraints may vary from task to task. Assessing where this balance point lies, by monitoring trade-offs between gaze and working memory use, may be a more accurate way in which to characterize memory *usage* rather than more traditional paradigms investigating working memory *capacity* (Vogel et al., 2001).

Change Detection

Poor performance in change detection experiments is often interpreted as evidence for sparse storage of internal visual representations (O'Regan, 1992; O'Regan & Noe, 2001; Rensink, 2000). Although there have been other suggestions as to what might cause change blindness, such as failing to encode the new information after the change or failing to compare old and new information (Simons, 2000; Simons & Rensink, 2005) the current experiment reveals that the particular cause will depend critically on the experimental context (Hayhoe, 2000; Hayhoe et al., 1998; Triesch et al., 2003). In the present experiment, we were able to manipulate sensitivity to changes by varying the task-relevance of the changed feature. We were also able to manipulate whether subjects used prechange or postchange information in the sorting task when they missed the change. This suggests that the information in memory is highly context sensitive, and this in turn will influence what changes are noticed.

Note that the strategy in the present experiment was to make changes infrequently in an effort to preserve the primary task goal of sorting bricks and to allow us to probe the information used to perform the task. This may account for the rather low rates of change detection, despite the fact that the brick was of "central interest" during the entire trial and subjects were told beforehand that changes would occur. Experiments where the primary goal is to detect changes will reflect a different set of constraints, such as subjects' prior expectations about changes. Sensitivity to changes under those laboratory conditions has also been used as evidence for detailed and robust long-term memory representations of objects present in natural scenes, where visual information accumulates in working memory across successive fixations (Henderson & Hollingworth, 1999; Hollingworth, 2004; Hollingworth, Williams, & Henderson, 2001). Is such robust memory representative of how we accumulate information during our everyday experience? The extreme task sensitivity of memory use revealed in the present experiment suggests that robust storage might not be

general. It is unlikely that observers typically attempt to store as much information as they are able during each fixation, as they are typically instructed to do in many memory experiments. Instead, long-term memory representation of scenes may reflect the history of goal-directed interactions by the subject.

Task Control

Vision is not simply the process by which the brain represents the visual environment. A central feature of vision is its versatility in performing complex tasks and assisting in myriad decisions, even in a task as simple as picking up and putting down an object. Performance in the present experiment suggests a model of vision in which visual operations (such as directing gaze, acquiring select visual information and storage in working memory) are controlled by learnt procedures that somehow orchestrate the sequence of visual operations required to complete a task. For example, in a single trial, subjects have the goal of sorting a brick. This goal may be organized as a sequence of subtasks, such as the pick-up or put-down operation for each brick, each lasting 2 to 3 seconds. These subtasks will require yet smaller tasks, such as using an eye movement to acquire an individual brick feature (Ballard et al., 1997). Feature states of the brick that are not immediately relevant may simply not be evaluated, and thus their change will escape detection. This way of thinking about vision suggests that the purpose of vision, and visual attention, is not to construct internal representations for general-purpose use. Instead, visual attention is best understood in the context of what information must be selected and evaluated for the purpose of guiding behavior (Allport, 1989). For example, the visual operations performed at time of pick-up may include not a single operation, in which all features are automatically acquired, but rather a process involving up to four discrete tasks: categorizing color, height, width, and texture into each of two possible feature states (e.g., red/blue, tall/short, etc.). These individual discriminations, or feature judgments, may relate to processes required in traditional dual-task paradigms in which subjects are shown similar stimuli in all trials but are asked to make multiple decisions (e.g., parity and/or magnitude judgments when shown numbers). Rather than performing multiple operations simultaneously, in parallel, it is often more efficient to perform them in a serial manner due to limits in attentional capacity (Logan & Gordon, 2001). This conceptualization of performance is similar to some formal models of executive control in which high-level decision processes affect lower level sensory selection (Bundesen, 1990; Logan, 2004; Logan & Gordon, 2001). In these models, the task defines the sensory parameters to be used, and this constrains what information is acquired (see also Figure 1). Visual information that is not immediately relevant or unlikely to be relevant is simply not evaluated, and thus never gains access to working memory. This general framework is also consistent with many theoretical ideas on how cognitive systems may be fundamentally organized (Ballard et al., 1997; Newell, 1990; Ullman, 1984) and can also be used to model complex behavior, such as navigating a cluttered sidewalk (Sprague, Ballard, & Robinson, in press).

Learning Task Control

The way in which tasks exert control over the acquisition of visual information must clearly be learnt (Hayhoe & Ballard,

2005). One contributing element to this learning may be the reward or costs that observers learn to associate with alternate courses of action. For example, recent work in the neurophysiological basis of eye movements have revealed that the saccadic eye movement circuitry is sensitive to the reward structure of the task (Hikosaka, Takikawa, & Kawagoe, 2000; Ikeda & Hikosaka, 2003; Platt & Glimcher, 1998; Sugrue, Corrado, & Newsome, 2004). In these experiments, monkeys are often free to choose between visually guided alternative behaviors (direct gaze to the target on the left or to the right). The monkey's behavior can be generalized as seeking a strategy in which the direction of gaze is chosen by virtue of the reward they have learned to expect to receive following their behavior, and monkeys are quite sophisticated at learning how to maximize their overall gain (e.g., drops of juice). Hayhoe and Ballard (2005) have argued that this sensitivity to reward may serve as a substrate for mediating the tight linkage between fixations and task structure in natural behavior. Similar to tasks in the reward literature, subjects in the present experiment were allowed to make self-directed visually guided decisions. Whereas monkey subjects may be maximizing their expected reward, human subjects during brick-sorting may be making their decisions based on the expected cost of each alternative (e.g., store visual information or refixate this object later). It is also interesting to note that sensitivity to reward is often observed in cortical areas traditionally associated as modulating attention and that distinguishing between the two may be difficult (Maunsell, 2004). Thus, it seems reasonable to consider that this reward circuitry may serve as a substrate to mediate the effects of task structure and visual attention seen in the present experiment.

Conclusion

The present experiment was designed to test the acquisition and storage of visual information throughout a naturalistic task. Our method allowed subjects to freely organize their behavior and cognitive operations in a manner in which they felt most comfortable. The pattern of eye movements suggested that subjects either stored visual information or delayed acquisition of this information depending on both the number of features used in the task and the degree to which this information could be expected to be relevant. We tested working memory more directly by making occasional feature changes to the object the subject was sorting. Rates of change detection and sorting performance during missed trials were consistent with the interpretation that observers have a selective visual representation of the brick, which is largely dictated by the immediate demands of the task. This representation is fleeting, and observers make trade-offs between storing either this information in working memory or acquiring it later with a redirecting gaze. How subjects settle on an operating point for this trade-off and the set of constraints that need to be satisfied are important questions for understanding how natural behavior emerges from elemental perceptual operations.

Rather than simply documenting yet another example of an egregious failure to detect a visual change, our intention in the present experiment was to exploit the phenomenon of change blindness as a probe to examine what visual information was being used in the course of natural behavior. Thus, we presented different visual information at different stages of a sorting task, allowing us to assess the time at which this information was acquired. By

allowing subjects freedom to guide themselves through an extended task, we can gain insight into a more plausible account as to how the brain decides where to look, what to attend, and what to remember.

References

- Allen, R. J., Baddeley, A. D., & Hitch, G. J. (2006). Is the binding of visual features in working memory resource-demanding? *Journal of Experimental Psychology: General*, *135*, 298–313.
- Allport, A. (1989). Visual attention. In M. Posner (Ed.), *Foundations of cognitive science*, (pp. 683–726). Cambridge, MA: MIT Press.
- Asaad, W. F., Rainer, G., & Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron*, *21*, 1399–1407.
- Ballard, D. H., Hayhoe, M., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*, 66–80.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, *20*, 723–742, discussion, 743–767.
- Boettiger, C. A., & D'Esposito, M. (2005). Frontal networks for learning and executing arbitrary stimulus-response associations. *Journal of Neuroscience*, *25*, 2723–2732.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*(4), 523–547.
- Cave, K. R., & Bichot, N. P. (1999). Visuospatial attention: Beyond a spotlight model. *Psychonomic Bulletin and Review*, *6*(2), 204–223.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. *Journal of Neuroscience*, *11*, 2382–2402.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
- Cowan, N. (2000). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–114, discussion, 114–185.
- Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*, 415–423.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
- Droll, J. A., Hayhoe, M. M., Triesch, J., & Sullivan, B. T. (2005). Task demands control acquisition and storage of visual information. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 1416–1438.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., & Van Essen, D. C. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the USA*, *102*, 9673–9678.
- Fu, W., & Gray, W. D. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology*, *52*, 195–242.
- Gordon, R. D., & Irwin, D. E. (2000). The role of physical and conceptual properties in preserving object continuity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 136–150.
- Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, *391*, 481–484.
- Gray, W. D., Schoelles, M. J., & Sims, C. R. (2005). Adapting to the task environment: Explorations in expected value. *Cognitive Systems Research*, *6*, 27–40.
- Haxby, J. V., Petit, L., Ungerleider, L. G., & Courtney, S. M. (2000). Distinguishing the functional roles of multiple regions in distributed neural systems for visual working memory. *NeuroImage*, *11*(5 Pt 1), 380–391.

- Hayhoe, M. M. (2000). Vision using routines: A functional account of vision. *Visual Cognition*, 7(1/2/3), 43–64.
- Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, 9, 188–194.
- Hayhoe, M. M., Bensinger, D. G., & Ballard, D. H. (1998). Task constraints in visual working memory. *Vision Research*, 38, 125–137.
- Hayhoe, M. M., Shrivastava, A., Mruczek, R., & Pelz, J. B. (2003). Visual memory and motor planning in a natural task. *Journal of Vision*, 3, 49–63.
- Henderson, J. M., & Hollingworth, A. (1999). The role of fixation position in detecting scene changes across saccades. *Psychological Science*, 10(438–443).
- Hester, R., & Garavan, H. (2005). Working memory and executive function: The influence of content and load on the control of attention. *Memory & Cognition*, 33(2), 221–233.
- Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiological Reviews*, 80, 953–978.
- Hollingworth, A. (2004). Constructing visual representations of natural scenes: The roles of short- and long-term visual memory. *Journal of Experimental Psychology: Human Perception & Performance*, 30, 519–537.
- 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.
- Hollingworth, A., Williams, C. C., & Henderson, J. M. (2001). To see and remember: Visually specific information is retained in memory from previously attended objects in natural scenes. *Psychonomic Bulletin & Review*, 8, 761–768.
- Ikeda, T., & Hikosaka, O. (2003). Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron*, 39(4), 693–700.
- Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parrasuraman & D. R. Davies (Eds.), *Varieties of Attention*, (pp. 29–61). Orlando, FL: Academic Press.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 25, 175–219.
- Land, M. F. (2004). The coordination of rotations of the eyes, head and trunk in saccadic turns produced in natural situations. *Experimental Brain Research*, 159, 151–160.
- Land, M. F., & Hayhoe, M. (2001). In what ways do eye movements contribute to everyday activities? *Vision Research*, 41, 3559–3565.
- Land, M., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, 28, 1311–1328.
- Lavie, N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, 133, 339–354.
- Li, W., Piech, V., & Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience*, 7, 651–657.
- Logan, G. D. (2004). Cumulative progress in formal theories of attention. *Annual Review of Psychology*, 55, 207–234.
- Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, 108, 393–434.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Mack, A., & Rock, I. (1996). *Inattention blindness*. Cambridge, MA: MIT Press.
- Marois, R., Yi, D. J., & Chun, M. M. (2004). The neural fate of consciously perceived and missed events in the attentional blink. *Neuron*, 41, 465–472.
- Maunsell, J. H. (2004). Neuronal representations of cognitive state: Reward or attention? *Trends in Cognitive Sciences*, 8, 261–265.
- McAdams, C. J., & Maunsell, J. H. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *Journal of Neurophysiology*, 83, 1751–1755.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review Neuroscience*, 24, 167–202.
- Mitchell, K. J., Johnson, M. K., Raye, C. L., & D’Esposito, M. (2000). fMRI evidence of age-related hippocampal dysfunction in feature binding in working memory. *Brain Research: Cognitive Brain Research*, 10, 197–206.
- Monsell, S. (1978). Recency, immediate recognition, and reaction time. *Cognitive Psychology*, 10, 465–501.
- Most, S. B., Scholl, B. J., Clifford, C. W., & Simons, D. J. (2005). What you see is what you set: Sustained inattention blindness and the capture of awareness. *Psychological Review*, 112, 217–242.
- Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattention blindness. *Psychological Science*, 12(1), 9–17.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Noles, N. S., Scholl, B. J., & Mitroff, S. R. (2005). The persistence of object file representations. *Perception & Psychophysics*, 67, 324–334.
- O’Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46(3), 461–488.
- O’Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–973, discussion, 973–1031.
- Platt, M. L., & Glimcher, P. W. (1998). Response fields of intraparietal neurons quantified with multiple saccadic targets. *Experimental Brain Research*, 121, 65–75.
- Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research*, 40, 1469–1487.
- Roelfsema, P. R., Lamme, V. A., & Spekreijse, H. (2000). The implementation of visual routines. *Vision Research*, 40, 1385–1411.
- Schmidt, B. K., Vogel, E. K., Woodman, G. F., & Luck, S. J. (2002). Voluntary and automatic attentional control of visual working memory. *Perception & Psychophysics*, 64, 754–763.
- Schwartz, M. F. (1995). Re-examining the role of executive functions in routine action production. *Annals of the New York Academy of Sciences*, 769, 321–335.
- Serences, J. T., Schwarzbach, J., Courtney, S. M., Golay, X., & Yantis, S. (2004). Control of object-based attention in human cortex. *Cerebral Cortex*, 14, 1346–1357.
- Shinoda, H., Hayhoe, M. M., & Shrivastava, A. (2001). What controls attention in natural environments? *Vision Research*, 41, 3535–3545.
- Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, 7(1/2/3), 1–15.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28, 1059–1074.
- Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin and Review*, 5, 644–649.
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9, 16–20.
- Sprague, N., Ballard, D., & Robinson, A. (in press). Modeling attention with embodied visual behaviors. *ACM Transactions on Action and Perception*.
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, 304, 1782–1787.
- Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 14, 107–141.
- Triesch, J., Ballard, D. H., Hayhoe, M. M., & Sullivan, B. T. (2003). What you see is what you need. *Journal of Vision*, 3, 86–94.

- Triesch, J., Sullivan, B., Hayhoe, M., & Ballard, D. (2002). Saccade contingent updating in virtual reality. In *Proceedings of the 2002 symposium on eye tracking research & applications* (pp. 95–102). New York: ACM Press.
- Ullman, S. (1984). Visual routines. *Cognition*, *18*, 97–159.
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *27*, 92–114.
- Wallis, G., & Bulthoff, H. (2000). What's scene and not seen: Influences of movement and task upon what we see. *Visual Cognition*, *7*(1/2/3), 175–190.
- Wheeler, M. E., & Treisman, A. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, *131*, 48–64.
- Yantis, S., Schwarzbach, J., Serences, J. T., Carlson, R. L., Steinmetz, M. A., Pekar, J. J., et al. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Neuroscience*, *5*, 995–1002.

Received July 5, 2006

Revision received February 15, 2007

Accepted March 30, 2007 ■

New Editors Appointed, 2009–2014

The Publications and Communications Board of the American Psychological Association announces the appointment of six new editors for 6-year terms beginning in 2009. As of January 1, 2008, manuscripts should be directed as follows:

- *Journal of Applied Psychology* (<http://www.apa.org/journals/apl>), **Steve W. J. Kozlowski, PhD**, Department of Psychology, Michigan State University, East Lansing, MI 48824.
- *Journal of Educational Psychology* (<http://www.apa.org/journals/edu>), **Arthur C. Graesser, PhD**, Department of Psychology, University of Memphis, 202 Psychology Building, Memphis, TN 38152.
- *Journal of Personality and Social Psychology: Interpersonal Relations and Group Processes* (<http://www.apa.org/journals/psp>), **Jeffrey A. Simpson, PhD**, Department of Psychology, University of Minnesota, 75 East River Road, N394 Elliott Hall, Minneapolis, MN 55455.
- *Psychology of Addictive Behaviors* (<http://www.apa.org/journals/adb>), **Stephen A. Maisto, PhD**, Department of Psychology, Syracuse University, Syracuse, NY 13244.
- *Behavioral Neuroscience* (<http://www.apa.org/journals/bne>), **Mark S. Blumberg, PhD**, Department of Psychology, University of Iowa, E11 Seashore Hall, Iowa City, IA 52242.
- *Psychological Bulletin* (<http://www.apa.org/journals/bul>), **Stephen P. Hinshaw, PhD**, Department of Psychology, University of California, Tolman Hall #1650, Berkeley, CA 94720. (Manuscripts will not be directed to Dr. Hinshaw until July 1, 2008, as Harris Cooper will continue as editor until June 30, 2008.)

Electronic manuscript submission: As of January 1, 2008, manuscripts should be submitted electronically via the journal's Manuscript Submission Portal (see the website listed above with each journal title).

Manuscript submission patterns make the precise date of completion of the 2008 volumes uncertain. Current editors, Sheldon Zedeck, PhD, Karen R. Harris, EdD, John F. Dovidio, PhD, Howard J. Shaffer, PhD, and John F. Disterhoft, PhD, will receive and consider manuscripts through December 31, 2007. Harris Cooper, PhD, will continue to receive manuscripts until June 30, 2008. Should 2008 volumes be completed before that date, manuscripts will be redirected to the new editors for consideration in 2009 volumes.