

# Task Demands Control Acquisition and Storage of Visual Information

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

Jochen Triesch  
University of California, San Diego

Brian T. Sullivan  
University of Rochester

Attention and working memory limitations set strict limits on visual representations, yet researchers have little appreciation of how these limits constrain the acquisition of information in ongoing visually guided behavior. Subjects performed a brick sorting task in a virtual environment. A change was made to 1 of the features of the brick being held on about 10% of trials. Rates of change detection for feature changes were generally low and depended on the pick-up and put-down relevance of the feature to the sorting task. Subjects' sorting decision suggests that changes may be missed because of a failure to update the changed feature. The authors also explore how hand and eye behavior are coordinated for strategic acquisition and storage of visual information throughout the task.

*Keywords:* change blindness, eye movements, task demand, visual working memory, visual attention

A central challenge in research on perception is to understand the principles that guide the selection of information from visual scenes. Scenes are typically composed of a complex array of stimuli, and the brain is fundamentally limited in its capacity to process and store this information. What guides the selection process? Although there is evidence that attention is attracted to locations that are salient by virtue of their stimulus properties, such as luminance or chromaticity (Itti & Koch, 2000, 2001; Parkhurst & Niebur, 2003, 2004), it is clear that, in general, the demands of ongoing visually guided behavior must account for much of the variance in attentional deployment. For example, Land, Mennie, & Rusted (1999) and Hayhoe, Shrivastava, Mruczek, & Pelz (2003) showed that when subjects perform tasks such as making tea or sandwiches, most of the fixations are tightly linked to the immediate demands of the task. Fixations in these tasks are overwhelmingly directed to areas relevant for guiding actions, such as grasping and moving objects, and very few fixations fall on objects that are irrelevant. Thus, direction of gaze can be an informative indicator of what information a subject is using in a scene. However, although fixation position and attention are tightly linked, the

mere presence of gaze at a particular location in the visual field does not reveal the variety of brain computations that might be operating at that moment. For example, when reaching to pick up a jar of peanut butter, are subjects processing information only about jar size and orientation to control the grasping movement, or is the jar represented as a more complete integrated structure in the context of an extended visual scene? Our goal in the present article is to define more precisely the extent to which the task constrains the acquisition of visual information from the scene.

A variety of experiments have indicated that the visual information acquired during a fixation may be quite specific. In an experiment by Ballard, Hayhoe, and Pelz (1995), observers copied simple colored block patterns 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 patterns twice, once before picking up a matching block and once before placement. Given the requirements of the task, a reasonable hypothesis is that block color is acquired during the first fixation and that the next fixation on the block is to acquire its location. A subsequent experiment in which changes were made to the block colors at different stages of the task supported the interpretation that the first and second fixations on a model block subserved different visual functions (Hayhoe, Bensinger, & Ballard, 1998; Hoffman, Landau, & Pagani, 2003). On the basis of these experiments, Hayhoe (2000) argued that vision involves the ongoing execution of special-purpose visual routines that depend on the immediate behavioral context and extract only the particular information required at the moment. The idea of visual routines was first introduced by Ullman (1984). The essential property of a routine is that it instantiates a procedure for acquiring specific information called for by the current cognitive agenda. Selection of just the task-specific information from a scene is an efficient strategy. Task-specific strategies not only circumscribe the information that needs to be acquired but also allow the

---

Jason A. Droll, Mary M. Hayhoe, and Brian T. Sullivan, Department of Brain and Cognitive Sciences, and Center for Visual Science, University of Rochester; Jochen Triesch, Department of Cognitive Science, University of California, San Diego.

This work was supported by National Institutes of Health/Public Health Service Research Grants EY07125, P41RR09283, and EY05729. We thank Keith Parkins for programming the virtual environment. Peter Skirko, Jeff Pelz, Keith Parkins, and Paul Illardi made significant contributions to the Fixation Finder software, which implemented algorithms developed by Constantin Rothkopf. We also thank Dana Ballard for discussions involving the development of these ideas.

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

visual system to take advantage of the known context to simplify the computation (Ballard, Hayhoe, Pook, & Rao, 1997).

Accumulating psychophysical and physiological evidence supports the notion that the visual system capitalizes on task context to selectively process information. The phenomenon of inattention blindness, introduced by Mack and Rock (1996), strikingly demonstrates failure to process unexpected visually salient objects when those objects are not required by the task. Subsequent experiments have shown that this blindness can occur even for periods of several seconds during sustained attention to other objects in the scene (Most et al., 2001; Simons & Chabris, 1999). Evidence that this selectivity may be specific to the task-relevant features of an object was demonstrated by Most et al. (2001). Detection of an unexpected cross passing across the scene during a tracking task depended on whether the cross matched the features of the attended targets. This supports the suggestion that subjects can adopt an attentional control setting that determines which features will gain access to neural processes responsible for visuomotor control, decision-making, and high-level visual analysis (Folk, Remington, & Johnston, 1992; Marois, Yi, & Chun, 2004; Shinoda, Hayhoe, & Shrivastava, 2001). These results are not easily explained with traditional models of vision that postulate early preattentive feature analysis followed by a late attentive analysis of objects. Neurophysiological evidence suggests that activity in even low-level cortical areas, traditionally considered to be preattentive, can be profoundly influenced by the demands of the task (Huk & Heeger, 2000; Li, Piech, & Gilbert, 2004; Ress & Heeger, 2003). For example, the spatial tuning of neurons in primary visual cortex of behaving monkeys depends on the required perceptual judgment (Crist, Li, & Gilbert, 2001). Roelfsema, Lamme, & Spekreijse (2000) demonstrated that the firing of neurons in the primary visual cortex (V1) depends not only on stimulus features in the receptive field but also on the task required of the animal. They also showed that addition of color specificity to a line-tracing task required increased processing time in V1 neurons (Roelfsema, Khayat, & Spekreijse, 2003). Thus, the extraction of even simple visual information requires active computation. Task-related effects are even more pervasive in higher cortical areas. For example, neurons in lateral intraparietal area respond to the presence of an object in their receptive field only when the object is behaviorally relevant (Gottlieb, Kusunoki, & Goldberg, 1998). Although it is tempting to consider selective neural activity as a correlate of visual salience, or constructing internal representations of the external world, it may be more appropriate to consider task-related neural enhancement as the instantiation of specific procedures, or elementary operations, required to accomplish a task. Understanding neural activity in this manner would require evidence of neurons whose activity was associated not just with the visual qualities of stimuli but also with the relationship the stimuli had with guiding behavior or decision making. In fact, such neural sensitivity can be found in a variety of brain areas involved in visuomotor transformations. Activity in neurons in dorsolateral prefrontal cortex (PFC) appear to code the conjunction of specific visual qualities with the learned motor response. For example, some neurons show selective responses for an image of a particular object but only when that image instructs an eye movement in a particular direction (Asaad, Rainer, & Miller, 1998; Miller & Cohen, 2001). Thus, the demands of the

task seem to be an intrinsic component of the brain's representational structure of visual information.

The idea that task relevance guides top-down selection of even simple feature information contrasts with the idea that visual information is represented and stored in the form of object files (Kahneman, Treisman, & Gibbs, 1992). This theory posits that when attention is directed to an object in a scene, a temporary representation called an *object file* is created and held in visual short-term memory and about three or four object files and their spatial locations can be held in memory at any time (Gordon & Irwin, 1996; Irwin & Andrews, 1996). Object file theory is consistent with claims that the units of short-term visual memory are integrated objects, not simple features (Luck & Vogel, 1997; Vogel, Woodman, & Luck, 2001). Another suggestion is that there is no cost to represent additional features, and thus, objects may be encoded in their entirety (Duncan, 1984, 1993). A similar proposal is the concept of object-based attention, which posits that attending to one feature of an object facilitates the representation of other features of the same object (Scholl, 2002). However, it is not clear how the concept of object files or object-based attention might extend to natural behavior, in which task demands are dynamic and specific to the immediate needs of the observer. One important consideration may be that different experimental paradigms impose different kinds of (often implicit) demands on subjects and do not necessarily reveal invariant or general properties of vision. Thus if a subject needs to pick up an object of a certain color, other information about the object, such as its height, may not be encoded, whereas in a memory test, subjects might encode all the features, as they are implicitly relevant. One of our goals in the current experiment, therefore, was to examine whether the task controls the acquisition and consequent representation of specific feature information from visual objects in this manner.

If the specific information requirements of the task are the critical factors controlling visual information acquisition, it is also likely that the task will similarly control the information that is stored in visual working memory. In natural behavior, observers select information from the environment for a particular purpose, and simple visually guided behaviors such as picking up objects and moving them around span several seconds and involve a series of gaze positions. Thus it is natural to suppose that the task is an important determinant of the visual representations that are stored over periods of a few seconds during ongoing behavior. It is generally accepted that the visual representations stored across saccades are limited to a small number of objects and that the attended items in prior views are the ones most likely to be remembered (Hollingworth, 2004; Irwin & Andrews, 1996; Rensink, O'Regan, & Clark, 1997). However, rather than passively representing the most recently attended objects in working memory, it may be necessary to specify exactly what tasks, or visual computations, the observer is engaged in from moment to moment in order to characterize the contents of working memory. For example, during tasks requiring active memory, many neurons in PFC store information of stimulus history throughout subsequent presentations of distracting stimuli (Miller & Desimone, 1994; Miller, Erickson, & Desimone, 1996). Thus, the task demand determines both what is attended and what is remembered.

A common approach to evaluating the contents of working memory in humans has been performance in change detection tasks. Change detection tasks typically present two successive

images that differ in the properties of a single object. The two images are interrupted with a mask, a mud splash, a blank screen, or an eye movement to disguise the transient that may otherwise be detected by low-level visual mechanisms. Reports of change blindness have emphasized the finding that subjects are notoriously poor at detecting changes, and this failure to notice changes has traditionally been interpreted as failure to retain information in working memory (O'Regan, Rensink, & Clark, 1999; Rensink, 2000a, 2000b, 2000c). However, because typical change detection paradigms do not clearly define the task demands, the specific use of working memory is unclear. Even when gaze is monitored, it is not clear what specific computations are being performed. In the present experiment, our manipulation of visual changes served a slightly different purpose. Whereas most change blindness paradigms test the capacity of visual memory, our aim was to describe the usage of working memory during more ordinary behavior. A salient feature of the experimental paradigm is that subjects' primary chore was to perform a sorting task involving the selection and pick-up of an object followed by a decision of where to put it down. We consider this task to be representative of the visuomotor demands required during many natural behaviors (Hayhoe et al., 2003; Land et al., 1999; Pelz, Hayhoe, & Loeber, 2001). This task was used to explore the hypothesis that visual representations are task specific. Hand and eye movements were measured for the purpose of inferring what microtasks, or visual operations, were performed throughout this task. We were then able to evaluate the consequence of a feature change on eye and hand movements in addition to explicit change detection. Thus, in this article we examine what information is being used and stored in the simple act of picking up and putting down an object.

Some progress on understanding the specificity and duration of task-relevant visual representations during pick up and put down has been made using a paradigm developed by Triesch, Ballard, Hayhoe, & Sullivan (2003). In a virtual reality environment, subjects picked up each of five short or tall bricks on one side of a table and placed them on one of two conveyor belts. On 10% of trials, brick height changed while the brick was being carried to the belts. One group of subjects was told to place each successive brick on the front conveyor belt. Thus, brick height was always irrelevant. A second group of subjects was told to first pick up the tall bricks and place them on the front conveyor belt and then to place the short bricks on the same front conveyor belt. For this condition, subjects needed to attend to brick height only for pick-up, because the final location of put-down was always the same. Changes in brick height were rarely detected in the first condition but were more likely to be noticed in the second. This suggests that subjects were more likely to encode brick height when it was relevant to the pick-up task. In addition to the first two conditions, a third group of subjects was told to pick up the tall bricks and place them on the front conveyor belt and then to pick up the short bricks and place them on the back conveyor belt. Thus, for this third group, brick height was relevant for both pick-up and put-down. Rates of change detection were most frequent in this third condition when height was relevant for both pick-up and put-down. One interpretation of these results is that in the second condition brick height was unlikely to be retained in visual memory even for a few hundred milliseconds after pick-up but was more frequently retained when it was needed later in the task for put-down, as required in the third condition.

Our goal in this experiment was to identify more precisely the information acquired and held in memory and provide a more definitive test of the task-dependence hypothesis. In the Triesch et al. (2003) experiment the only feature of the bricks that was used in the task was height. Is it the case that brick features that were not required, such as color and shape, were never encoded? This would mean that fixating an object and attending to it would not necessarily bind the features of the brick into some object representation, or object file. The current experiment extended the Triesch et al. (2003) study and is a stronger test of the hypothesis that acquisition and memory for individual object features is sensitive to the demands of the task. Another way we wished to extend the Triesch et al. (2003) experiment was to control more precisely the use of memory in performing the task. In that experiment, a variety of strategies were available to the subject, some of which required little or no memory for brick features. For example, subjects were able to decide where to put down the brick at the time of pick-up and to execute this movement without need for memory of brick size.

In the current experiment, subjects performed a brick sorting task similar to that used in Triesch et al.'s (2003) study but with several modifications. First, rather than occasionally changing only the relevant feature, changes were made to features that were either relevant or irrelevant to the sorting task. Rates of change detection were used to assess the degree to which information on each feature value was stored throughout the pick-up and put-down decisions. Second, although subjects were able to predict which feature dimension is relevant for the sorting decisions (e.g., color, height, width, texture), they were not able to predict the relationship between a feature value (e.g., red, tall, wide or thin striped) and whether that brick should be selected for pick-up or where it belonged for put-down. Instead, during each trial, cues in the scene informed the subject which bricks were appropriate for pick-up (e.g., "pick-up a tall brick") and where to place the brick (e.g., "Tall bricks go on the right belt, short bricks go on the left belt.") on that particular trial. Because each trial could require a different feature value for pick-up selection, and impose different rules for put-down, we ensured that subjects were performing pick-up and put-down decisions at different times within each trial. Third, subjects performed two blocks of trials. In one trial block, the same feature was used for pick-up and put-down. In a second trial block, one feature was used for pick-up, and a second feature was used for the put-down decision. We were interested in the following questions: Will subjects retain information on a feature used for pick-up when it is no longer needed? Will subjects acquire the put-down feature before it is needed or delay acquisition until absolutely necessary?

## Method

### *Equipment*

Subjects wore a Virtual Research (Aptos, CA) V8 head mounted display, shown in Figure 1. The helmet was equipped with a magnetic head tracking device (Fastrack; Polhemus, Colchester, VT) that measured the head's position and orientation with respect to a fixed reference frame. The magnetic tracker operated at 120 Hz with a 6-ms internal latency. This information was passed on to the graphics engine to determine the view-point(s) from which to render the virtual scene with 1–2 frame latency. The visual display was generated by a Silicon Graphics computer at a rate of 60





Figure 1. Brian T. Sullivan demonstrating the use of virtual reality goggles and Phantom (SensAble Technologies, Woburn, MA) haptic feedback.

Hz and was rendered in stereo on two LCD screens in the headset with  $640 \times 480$  pixel resolution.

Two devices monitored the movements of the eyes. On the left side of the helmet, an Applied Science Laboratory (ASL) 501 video-based eye tracker monitored the position of the eye with 60 Hz temporal resolution and approximately  $1^\circ$  in accuracy. Eye position was calibrated by having the subject look at each of nine points on a  $3 \times 3$  grid. Eye, head, and gaze directions were recorded throughout the experiment and saved in each data file. In addition to the data stream, a video record of the scene, with eye position superimposed, was captured using a Hi-8 video recorder. An image of the left eye was also included in the video record to monitor track loss. Because of real-time delays in the ASL signal, it was necessary to use a different eye tracker to trigger changes during saccades. To do this, a

limbus eye monitor (ASL 210) mounted on the right eyepiece monitored the velocity of the eye with 1000 Hz resolution. The overall latency for scene updating is less than a mean time of 50 ms, with a range of 40–55 ms, when triggered by saccades with  $15^\circ$  amplitude. This short latency for scene updating allows for changes to be displayed before the beginning of the following fixation. This system for doing saccade-contingent updating in virtual reality and evaluation of its performance is more thoroughly described in Triesch, Sullivan, Hayhoe, & Ballard's (2002) study.

Force feedback from physical interaction with objects in the environment is given with two haptic stimulation devices that allow subjects to grasp objects while experiencing realistic forces. Two Phantom-3 devices from SensAble Technologies (Woburn, MA) were used in opposition—one for the index finger and one for the thumb (see Figure 1). Force feedback was provided to the subject at a rate of 1 kHz. Thumb and index finger position are represented visually as small red spheres displayed in the virtual world (see Figure 2). Thimble position was recorded throughout the duration of the experiment at a rate of 60 Hz, with a real-time delay of 8 ms. Despite some equipment with higher sampling rates, all data were recorded at 60 Hz. The usable work space volume was  $55 \text{ cm} \times 55 \text{ cm} \times 40 \text{ cm}$ . The geometry of the visual stimulus was matched to the physical geometry of the workspace. In other words, size measurements of the virtual bricks reflected the height and width dimensions as if the brick were an actual object in the real world.

### Sorting Task

The basic task of the current experiment was to select one brick from an array and to sort this brick onto one of two conveyor belts. The bricks were defined by several features, and a pick-up cue indicated which feature value was relevant for a particular trial. After picking up the brick, a put-down cue was displayed to guide the sorting decision. The brick was placed on the appropriate conveyor belt, removing the brick from the scene, which initiated a new trial with a new pick-up cue and array of bricks. Thus, because the put-down cue was presented after pick-up, the put-down decision was separated in time and space from pick-up, and the representations of the relevant object feature needed to be stored until the put-down decision was made. In one condition, subjects performed a task




Figure 2. Scene during a single trial of the *One Feature* condition 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 themselves, (D) decides on which conveyor belt the brick belongs based on a put-down cue, (E) guides the brick to the conveyor belt, and (F) sets the brick on the belt where the brick is carried off.

in which only one feature dimension was relevant for both pick-up and put-down (e.g., color). In another condition, different features were used for pick-up and put-down (e.g., color for pick-up, height for put-down). In a small fraction of trials, after the subject picked up the brick, but before put-down, a change was made to one of the brick features. The changed feature was either relevant or irrelevant to the current sorting task.

Figure 2 shows the scene visible to the observer within the helmet. In both conditions, an array of five bricks appeared on a table (see Figure 2a). Each brick was defined by four features: color, width, height, and texture. Each feature was in one of two different states. Brick color was either red or blue; width was either wide or thin (7.7 and 6.0 cm); height was either tall or short (10.0 and 8.0 cm); and texture was either thick or thin diagonal stripes (0.50 cycles/cm and 0.25 cycles/cm). These differences in feature values were very clearly discriminable. At least one instance of each feature state appeared within the array of bricks. The example shown in Figure 2 is taken from a block of trials in which color was relevant. Above the bricks a pick-up cue indicated which feature state the subject should select (see Figure 2a). In this particular trial the subject is instructed to pick-up a red brick because the red rectangle is highlighted by the yellow arrows. The relevant feature value was varied pseudorandomly from trial to trial with the constraint that each pick-up instruction occurred five times in a block of 10 trials. (The relevant feature value was always presented on the bottom half of the pick-up cue.) Subjects were instructed to select any brick with the relevant feature value regardless of the value of the other three features. After the subject lifted the brick, the pick-up cue disappeared from the scene. The subject then carried the brick toward the put-down area. Immediately after pick-up the put-down cue was displayed in between two conveyor belts (see Figure 2d). 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, the cue in Figure 2d instructs the subject to place the brick on the right belt because the brick is red. (When width, height, or texture was relevant for pick-up or put-down, the cues consisted of two black rectangles differing in their width or height or had two rectangles with thin or thick diagonal stripes.) Similar to the pick-up cue, the put-down cue (e.g., "wide on right/thin on left" or "wide on left/thin on right") was determined pseudorandomly across trials. After the subject placed the brick on the conveyor belt, the brick was carried out of the scene and the put-down cue disappeared, indicating the end of a trial. The next trial began immediately, presenting five new bricks and a new pick-up cue. This general trial structure, requiring brick selection and placement depending on object features, constituted the sorting task. Apart from the constraints mentioned, the details of task performance were entirely at the discretion of the subject. Subjects were encouraged to perform the task at a pace they considered comfortable and natural.

In the *One Feature* condition, only one of the four features was relevant for both pick-up and put-down. The same feature remained relevant for the entire block of 80 trials. Because in any given trial the same feature was relevant for both pick-up and put-down, subjects could either retain information on the relevant feature in working memory throughout the trial or reacquire this information during the put-down decision. The decision of whether to store this information after pick-up may depend on whether it is needed later for put-down. To investigate whether subjects retained information, we had subjects perform another block of 80 trials in which one feature was used for pick-up and a second feature was used for put-down (the *Two Feature* condition). For example, brick color might be relevant for pick-up and brick height relevant for put-down. By making different features relevant for pick-up and put-down, our aim was to test (a) whether information that was no longer needed for sorting would be stored and (b) whether subjects acquired the put-down feature at the time of pick-up (before it is needed) or delayed acquisition until after inspection of the put-down cue.

The same two features were used for pick-up and put-down, respectively, for a block of 80 trials. In the *Two Feature* condition the same

feature was relevant for pick-up as had been used in the *One Feature* condition, but one of the three remaining features was used for put-down. Seventy-two subjects were divided into four groups of 18. For each group, a different feature was relevant in the *One Feature* condition. Each of the four groups of 18 subjects were further subdivided into three groups. Each of these subgroups used one of the remaining three features for put-down in the *Two Feature* task. Because of the large number of conditions, and the laborious nature of data collection, subjects always performed the *One Feature* condition followed by the *Two Feature* condition. This puts constraints on interpretations of the *Two Feature* condition because of possible carry-over effects from the first condition. The two conditions were separated by a break, during which the eye tracker was recalibrated and subjects were reminded that future trials would require a different feature for put-down. Before performing the experiment, subjects had eight practice trials of the *One Feature* task and eight trials of the *Two Feature* task.

### Change Detection Task

To test whether task relevance influences visual information acquisition, we had subjects perform a change detection task concurrent with the sorting task. A potential criticism of the Triesch et al. (2003) experiment is that subjects were not told that any changes might occur, or what to do if they did, so the perceived changes were likely to be underreported. This is a difficult problem to deal with, because asking subjects to detect changes alters the task and the distribution of attention. The advantage is that subjects know what to do when a change occurs. Thus, 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 in between the conveyor belts (see Figure 2d–2f). Thus, the overt movement to the trash bin served as an explicit measure of feature change detection. To minimize intrusion of the secondary change detection task, we ensured that changes occurred in less than 10% of trials.

To ensure that subjects were familiar with the trash bin and comfortable with the instructions, yet naive at detecting actual changes, sometime between the 10th and 16th practice trial the experimenter interrupted the subject 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, suggesting comprehension of the change detection task.

Following the practice trials, we recalibrated the eye tracking equipment to prevent any tracking errors. At this time subjects were reminded that, similar to the format of the practice trials, they would perform 80 trials in which they would pick up and put down bricks on the basis of the value of one feature, pause briefly to recalibrate the eye tracker, and then perform another 80 trials in which the same feature was used for pick-up but the second feature would be used for put-down. Subjects were reminded a third time to place any brick that changed in any of the four features into the trash bin.

Up to eight changes could occur within both the *One Feature* and the *Two Feature* tasks. Within each block of 80 trials, up to two changes could occur for each of the four brick features. Change trials were identified at the beginning of the experiment; one change trial was assigned within each group of 10 trials and no two consecutive change trials were permitted. The order of the feature changes was randomized. For a designated change trial to induce a feature change, two conditions needed to be met. First, the brick in hand needed to be within a change zone, that extended from the halfway point between the conveyor belts and the brick array to the back edge of the conveyor belts. Second, the subject needed to make a saccade while the brick was in the change zone. The position requirement was intended to separate the time at which the brick would be fixated for pick-up and put-down. The requirement for a saccade was to mask the transient gen-

erated by the change. Note that because subjects were instructed to grasp the bricks with their finger on the back of the brick and their thumb on the front, changes in brick height or width were not accompanied by force changes on the subjects' fingers. Thus any changes in width or height did not conflict with grasp posture.

### Subjects

Seventy-two subjects participated in the experiment for \$10 per hour. Experimental sessions typically lasted about 1 hr. Subjects were recruited through posters around the University of Rochester. Their ages ranged from 18 to 39 years. A small number of additional subjects were excluded because of poor eye tracking or failure to follow instructions.

### 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° per second. This high threshold was used because of the noise present in the tracking signal. All recorded fixations needed to meet the additional criteria of having angular velocity less than 50° per second for at least 50 ms. Successive fixations occurring less than 50 ms apart, and with a displacement of less than 1.5°, were consolidated. 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 with 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. Eye position was monitored for all 72 subjects. Automated fixation analysis is sensitive to noise in the tracking signal, so it was necessary to screen subjects on the basis of 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 Fixation Finder. Analysis of the video records resulted in the selection of 43 subjects for whom eye position was judged adequate for automated analysis. On the basis of the video records, we are confident that the results reported for these 43 subjects are representative of all 72 subjects. Although these 43 subjects are not equally distributed across the four main subject groups using different features for sorting, none of the eye movement analysis involves the comparison between these four groups (color = 10 subjects, width =

8 subjects, height = 10 subjects, texture = 15 subjects). Further details of eye position analysis are included in the relevant sections of Results.

## Results

There are a number of aspects of subjects' performance that need to be considered. We first review general properties of performance during normal trials without feature changes to establish the behavioral context of the sorting task. Next we consider both explicit and implicit performance on the change detection task and the effect of relevance of the change to the sorting task. Finally, we address effects of recent trial history on hand and eye movements.

### 1. Performance on Normal Trials

#### 1.1 Sorting Accuracy

Although the primary measure was performance on the change detection task, performance on the sorting task was also monitored to ensure that subjects had no difficulty comprehending task instructions and discriminating brick features. Subjects performed the sorting tasks with high accuracy. In both the *One Feature* and the *Two Feature* tasks, subjects selected an improper brick or put the brick down on the incorrect belt in less than 1% of trials.

#### 1.2 Pattern of Hand Movements

Subjects performed the sorting task with a predictable pattern of hand and eye movements. Hand movements throughout all trials were monitored by recording the position of the left thimble, usually attached to the thumb. For convenience we refer to this simply as *hand position*. The position and velocity of the hand throughout an example trial is shown in Figure 3. After picking up a brick, subjects carried it toward the put-down cue and then slowed their movement before moving their hand to the belt. We categorized each stage of movement as having occurred during pick-up, sorting, or put-down, depending on hand position (see Figure 3A). Pick-up movements occurred anywhere behind an invisible vertical plane 10 cm in front of, and parallel to, the plane bisecting the brick array. Put-down movements occurred within a three-dimensional radius 7 cm from the center of the brick where




Figure 3. A: Left thimble position as seen from above. B: Left thimble velocity throughout a sample trial. Movements were separated into three stages on the basis of thimble position: pick-up (black crosses), sorting (green squares), and put-down (red circles). sec = seconds.



the finger last touched the brick. Sorting movements included any movement in between pick-up and put-down. The velocity profile for the hand movement and corresponding categorization is shown in Figure 3B. Instantaneous velocity was calculated by using the three-dimensional distance the right thimble traveled across each 60-Hz sample. Note that the dip in the velocity profile corresponded to the time at which subjects fixated either the put-down cue or the brick in hand, presumably while making the decision of where to place the brick. Although this temporary decrease in velocity during sorting occurred during nearly all trials, the hand came to a complete rest in only a small fraction of trials. Generally, the hand movement was continuous throughout each trial. Factors influencing the duration of the hand movement during sorting and put-down are considered in Sections 2 and 3.

### 1.3 Pattern of Fixations

Another characteristic aspect of performance was the order in which areas of the scene were fixated. To illustrate this behavior, we analyzed eye position during all trials for 43 of the 72 subjects as described in the Method section. We identified the start and end time for fixations directed to five areas of interest: the pick-up cue, bricks in the array, the brick being carried, the put-down cue, and the conveyor belts. The start and end times to each of these areas were then normalized with respect to the total duration of that trial. Thus time is expressed as a percentage of total trial duration. Normalized times for each area were averaged across trials for a single subject and then averaged across subjects. This gave the probability of fixating each area as the trial progressed. Figures 4A–4E plot the probability of fixating each area throughout the course of a trial during the *One Feature* condition. Only trials with no feature changes were included for this analysis. Early in the trial, subjects were most likely to fixate the pick-up cue, as shown in Figure 4A. They next fixated one of the five bricks in the array, as shown in Figure 4B. Note that subjects preferentially fixate the central bricks in the array. After selecting which brick to pick up, fixations were then directed to the brick in hand (first peak in Figure 4C). Subjects then fixated the put-down cue (see Figure 4D) and then refixated the brick being carried (see the second peak in Figure 4C) while guiding the brick to the left or right conveyor belt (see Figure 4E). (For the analysis in Figure 4, fixations that landed on the brick in hand while occluding the conveyor belt were counted as fixations to both the brick in hand and the conveyor belt. Thus, toward the end of the trial, the distributions for fixating the brick in hand and the belt overlap.)

Observations during natural behavior have suggested that the target of fixations are indicative of the immediate needs for the task. The predictable pattern of fixations in Figure 4 suggests that subjects segmented each trial into a sequence of subtasks that required eye movements to obtain the necessary information. For example, once a brick was selected and picked up, the subject needed to make a decision of where to put down the brick. The subtask of planning where to place the brick would presumably require a fixation to the put-down cue. Once subjects knew the put-down rule, they needed information on the feature value to make the sorting decision. This information might either be stored in working memory or be acquired with a refixation to the brick in hand. Return fixations to the same object in a scene have been interpreted as evidence that the required visual information was

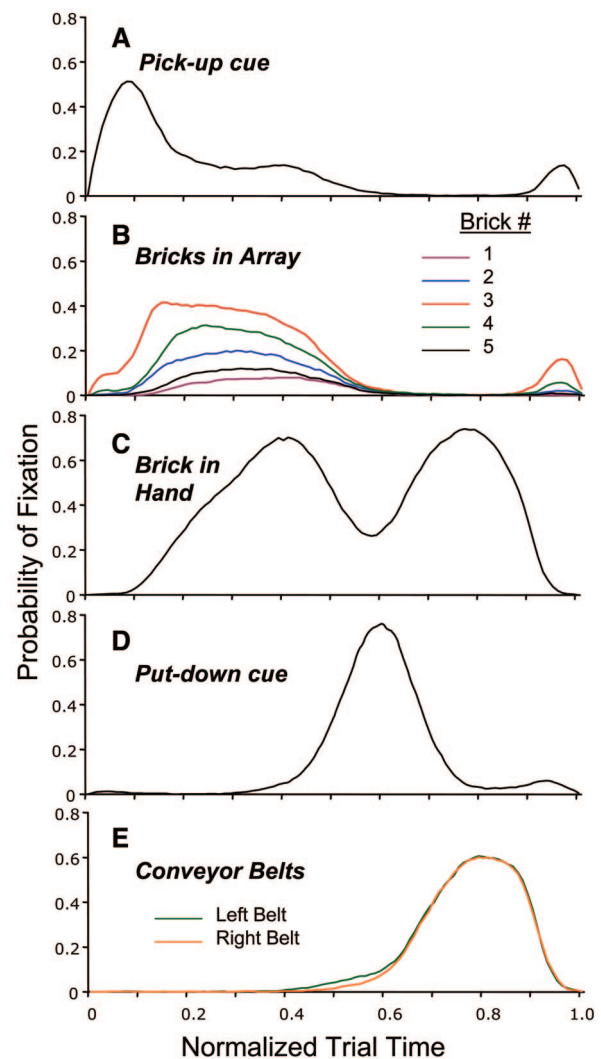


Figure 4. Probability of fixating each area in the scene throughout each trial during the *One Feature* condition. Trial length is normalized, and probability is averaged across at least 72 trials in which no change occurred during the *One Feature* condition and then across 42 subjects.

not present in working memory (Ballard et al., 1995). Thus, subjects' refixations on the brick after fixating the put-down cue may indicate whether the put-down feature was in working memory. Because the fixation time histogram in Figure 4 averages across a normalized timeline, the order in which subjects fixate different areas is not explicit. Therefore, we analyze fixation sequences in the following section.

### 1.4 Fixation Sequences

To more explicitly examine fixation sequences, we tabulated the probability of fixating each area following a fixation to each of the other areas. This tabulation is expressed in the transition matrix in Table 1 for each experimental condition. A particular fixation directed to one of the five areas was categorized by row. That particular fixation can be redirected to any of the remaining four areas in the scene, categorized by column. (Multiple fixations

Table 1  
Transition Probabilities for Gaze Location Within a Trial in the  
One Feature and Two Feature Tasks

| Current fixation        | Next fixation |     |     |     |     |
|-------------------------|---------------|-----|-----|-----|-----|
|                         | 1             | 2   | 3   | 4   | 5   |
| <i>One Feature task</i> |               |     |     |     |     |
| 1. Pick-up cue          | —             | .84 | .05 | .10 | .01 |
| 2. Array                | .43           | —   | .10 | .42 | .05 |
| 3. Brick                | .03           | .01 | —   | .52 | .44 |
| 4. Put-down cue         | .00           | .01 | .19 | —   | .80 |
| 5. Conveyor belt        | .01           | .06 | .04 | .89 | —   |
| <i>Two Feature task</i> |               |     |     |     |     |
| 1. Pick-up cue          | —             | .87 | .06 | .06 | .00 |
| 2. Array                | .44           | —   | .10 | .42 | .04 |
| 3. Brick                | .01           | .06 | —   | .48 | .45 |
| 4. Put-down cue         | .00           | .01 | .29 | —   | .70 |
| 5. Conveyor belt        | .03           | .09 | .06 | .83 | —   |

*Note.* Values in each row represent the percentage of fixations that were directed to each of the four areas given a fixation to the previous area. Values are averaged across trials for each subject and then across subjects. Brick = brick in hand. The largest difference between the two task conditions is the transition from put-down cue to either the brick in hand or the conveyor belt.

within the same area are not considered.) A slightly different set of criteria was used to categorize fixations for this analysis (see Appendix A). Thus, values from this analysis cannot be directly compared with values plotted in Figure 4 (see Section 1.3). (Standard errors of the mean across subjects ranged between 0.1%–3.4%, with an average of 1.3%.)

Table 1 shows that subjects' fixations were directed to areas in the scene that were relevant for the immediate subtask. A fixation to the pick-up cue was most commonly followed by a fixation to the array of bricks. A fixation to the array of bricks was followed by a fixation to (a) either the pick-up cue or to the put-down cue, (b) occasionally to the brick in hand, and (c) rarely to the belts. (Trials typically ended with fixations to the belt area; occasionally, subjects would refixate the put-down cue before the trial ended.) Stepwise traces through the table can be used to generate a set of fixation sequences that are similar to those observed throughout entire trials.

When do subjects acquire the information necessary for put-down? To examine this, we concentrated on fixation transitions following the fixation on the put-down cue. We reasoned that if subjects had retained the put-down-relevant feature in working memory, then the decision of where to place the brick could be made immediately after acquiring the put-down cue. Conversely, if the feature was not in working memory, subjects may need to refixate the brick before making the sorting decision. In both the *One Feature* and *Two Feature* conditions, subjects most frequently directed their gaze from the put-down cue to the belt area, suggesting that working memory was most commonly used as the basis of the sorting decision. However, note that the largest difference between the transition matrices in Table 1 is the probability of fixating the brick in hand following a fixation on the put-down cue. Fixations back to the brick in hand following a fixation on the put-down cue are more common in the *Two Feature*

task (29%) than in the *One Feature* task (19%; paired *t* test,  $p < .01$ ). (Conversely, fixations to the belt were more common following a fixation on the put-down cue in the *One Feature* task [80%] than in the *Two Feature* task [70%; paired *t* test,  $p < .01$ ].) This suggests that when a feature had not yet been required for a task, it is less likely to be retained in working memory than is a feature that has already been used.<sup>1</sup> Thus, although the preference to commit to a sorting decision using working memory suggests that subjects can retain feature information, the increased frequency of refixating the brick when the task requires new information suggests that information is occasionally acquired on a need-to-know basis just in time.

## 2. Performance on Change Trials

A primary purpose of the experiment was to evaluate what information was selected and what information was stored in working memory during visuomotor tasks. To this end, we used feature changes as a probe for working memory use. In trials with a feature change, the change most often occurred following pick-up, either as the subject was making a saccade toward the put-down cue or on the saccade from the put-down cue back to the brick. Thus, during a change trial, different information was available during fixations to guide pick-up (see the first peak in Figure 4C) than on subsequent fixations used during put-down (see the second peak in Figure 4C). If a subject detects a feature change, this suggests storage of visual information across the intervening period in which the put-down cue was fixated. Greater sensitivity to changes to features when they are relevant to the sorting task would suggest that information in working memory is specific to the immediate needs of the subject.

### 2.1 Effect of Task on Change Detection

*One Feature condition.* We first examined whether the task relevance of a feature influenced subjects' ability to detect a change in that feature. For each subject we calculated the percentage of changes detected, and we averaged this performance across features and subjects. Subjects reported feature changes when in fact no change occurred (false alarm) in less than 1% of trials. Change detection for relevant and irrelevant features in the *One Feature* task is plotted in Figure 5A. In the *One Feature* task, rates of change detection were higher for task-relevant (37%) than for task-irrelevant (18%) features,  $F(1, 66) = 18.37$ ,  $p < .01$ . These results suggest that the task is an important factor guiding the acquisition and consequent memory representation of objects. The results suggest that individual features of objects are represented preferentially, in contrast with the idea that visual information is represented in the form of object files containing integrated object properties (Irwin & Andrews, 1996). Object file theory predicts that irrelevant information on one feature (e.g., color) is acquired and retained with equal strength when information on another feature (e.g., shape) is relevant to the task. Note that overall rates of change detection are quite low. However, this does not necessarily mean that subjects were not retaining the changed feature. This issue is discussed in Section 2.4.

<sup>1</sup> Subjects may of course acquire information in peripheral vision without an overt fixation.






Figure 5. Average rate of change detection for relevant and irrelevant features. A: *One Feature* task; same feature relevant for pick-up and put-down. B: *Two Feature* task; different features relevant for pick-up and put-down. Error bars represent standard error of the mean between subjects.

The pattern of change detection revealed in Figure 5A raises three questions. First, what is the lifetime of an acquired feature? The lifetime of visual representations across multiple fixations is a matter of considerable debate. Hollingworth and Henderson (2002) found evidence to suggest that object representations can be stored over at least nine fixations during scene viewing when preparing for an upcoming memory test. Yet when subjects were engaged in a block-copying task, Ballard et al. (1995) observed frequent return fixations onto the same area of a scene, and Hayhoe et al. (1998) observed high rates of change blindness in this same task, suggesting that information acquired during visuomotor tasks is discarded after it has served its immediate purpose. If a feature was used for pick-up and was not required for put-down, would subjects retain the pick-up feature simply because it had been relevant in the past, or would it be discarded because it was not immediately needed? Second, can information on a presently irrelevant feature be acquired and stored for future use? In the *One Feature* task, subjects biased the acquisition and retention of information to only one feature that was relevant throughout the duration of every trial. Is the information acquired during pick-up restricted to the information relevant to the immediate microtask of pick-up selection? Or, do subjects also acquire and store a second brick feature that would be used in a later microtask, such as the put-down decision? Third, if subjects can acquire and store a second object feature, how would this influence the acquisition of the remaining two irrelevant features?

Performance in the *One Feature* task suggests that irrelevant features were less likely to be stored or updated. However, the *One Feature* task required the acquisition of only one feature. An object file may not have been formed in working memory because there were an insufficient number of acquired features. As a greater number of object features become relevant, the representation of the entire object may become more robust. Thus, if subjects are holding multiple relevant features in working memory, this may improve memory for the remaining irrelevant features. The three questions detailed above are addressed in the following section.

**Two Feature condition.** In the *Two Feature* condition subjects used the same feature for pick-up as had been used in the previous *One Feature* condition, but a second feature was used for the put-down decision. Again, subjects reported feature changes when in fact no change occurred (false alarm) in less than 1% of trials. Change detection performance in the *Two Feature* condition is

shown in Figure 5B. In this condition, change trials were classified as either relevant for pick-up, relevant for put-down, or irrelevant. Rates of change detection differed between conditions of task relevance,  $F(2, 62) = 7.00, p = .01$ . Changes were detected more often when the change was relevant to pick-up (48%) or put-down (46%), than when they were irrelevant (28%),  $F(1, 54) = 22.80, p < .01$ ;  $F(1, 53) = 10.97, p < .01$ . As in the *One Feature* condition, higher rates of change detection for relevant features than irrelevant features suggests that object features are acquired and stored preferentially, depending on their task relevance. This again is in contrast to theories that posit integrated object representations as a result of deployment of attention to an object. The pattern of results also addresses the points raised above. First, changes to pick-up relevant features were detected at similar rates during the *One Feature* and *Two Feature* conditions. This suggests that once a feature is used for one subtask, it is not discarded immediately. Storing a previously used feature regardless of its immediate irrelevance argues against the strongest interpretation of a microtask model in which information in working memory pertains only to the immediate task. Second, changes to a feature used for put-down in the *Two Feature* condition are more likely to be detected than when this same feature was irrelevant in the *One Feature* task (see Appendix B). To detect a change, subjects must acquire and store information on the feature before the change. Detection of a put-down-relevant change suggests that subjects acquire information on the put-down-relevant feature before it is immediately required for the sorting decision. Thus, subjects can acquire not just features immediately relevant for the pick-up task but also other features that are expected to be relevant in the short-term future. Third, change detection was not significantly different for irrelevant features between the *One Feature* and *Two Feature* conditions (see Appendix B). The selective increase in rates of change detection for put-down-relevant features suggests that acquisition of the put-down feature before the change did not result in an obligatory acquisition and storage of the remaining irrelevant brick features. This is in contrast with results predicted from positing the use of object files (Kahneman et al., 1992). Lastly, it should be noted that rates of change detection for pick-up and put-down-relevant features were not significantly different,  $F(1, 54) = 0.25, p = .62$ . This suggests that the pick-up and put-down features may be stored with similar strength throughout each trial.

## 2.2 Effect of Feature on Change Detection

The analysis thus far has focused on how task goals influence the acquisition and storage of object features. Another potential influence on acquisition of visual information is salience. The specific brick feature values were chosen for the experiment with no particular criteria other than that they be clearly distinguishable from one another (e.g., red vs. blue). Near-perfect performance in the sorting task suggests that subjects had no difficulty distinguishing between feature values. Nevertheless, each feature dimension, or differences between the values within each feature dimension, may have intrinsic differences in detectability or salience. Visual salience is known to attract attention (Parkhurst & Niebur, 2003, 2004; Yantis & Egeth, 1999; Yantis & Hillstrom, 1994) and thus may consequently influence change detection (Cole, Kentridge, & Heywood, 2004). To investigate the potential influence of visual

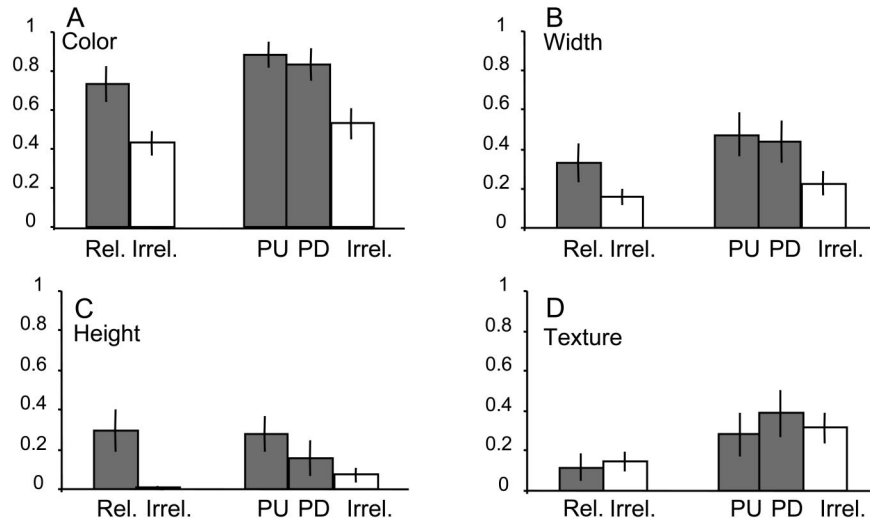


Figure 6. Effect of task relevance for the different features. Relevant (Rel.) changes are represented by gray bars; irrelevant (Irrel.) changes are represented by white bars. Each graph plots performance in the *One Feature* task (left) and *Two Feature* task (right). PU = pick-up; PD = put-down.

salience on change detection, we examined rates of change detection for each individual feature. Figure 6 shows the effect of task relevance on each of the four brick features for both the *One Feature* and *Two Feature* conditions. There was an effect of feature in both the *One Feature* and *Two Feature* tasks for relevant and irrelevant features,  $F(3, 66) = 7.94, p < .01$ . Note that the lack of a task-relevant effect for texture changes may have been due to a floor effect. The fact that the changed feature value influenced rates of change detection suggests that a strictly top-down model of highly specific task driven control cannot fully explain how visual information is acquired.

### 2.3 Fixation Strategies and Change Detection Performance

Another potential influence on performance in the change detection task is the frequency with which subjects fixate the brick before and after the change. Subjects invariably fixate the brick while picking it up, before the change. Table 1 reveals that subjects refixated the brick on 19% to 29% of trials following fixation to the put-down cue. During change trials, these refixations occurred after the change. Are these refixations associated with increased likelihood of detecting changes? One way to address this is to compare rates of refixation during change trials to rates of change detection. However, such an analysis would yield ambiguous results. Because of the frequent proximity of the brick to the fovea as the subject fixated the put-down cue, it is possible that some changes were detected using peripheral vision. If a change was detected in peripheral vision, it would be natural to expect fixations to the brick while it is being held in between the conveyor belts, as the brick is guided to the centrally located trash can in response to subjects' detection of the change. Thus, refixations during change trials may be either a contributing cause of detecting a change or a consequence of detecting a change. To avoid this complication, we looked at the correlation between performance in the change detection task and the probability of refixating the brick

in normal trials. This is plotted in Figure 7. There was a significant relationship between performance in the change detection task and the probability of fixating the brick following a fixation to the put-down cue ( $p < .01$ ). Thus, this behavior may have facilitated change detection. However, this relationship was moderate and accounted for only 15%–18% of the variance in performance. Note that the different intercepts of the two regression lines reflect the overall increased rate of change detection in the *Two Feature* condition. Further analysis of refixations to the brick in hand is described in Sections 3.1 and 3.2.

### 2.4 Sorting Performance on Missed Trials

Change detection requires that subjects represent both the pre- and postchange information, whereas missed changes suggest fail-




Figure 7. Probability of looking at the brick following a fixation on the put-down cue versus performance in the change detection task. *One Feature* task (open circles and dashed line;  $R^2 = 0.156, p < .01$ ) and *Two Feature* task (solid circles and line;  $R^2 = 0.176, p < .01$ ).

ure either to store or to acquire the changed information. The high rate of missed changes is surprising considering subjects' accurate performance on the sorting task, which demonstrates that subjects have no difficulty in acquiring the feature values needed for pick-up and put-down. We sought to understand why subjects were so insensitive to salient feature changes to an object in hand that they were using to guide their behavior. Subjects' insensitivity to changes may have been due to (a) failure to encode and retain the prechange stimulus, (b) failure to encode the postchange stimulus, or (c) failure to compare the retained prechange feature and the subsequent postchange feature (Simons, 2000). We examined subjects' put-down behavior when subjects missed a change in the state of the feature that was relevant for put-down (e.g., missed width changes when sorting on the basis of width). If misses were caused by failure to retain the prechange feature, subjects would be expected to acquire the postchange feature value and consequently sort the brick by the new value. However, if misses were caused by failure to encode the new postchange feature value, subjects would be expected to use their memory for the prechange feature and consequently sort the changed brick by the old value.

Sorting performance when subjects missed a change in the put-down-relevant feature is plotted in Figure 8. Bricks with changes relevant to sorting were most commonly sorted on the basis of their prechange feature. 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 both the *One Feature* and *Two Feature* conditions, 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 is consistent with the observation in normal trials in which subjects most frequently fixate the belt area after the put-down cue rather than refixating the brick in hand. We interpreted that result as evidence that subjects use working memory for the sorting decision rather than acquiring the relevant feature at the last possible moment. Because subjects are using the prechange feature value, the failure to detect the change seems to be caused

by a failure to update the new feature and not from either failure to remember the initial value or failure to compare the two values.

However, there were differences between the *One Feature* and *Two Feature* conditions. Sorting by the new feature in the *One Feature* task occurred in only 8% of missed put-down-relevant change trials (6 of 78), whereas subjects sorted with respect to the new feature in the *Two Feature* task in 24% of trials (16 of 66; chi-squared,  $p < .05$ ). 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. Thus, when a second feature was not needed until several hundred milliseconds later, after pick-up, there was some tendency to delay the acquisition of this feature until the sorting cue had been inspected. This is also consistent with normal trials, in which subjects were more likely in the *Two Feature* than in the *One Feature* condition to refixate the brick in hand before guiding it to the belt (see Table 1). However, performance in the change detection task did not reveal this tendency, which would have resulted in infrequent detection of the put-down-relevant feature change. Instead, increased rates of change detection for the put-down feature suggested early acquisition, as discussed above (see Figure 5). Thus, the analysis of fixation strategies and sorting behavior may provide a more complete description of the underlying strategies not captured in the explicit change detection task.

One possible reason for subjects' tendency to sort bricks by the old feature is a failure to fixate the brick after the change. As shown in Table 1, subjects looked back to the brick after the put-down cue, but before put-down on the conveyor belt, in only 19%–29% of trials. Subjects also looked at the brick while placing it on the belt, as shown by second peak in Figure 4C. We examined total fixation duration on bricks before and after a feature change, regardless of where the brick in hand was in relation to the belt area. Before a feature change, subjects fixated the brick an average of 1,144 ms. Following a feature change, fixations to a brick were recorded on 78% of change trials. On trials in which one or more brick fixations were recorded following a change, the average total duration of brick fixations following the change was 955 ms. Thus, in the majority of trials, subjects had considerable time to evaluate the quality of the new put-down-relevant feature, even though they were not using this information for their sorting decision (see Figure 8).

These measurements of fixation duration on bricks following a change are a conservative estimate. The state of the feature was often available to the subject even during times at which no fixation was recorded by the automated analysis. Because precise placement on the belts is required for the brick to be carried away, many subjects fixate the edge of the belts during put-down and bring the brick very close to the fixation location. Second, brick fixation may not have been detected by the program on some trials because track loss was most common when subjects' eyes were directed down to the belts, rendering a precise measure of fixation location impossible. Thus, although 22% of trials had no recorded fixation on the brick following the change, inspection of the video records revealed that the changed feature was almost always close to fixation.

Does sorting by the new feature require some additional computation by the subject? This additional computation may be revealed by some aspect of behavior that differentiates trials in which subjects sort the changed brick by the new feature. We

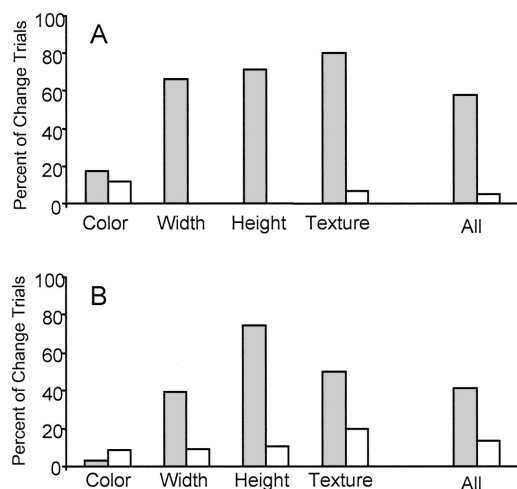


Figure 8. Percentage of put-down-relevant change trials sorted by the old (gray) versus the new (white) feature in the (A) *One Feature* task and (B) *Two Feature* task. Percentages on the y-axis represent all change trials, including detected changes (not shown).



looked first at the duration of the hand movement during the sorting phase of the experiment. When subjects sorted by the new feature, the hand movement during sorting lasted 972 ms longer in the *One Feature* task and 1,803 ms longer in the *Two Feature* task than when they sorted by the old feature ( $t$  test,  $p < .01$ ). Because there are only a small number of trials when subjects sort by the new feature, we examined the video records. It appears that these rather large increases in hand movement duration are a consequence of a correction or alteration in the planned movement (e.g., a double-take). In a similar fashion, we examined the duration of fixations on the bricks following the change. For both the *One Feature* and *Two Feature* tasks, trials in which subjects sorted the brick by the new feature were longer in duration following the change than trials in which subjects sorted the brick by the old feature. However, the differences were not statistically significant, possibly as a consequence of the small number of trials (6 and 15 in the *One Feature* and *Two Feature* tasks, respectively). Thus it appears that it is most efficient for subjects to sort by the old feature and that sorting by the new feature requires information typically not acquired.

### 2.5 Implicit Effects of Changes

Eye movements during normal trials (see Section 1.3), performance in the explicit change detection task (see Section 2.1), and sorting behavior for changed bricks (see Section 2.4) provide converging evidence that subjects are selectively acquiring task-

relevant object features for the sake of sorting the bricks, rather than storing integrated object representations. Another potential source of evidence is implicit behavior following a change. Subtle differences in behavior, such as fixation duration, are often interpreted as evidence for residual representations of the prechange feature (Fernandez-Duque & Thornton, 2003; Hollingworth & Henderson, 2002; Mitroff, Simons, & Franconeri, 2002; Thornton & Fernandez-Duque, 2000, 2002). Given that subjects had such low rates of explicit change detection in the present experiment, we considered implicit hand and eye behavior as another candidate for revealing what information subjects were representing during the task.

**Hand movements.** Hand movements during the sorting task were stereotypical across subjects and trials (see Figure 3). After subjects carried a brick toward themselves, their hands often paused briefly while they evaluated the put-down cue, and they occasionally refixated the brick before putting the brick down on either conveyor belt. Despite this consistent pattern, there was variability in the duration of the sorting and put-down stage of the hand movement. We hypothesized that cognitive processes contributing to variability in hand movement duration would be sensitive to feature changes despite a failure to explicitly report the change. Thus, we examined hand movement durations for normal and missed trials in each experimental condition.

Figure 9 shows the cumulative distributions of hand movement duration during sorting and put-down for all normal and missed

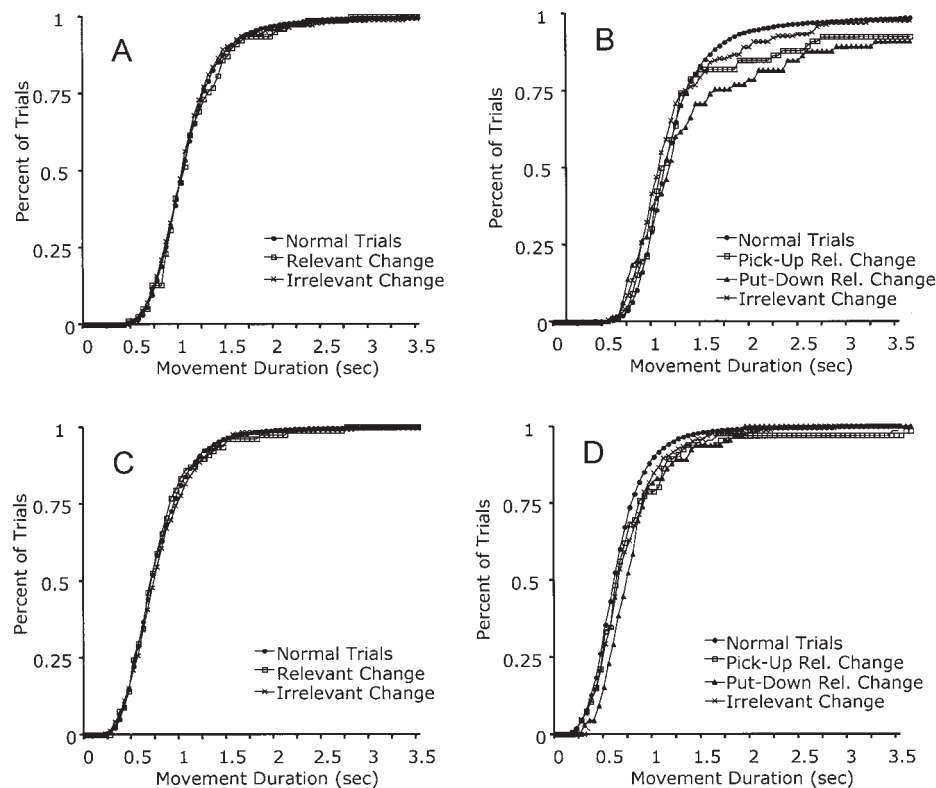


Figure 9. Duration of hand movement during sorting and put-down for all normal and missed change trials during the *One Feature* task (A and C) and *Two Feature* task (B and D) during sorting (A and B) and put-down (C and D). sec = seconds; Rel. = relevant.

change trials for both experimental conditions. In the *One Feature* task, there was no difference in the distribution or mean movement duration between normal trials and trials with missed changes during either sorting or put-down. However, in the *Two Feature* task, there were differences in the mean movement duration. This difference was most clear for changes relevant to put-down during the sorting movement in the *Two Feature* condition (see Figure 9B). For example, 90% of normal trials have hand movements less than 1.55 s, whereas this movement duration accounts for only 75% of trials with missed put-down-relevant changes (see the triangles in Figure 9B). Within-subject *t* tests on the mean hand movement duration revealed that only put-down changes were significantly longer than normal trials (normal trials: 1,129 ms; pick-up trials: 1,289 ms; put-down trials: 1,419 ms; irrelevant trials: 1,163 ms; one-tailed *t* test,  $p = .01$ ). Put-down change trials were also longer than irrelevant change trials ( $p < .05$ ). Figure 9 reveals that the difference in mean hand movement duration in the *Two Feature* task derives primarily from a subset of trials. The distribution of movement duration in each of the trial types is identical in the shortest three quarters of trials in each category. However, there appears to be a subset of trials for which hand movement duration is substantially longer. Trials with long hand movement duration following a missed change to a put-down-relevant feature were frequently the same trials in which the changed brick was sorted by the new feature (see Section 2.4).

Though the primary effect of changes appears to be in the sorting phase of the experiment, there are also small differences during the put-down phase (see Figure 9D). Movement duration was longer for missed pick-up or irrelevant changes than for normal trials (normal trials: 620 ms; pick-up trials: 794 ms; put-down trials: 638 ms; irrelevant trials: 703 ms; one-tailed *t* test,  $p < .01$ ). Unlike the sorting phase (see Figure 9B), in the put-down phase, there was no difference between the three change conditions (pick-up, put-down, or irrelevant changes).

Thus, in the *Two Feature* condition, hand movement duration revealed sensitivity to feature changes despite subjects' failure to explicitly report the change. These longer hand pauses for changes were clearest during the sorting phase of the trial and when the change was relevant to the sorting decision.

*Eye movements.* A similar analysis was performed on the duration of fixations on the brick following the change. Because the exact position of the brick at the time of a change varied across trials, we required a common criterion of brick position with which to segment both normal and change trials. Thus, brick fixations were analyzed after the brick had been carried halfway between the array and the conveyor belts. Using the halfway point, rather than the precise time of the change, allowed comparison of fixation duration on normal trials, where no change occurred with that on change trials. Brick fixations after being carried past the halfway point during change trials overwhelmingly occurred following the feature change. Note that these fixations occurred during the sorting phase and act of put-down. Because of track loss, there were fewer subjects and change trials with adequate eye position data than trials used in hand movement analysis.

In the *One Feature* task, there was no significant difference in average brick fixation duration between normal trials and trials with a missed change. In the *Two Feature* task, bricks with missed changes were fixated for longer than bricks without a change (1,029 ms and 933 ms, respectively, one-tailed *t* test,  $p < .05$ ). Of

the three kinds of change trials (pick-up, put-down, irrelevant), only put-down-relevant change trials were significantly longer (put-down: 1,148 ms; normal: 933 ms; one-tailed *t* test,  $p < .05$ ). Note that the number of trials in these comparisons is quite small (pick-up:  $n = 27$ ; put-down:  $n = 26$ ; irrelevant:  $n = 79$ ). Thus, missed changes had a similar effect on both hand movement and fixation duration: no effect in the *One Feature* task but an effect following put-down-relevant changes in the *Two Feature* task.

### 3. Effect of Trial History

The analysis in Section 2 suggests that the immediate task demands influence the acquisition of visual information and its representation in visual working memory. These effects were revealed both in change detection performance and in hand and eye movements. We next considered the history of events across trials as another potential influence on the deployment of attention and consequent selection of information. We examined two different kinds of effect of previous trials. One was whether the detection of a change influenced subjects' eye and hand movements on subsequent trials. The other was whether repetition of the same pick-up or put-down rule influenced hand and eye movements.

#### 3.1 History of Change Detection

*Duration of hand movement and brick fixation.* Because subjects typically sort by the old feature, this suggests that the way they sample information from the scene depends on their expectation of stability of the object features (despite having been told to expect changes). After a trial in which a change is detected, subjects may revise their expectations, and this may be revealed in hand and eye movements. For each subject in each experimental condition, we isolated the five trials before and after a trial with a detected change. Changes in the duration of hand movements and brick fixations following a change trial, for both sorting and put-down, are plotted in Figure 10.<sup>2</sup> Figure 10A shows that the duration of the hand movement during sorting, in both *One Feature* and *Two Feature* tasks, increased by approximately 400 ms following a trial with a detected change. However, the strength of this effect quickly dissipated, with only modest effects in subsequent trials. Note that the sorting phase of the trial had the largest increase in movement duration. Movement during put-down had a fixed, and more modest, increase in duration across the five trials

<sup>2</sup> Hand movements were categorized as sorting or put-down as described in Section 1.2. Fixations were categorized in a comparable manner. However, in addition to fixations to the brick while it was being carried to the put-down cue, we also included fixations to this same brick while it was still in the array just prior to pick-up. We included these fixations because this is a critical time for feature acquisition. Fixations to the brick when the brick occluded the belt area were classified as put-down fixations. To identify possible changes in behavior following a change, and to reduce between-subjects variability, we do not show the mean duration across subjects. Instead, durations during the five trials prior to a change trial were averaged together and used as a baseline. The values in the figure are deviations from this baseline. These deviations were averaged within each subject across corresponding trials and then across subjects. Subjects occasionally made a gross change in their hand movement during a trial, sometimes correcting a potentially wrong sorting decision or picking up a

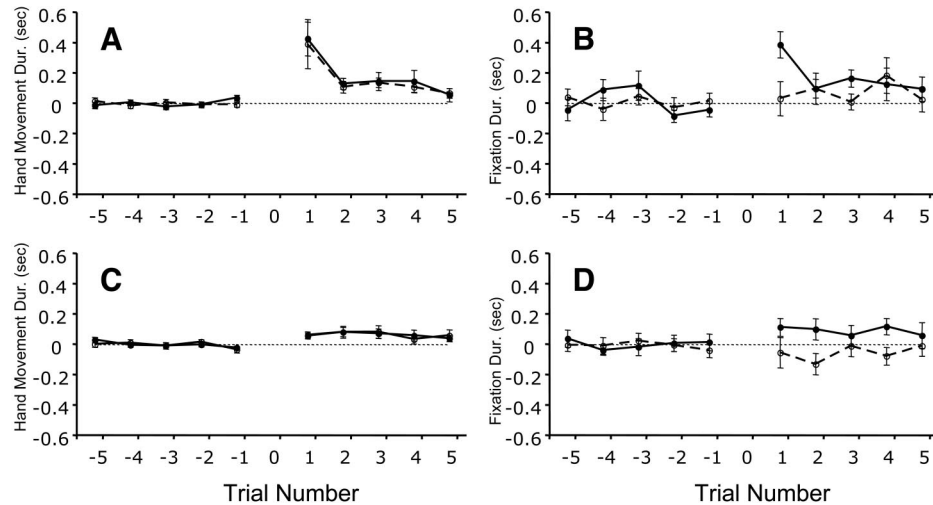


Figure 10. Changes in hand movement duration (A and C) and brick fixation duration (B and D) in trials following a detected change. Movements are categorized as occurring during sorting (A and B) or put-down (C and D). Dashed line represents *One Feature* condition; solid line represents *Two Feature* condition. Error bars represent standard error across subject mean. Dur. = duration; sec = seconds.

following a noticed change (see Figure 10C). When averaged across all five trials after the change, hand movements during both sorting and put-down in both tasks were significantly longer than in the preceding five trials ( $p \leq .001$ , one-tailed paired  $t$  test).

A similar 400-ms increase was observed in brick fixation duration in the first trial after a detected change but only in the *Two Feature* task (see Figure 10B). The five trials following a noticed change had significantly longer fixation duration during sorting in the *Two Feature* task ( $p < .01$ ), and approached significance for the *One Feature* task ( $p = .06$ ), despite not having a punctate effect in the first postchange trial. Only the *Two Feature* task had an increase in brick fixation duration during put-down in trials following a noticed change ( $p < .05$ ; see Figure 10D).

Thus, both hand and eye behavior differed in trials following a noticed change. This difference in behavior was characterized as longer hand movement duration and longer total fixation to the brick, particularly in the trial immediately following a noticed change, and particularly during the sorting phase of the trial. The significance of these results is addressed in the Discussion section.

**Fixation sequence.** In addition to fixating the brick for a longer duration following a trial with a noticed change, subjects may also modify the order in which they sample information from the brick. We examined the frequency of refixations on the brick after fixating the put-down cue (see Table 1), but before placement on the conveyor belt, in the five trials preceding and the five trials following a noticed change. In the *One Feature* condition, the probability of refixating the brick following a noticed change (35.2%) was not significantly different than in the five trials preceding the change (31.0%). In the *Two Feature* condition, subjects were more likely to refixate the brick in the five trials following a noticed change (43.6%) than in the five trials preceding the change (35.0%; paired  $t$  test,  $p < .05$ ). This finding indicates that subjects modify their behavior following a noticed change. This modification in behavior may be an attempt to increase the probability of detecting changes in the future, as discussed in Section 2.3 and shown in Figure 7.

Note that the average probability of refixating the brick is higher in this analysis than in Table 1. This difference may have been caused by two factors. First, a smaller number of trials were considered in the analysis for Figure 11. Second, the analysis for Figure 11 includes only subjects who detected changes. Section 2.3 provides evidence that subjects who detected changes had a different strategy involving more refixations to the brick before setting it down on the conveyor belt. Analyzing fixations from only these subjects may have increased the average frequency of refixations.

### 3.2 Repetition of Pick-Up and Put-Down Cues

Repetition of the stimuli or sorting rules across trials may also influence deployment of attention. Therefore, the second way in which we examined the effect of trial history was with respect to the pick-up and put-down sorting rules. For each subject we isolated trials with the first, second, and third consecutive presentations of a particular pick-up cue and/or put-down cue. We examined the influence of this history of task instruction on both fixation strategies and hand movements.

We looked at three kinds of repetition: repetition of the same pick-up instruction (e.g., “pick-up red”), the same put-down rule (e.g., “wide on right/thin on left”), or both (e.g., “pick-up red, wide on right/thin on left”). Figure 12 plots the effect of cue repetition on the duration of hand movement during sorting and on the probability of refixating the brick in hand. Figure 12A plots the duration of hand movement during sorting, as a function of trial number, when the put-down cue was repeated. On Trial 0, the pick-up cue differs from the previous trial. On Trials 1 and 2, the pick-up cue is the same as in the previous trials. Thus, Trials 1 and 2 reveal the effect of the cue repetition. Repetition of the put-down

dropped brick. These deviant trials were excluded from the analysis and accounted for less than 6% of trials.






Figure 11. Probability of refixating the brick in hand following a fixation to the put-down cue in the five trials before and after a trial in which a change was noticed. In the *Two Feature* condition, subjects were more likely to refixate the brick in hand in trials following a noticed change ( $p < .05$ ).

cue resulted in a small decrease in the duration of hand movements. Trials with repetition of both the pick-up and put-down cue also resulted in a decrease in movement duration for the *Two Feature* task (see Figure 12C). Because there are so few trials with two repetitions of both pick-up and put-down cue, only a single repeated trial is presented. We also analyzed movement duration during the put-down phase, following repetition of the put-down cue, but there was no effect. Repetition of the pick-up cue had no effect on movement duration during sorting or put-down (not

shown). Thus, the main finding for hand movements is that repeated presentations of the same put-down rule, not the same pick-up rule, shorten the duration of hand movements during the sorting phase. (See the caption of Figure 12 for statistically significant comparisons.) Note that a repetition of the put-down rule does not simply represent a repetition of the movement. For example, a subject selecting a wide brick on one trial and a thin brick on the following trial will make different movements even though the put-down rule (e.g., “wide on left/thin on right”) is the same. During the course of a single trial, the subject must remember not just brick features but rules guiding behavior. The reduction in hand movement duration may reflect the storage of both features and the put-down rule in a manner consistent with action binding in *event files* (Hommel, 2004). Repetition of the put-down cue may also facilitate the interpretation of this instruction, and thus reduce the duration of the hand movement.

Figure 12B plots the probability of refixating the brick in hand, as a function of trial number when the pick-up cue was repeated. In both the *One Feature* and *Two Feature* tasks, there is a small but significant reduction in the probability of a refixation for both the first and second repeated trials. In contrast to the hand movements, there was no influence of put-down cue repetition on refixations (not shown). However, when both the pick-up and put-down cues were repeated, there was again a significant effect. This is shown in Figure 12D. If, as discussed above, refixations are less frequent when the information is held in working memory, the decreased

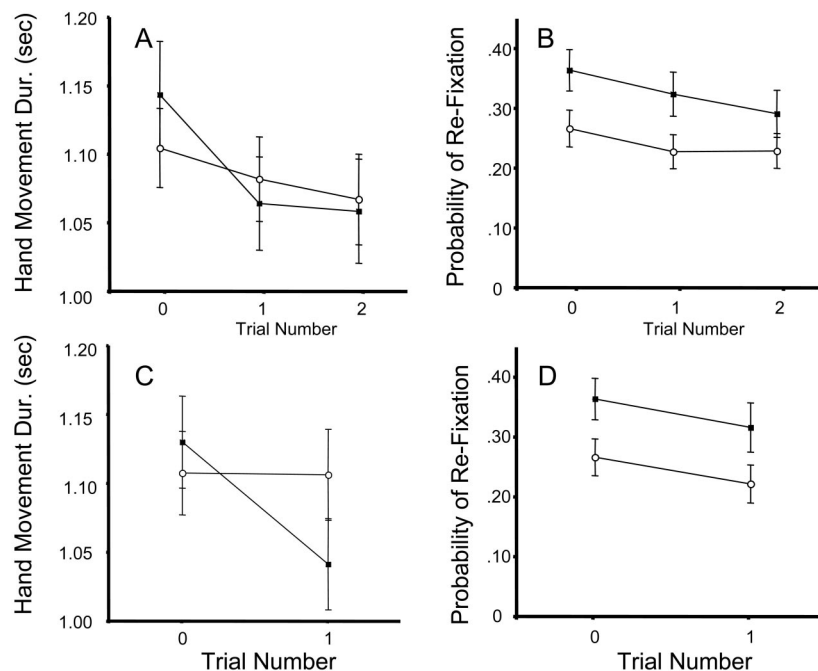


Figure 12. Effect of cue repetition on duration of hand movement and probability of brick refixation: Repetition of (A) put-down cue, (B) pick-up cue, and (C and D) both pick-up and put-down cue during *One Feature* (open circles) and *Two Feature* (solid squares) conditions. Error bars represent standard error between subjects. Significant differences (one-tailed  $t$  test,  $p < .05$ ) are as follows. A: *One Feature*, Trial 0 > Trial 1, Trial 0 > Trial 2; *Two Feature*, Trial 0 > Trial 1, Trial 0 > Trial 2; B: *One Feature*, Trial 0 > Trial 1; *Two Feature*, Trial 0 > Trial 1, Trial 1 > Trial 2, Trial 0 > Trial 2; C: *Two Feature*, Trial 0 > Trial 1; D: *One Feature*, Trial 0 > Trial 1; *Two Feature*, Trial 0 > Trial 1.

probability of refixating the brick with repetition of the pick-up instruction may indicate that history of the pick-up task (and consequent history of picking up a brick with the same feature value) across trials influences the storage of brick features in working memory. These results are considered more thoroughly in the Discussion section.

## Discussion

### *Synopsis*

In this study, we sought to elucidate the relationship between task goals, attentional selection, and working memory for objects and their features. Subjects performed a primary task of sorting bricks by their features and a secondary change detection task for features both relevant and irrelevant to sorting. This secondary change detection task was used as a probe to determine what information subjects were using in the primary sorting task. To summarize the main findings: Changes to features were detected more often when the feature was relevant to the sorting task than when the feature was irrelevant. The stimulus dimension of the changed feature also influenced change detection. Rates of change detection were generally quite low, regardless of task relevance. We tested possible causes of change blindness by examining sorting behavior for missed changes. Preference to sort changed bricks by the prechange feature suggests that subjects fail to acquire the new feature after the change. It also suggests that subjects used working memory for the put-down decision and not the feature value that was present at or near fixation at the time of put-down. Gaze and hand movements provided a more subtle assessment of the time course of acquisition and storage of visual information. Although the most common strategy is to acquire both pick-up and put-down features at the time of pick-up, subjects are more likely to delay acquisition of the put-down feature in the *Two Feature* task than in the *One Feature* task (see Table 1 and Section 2.4). Despite subjects' failure to explicitly report a change, a fraction of trials revealed implicit differences in subjects' hand and eye behavior (see Section 2.5). Subjects' behavior was sensitive to the history of events across trials. Subjects increased the duration of hand movements and brick fixations in trials following a noticed change (see Section 3.1). Repetition of pick-up and put-down cues across trials decreased both hand movement duration and the probability of refixating the brick (see Section 3.2).

### *Role of Task in Acquisition and Working Memory for Object Features*

The present experiment has a number of important implications for our understanding of both the contents and the persistence of information in visual working memory during ongoing behavior. The first implication concerns the nature of the information about an object that is captured during fixations. The notion of object files, an internal representation of all of the bound features of an object, has influenced much work on visual working memory (Hollingworth & Henderson, 2002; Irwin & Zelinsky, 2002; Kahneman & Treisman, 1984; Kahneman et al., 1992; Saiki, 2003; Treisman, 1988; Vogel et al., 2001; Wheeler & Treisman, 2002). However, the issue of the generality of object file representations has not been addressed. For the purpose of many memory tasks,

integrated object files may be required. This may not, however, be an invariant feature of working memory. The finding that subjects were more sensitive to a feature change when the feature was relevant to the sorting tasks argues against such general purpose internal representations. Instead, subjects seem to preferentially acquire specific feature information demanded by the task. Thus, for example, in an action like picking up an object, the visual computations and contents of working memory may be very different than the computations performed when recognizing an object or remembering it as part of a scene. We suggest that a multiplicity of different computations are required for normal functioning and that each of these computations requires different information depending on the specific goals of the subject. Thus, in this sense, all perception is goal driven.

Other paradigms have also demonstrated separate or enhanced processing or memory for attended features. For example, cortical areas with neurons that are selective for a particular feature have enhanced activity or selectivity for the feature when this feature is attended to (Desimone & Duncan, 1995; McAdams & Maunsell, 2000; Motter, 1994; Treue & Maunsell, 1999). There is also recent evidence to suggest that object features may be more selectively stored than previously thought (Wheeler & Treisman, 2002). The present findings show that similar specificity of object feature representations extends to dynamic natural tasks. Thus, integrated object representations are not an inevitable consequence of an eye fixation during ongoing behavior.

It has also been proposed that object files are automatically updated when subsequent fixations fall on the same object (Hollingworth & Henderson, 2002; Kahneman et al., 1992). The fact that subjects sorted the changed bricks by the prechange feature argues against an automatic updating process. Subjects instead preferred to use the contents of working memory for the sorting decision rather than the information available in foveal vision for three quarters of a second. Rates of change detection were generally low, a surprising result considering that subjects had no difficulty discriminating the features for the sake of the sorting task. Thus, the information acquired within a fixation may not relate in any simple way to information present in the local raw image. Rather, it is intimately connected to the ongoing task computations. For example, in Yarbus's (1967) classic experiments, fixations to the same object in response to different instructions may have reflected the acquisition of quite different information. This selective acquisition may be reflected in even low-level cortical areas with neural activity that depends not only on stimulus features but on task context (Ito & Gilbert, 1999; Roelfsema, Lamme, & Spekreijse, 1998).

The current findings strengthen previous evidence that task demands restrict acquisition of information to relevant object features. Much of this evidence also derives from paradigms in which subjects were free to organize their own behavior while engaged in visuomotor tasks. For example, in a block-copying task, subjects preferred to refixate an area rather than store this information internally (Ballard et al., 1995). These refixations during different stages of a copying sequence were hypothesized to be evidence for acquisition of individual object features relevant to the immediate needs of the subject. This hypothesis was strengthened with the finding that fixation duration on a changed block in the model depended on what information (color or spatial position) was required from the block for the immediate demands of the

copying task (Hayhoe et al., 1998). In a virtual brick sorting paradigm similar to the present experiment, the rate of detection of unexpected changes to brick height was associated with the use of height for the decision of pick-up and put-down (Triesch et al., 2003). However, the acquisition of irrelevant object features could not be determined because only the relevant feature was changed. The present finding, that changes to an object feature are more frequently detected when that feature is relevant for sorting, provides more substantive evidence for task-specific acquisition of object features.

Although the specific requirements of the task influence the acquisition and storage of visual information in working memory, the relationship is complex. Ballard et al. (1995, 1997) suggested that information is acquired just in time and not held in working memory unless needed. O'Regan and colleagues have also suggested that the external world is used as a repository for information, acquired upon demand (O'Regan, 1992; O'Regan & Noë, 2001). However, a strict interpretation of this idea is inconsistent with some aspects of the present findings. In the *One Feature* task, a memoryless visual system would have not only missed the feature changes but also used the information immediately available in the scene and consequently sorted changed bricks by their new feature. The preference to sort changed bricks by the old feature suggests that operations performed during subtasks may include the use of information acquired from previous fixations. One might argue that it makes sense to store the relevant feature in the *One Feature* task because it was required for the put-down decision. However, detection of changes in the pick-up feature was just as high in the *Two Feature* task, even though the pick-up feature was no longer needed. This suggests that information is not immediately discarded when it is no longer relevant to the immediate demands of the subtask.

Performance in the *Two Feature* task also suggests that subjects sometimes acquire information ahead of time, before it is absolutely necessary. Sensitivity to changes in the put-down feature was greater than to irrelevant features, even though the change occurred before it was immediately needed. This suggests that subjects were acquiring both pick-up and put-down-relevant features prior to the change. This is supported by the observations that subjects typically looked directly to the belt after fixating the put-down cue. However, subjects were more likely to refixate the brick before put-down in the *Two Feature* task than in the *One Feature* task, suggesting that in some cases, subjects delayed the acquisition of the put-down-relevant feature until they had acquired the put-down cue. The strategy of acquiring the put-down-relevant feature prior to the change may have been a consequence of the blocked trial design. Subjects could anticipate which feature was relevant for the upcoming put-down decision, and thus it may have been more efficient for subjects to acquire and store this information during pick-up rather than making an extra fixation later in the task. This argues that visual information extracted from an object may not be isolated to the immediate microtask. Instead, visual information that is expected to be relevant in the future may also be acquired and stored. A strategy of acquiring information expected to be relevant in the future may also be the source of look-ahead fixations in natural tasks (Hayhoe et al., 2003; Land et al., 1999; Pelz et al., 2001; Pelz & Canosa, 2001). Future work in our lab will address how expectations of task demand influence what information is selected for storage in working memory.

It is important to note that although the requirements in the sorting task are an important modulator of performance in the change detection task, the quality of the stimulus also plays a role. For example, changes in color were more frequently detected than changes in texture. We did not attempt to equate feature salience in this experiment, other than making sure that different feature states were easily discriminable. Whether this implies that bottom-up factors influence what information is acquired and stored, or whether the magnitude of the change is more intrinsically salient for some features than others, is unclear.

It is also important to recognize the possibility that subjects may have relied not entirely on visual working memory but may have also used verbal encoding and rehearsal to sort the bricks or to detect changes. Subjects in the present experiment were not required to perform a simultaneous verbal rehearsal task, as performed in other paradigms measuring the capacity of visual working memory (Luck & Vogel, 1997; Wheeler & Treisman, 2002). The traditional purpose of simultaneous verbal rehearsal tasks is to suppress the use of verbal strategies during tasks designed to isolate the use of visual memory. However, the sorting and change detection tasks in the present experiment seem to have prompted subjects to store information on only one or two features throughout a single trial, presumably well below the capacity of both visual working memory and verbal rehearsal. Nevertheless, it remains possible that subjects chose not to depend on their visual working memory to store this modest information but used verbal strategies as well. It is not clear how verbal rehearsal would influence performance in the sorting paradigm. If verbal strategies were used in the present experiment, this would be of interest in its own right. Any effort to use verbal strategies would suggest that while the capacity of visual working memory may include many features or objects, the brain may minimize its usage in the context of ongoing tasks.

### *Why Are Changes So Rarely Noticed?*

Overall rates of change detection were quite poor. This is surprising considering that the brick was the focus of attention throughout the trial and was fixated for close to a second before and after the change. Change blindness is often interpreted as evidence for limited memory representations from prior fixations (Blackmore, Brelstaff, Nelson, & Troscianko, 1995; Grimes, 1996; Irwin, 1991; O'Regan, 1992). However, subjects' tendency to sort changed bricks by the prechange feature shows that change detection was not limited by memory. Another suggestion is that internal representations are limited to objects of central interest that attract attention (Rensink et al., 1997). However, in the present experiment, it is easy to argue that the brick in the subject's hand is the object of central interest, yet change detection was still less than 40%. Hollingworth and Henderson (2002) have argued that change detection may require a fixation on the object both before and after the change. However, subjects in the present experiment met this requirement, fixating the brick for approximately 1 s before and after the change, but still failed to detect most changes. Thus, although change detection may necessitate both a memory representation of the prechange object and a fixation on the object following a change, these events are still not sufficient. Because working memory representations were the dominant source of information for the sorting decision despite fixations on the new



feature, we focus on the question: Why did subjects not notice the new feature?

One possible explanation is that the computations involved in sorting the brick and guiding it to the belt placed such tight constraints on visual processing that the new feature value was not encoded. For example, once the sorting decision had been made (using working memory), the only visual information necessary for put-down is the spatial position of the brick with respect to the conveyor belt. If subjects were acquiring only spatial information, they may be blind to any of the intrinsic features of the brick and thus miss the change. Several bodies of literature support the hypothesis that top-down signals enhance the processing of relevant, and suppress the processing of irrelevant, objects and their features (Desimone & Duncan, 1995; Gottlieb et al., 1998). Top-down signals may also underlie the phenomenon of inattention blindness, the failure to process salient visual objects in the context of a competing task set (Mack & Rock, 1996; Most, Scholl, Clifford, & Simons, in press; Simons & Chabris, 1999). Failure to process the changed feature in the present results may reflect the same underlying limitations. Processing limitations revealed in the attentional blink phenomenon might also be related to the present failure to notice changes. The attentional blink is thought to reflect the time required for the observer to modify his or her task set for a new target when targets are presented in quick succession in a rapid serial visual presentation (RSVP) task (Chun & Potter, 1995). As subjects in the present experiment are presumably attending to the put-down cue and sorting decision, they may not switch their task set to one appropriate for analyzing brick features for change detection. Stimuli not recognized during the attentional blink also fail to stimulate activity in frontal cortex (Marois et al., 2004), an area thought to be participating in change detection (Beck, Rees, Frith, & Lavie, 2001). Note that subjects may organize their behavior to modulate effects of attentional limitations. After a change was noticed, subjects increased the duration of their hand movements during sorting; the probability of refixating the brick before put-down also increased (see Figures 10 and 11). Higher rates of refixation in general were associated with higher rates of change detection (see Figure 7).

Another consideration in understanding why subjects fail to update their representation of the changed brick feature is that, although a red to blue change is a very salient difference for a single brick, both red and blue bricks were always present within a trial. Similarly, the only brick colors present across all trials were red and blue. Despite the intrinsic salience of a feature change within a brick, both of these color values were used exclusively throughout the experiment, and this may have impaired detection performance.

In principle, it is possible that subjects did successfully encode both the pre- and postchange features but failed to compare the two. Such failure has been attributed to change blindness in the real world (Levin, Simons, Angelone, & Chabris, 2002; Simons & Levin, 1998) and other explicit change detection paradigms (Mitroff et al., 2002). Subjects' tendency to sort by the prechange feature suggests that the postchange feature was not encoded and consequently could not be compared. However, it is possible that once the sorting decision had been made, the new feature was encoded, but subjects forgot the basis of their sorting decision and then failed to compare the two brick representations. Alternatively, failure to compare pre- and postchange values may be a conse-

quence of subjects' ignorance that changes might occur. In the present experiment, subjects were told three times to expect changes, although it is possible that they did not store this expectation throughout the entire experiment. Thus, although we cannot rule out the possibility that subjects failed to compare the pre- and postchange features, this account seems unlikely to explain failure to detect the majority of changes. The small probability of sorting by new suggests that subjects encoded the new feature value only rarely. This new feature information is a prerequisite for a successful comparison, so it seems unlikely that failure to compare accounts for a significant number of trials.

### *Implicit Behavior in Missed Trials*

Subjects' hand movements and fixations to the brick were often longer in duration when a change was made to a brick feature, despite subjects' failure to explicitly report the change (see Figure 9). Other change blindness paradigms have also found longer fixation duration on changed objects despite failure to report the change (Hayhoe et al., 1998; Hollingworth & Henderson, 2002). These longer fixations on changed stimuli have been cited as evidence for residual representations of the prechange feature.<sup>3</sup> The results of the present experiment demonstrate that changes in hand movements, as well as changes in fixation durations, can accompany missed visual changes.

The finding that changes in brick features increased the duration of both hand movements and brick fixations has consequences for our understanding of what information is used in sensorimotor transformations, serving the coordination of simple behaviors, and what information is used in more abstract cognitive processes, serving decision making and explicit change detection. Ballard et al. (1997) have suggested that visual computation can be considered as a hierarchy of processes that use increasing levels of abstraction and operate at proportionately longer time scales (see also Newell, 1990). For example, elementary visual computations, such as parallel search, may require 50–100 ms. Primitive physical acts, such as an eye movement, at the embodiment level, require 200–300 ms. Sequential operations, composed of these primitive acts, at the cognitive level, require a few seconds, comparable to the time constant of working memory. Using the put-down cue to guide the sorting decision can be considered a process at the cognitive level, because it requires a sequence of operations involving acquiring the put-down cue, translating the cue into a sorting rule, and applying this knowledge to information on the brick features, which, in turn, may be acquired with an eye movement or accessed from working memory. Explicit change detection can also be considered a process at the cognitive level, because it requires a sequence of operations involving acquisition, storage, and comparison of visual information. However, note that these two cognitive operations may impose different demands on what visual information is acquired and stored and how this

<sup>3</sup> Determining whether modifications in behavior following a visual change constitute implicit change detection is a difficult, if not controversial, process (Fernandez-Duque & Thornton, 2003; Mitroff et al., 2002). For example, it is reasonable to argue that, in trials with longer hand movements or brick fixations, subjects did in fact consciously perceive the change but failed to report it because they lacked confidence in their detection.

information is used. This may explain why, in the present experiment, we found longer hand movements on missed change trials, despite failure to report the change. Note that this increase in duration occurred mainly in the subset of trials in which the subjects sorted the brick by the new feature (see Section 2.4). This means that subjects were making a cognitive decision using the new information but still failed to report the change, perhaps because of a failure to perform a comparison operation, relating the new feature in the scene to the old feature stored in working memory. Thus, the operations involved in guiding the hand during sorting are not necessarily the same as those used in change detection. Lastly, the fact that irrelevant changes had only a small effect on implicit behavior, and that changes relevant to sorting had a substantially larger effect (see Section 2.5), strengthens the argument that acquisition of visual information and the access of this information to abstract cognitive processes is highly selective.

### *Modifying Behavior Strategies to Detect Changes*

The argument that change detection requires cognitive processes typically not used by other tasks, such as sorting bricks, has consequences for how results from other change detection paradigms are interpreted. Change detection tasks typically require subjects to view an image of a scene, or an array of stimuli, and to report any changes. Performance in the change detection task is used to evaluate the quality or capacity of visual processing during viewing. We applied a similar strategy in the present experiment: Subjects were asked to perform a change detection task concurrent to the task of brick sorting. We then used rates of change detection to infer what information was used during brick sorting. The implicit assumption behind this and other change detection experiments is that visual computations used in change detection are representative of visual computations used during ordinary behavior. However, hand and eye movements in the present experiment argue against this assumption. Following a noticed change trial, subjects dramatically altered their behavior (see Section 3.1). In the trial immediately following a noticed change, the duration of both hand movement and brick fixation increased by approximately 400 ms. This strongly suggests that subjects are making large changes in their behavior, and thus the underlying internal operations, when prompted to prioritize the change detection task. However, this radical change in behavior was short-lived. Hand and eye movements quickly returned close to normal in subsequent trials, although a slight change in behavior was found up to five trials following a noticed change. This quick return toward baseline behavior may have been caused by subjects' reprioritizing the sorting task. Following a noticed change, subjects did not radically alter their behavioral strategy throughout the remainder of the experiment. Although it is well within the capacity of visual attention and working memory to detect a change, the challenge is in prioritizing the acquisition, storage, and comparison of this information. This conflict is also made apparent in real-world change detection paradigms when subjects are not told to expect changes. Subjects are quite poor at detecting an exchange of strangers in the middle of a conversation (Levin et al., 2002; Simons & Levin, 1998). Presumably, these same subjects would be more likely to identify the change if they had been prompted to prepare for the change detection task and consequently prioritize storing information in working memory and perform frequent

comparisons. Thus, it is becoming less clear how conclusions drawn from tasks that explicitly instruct subjects to detect changes apply to an understanding of how visual information is used in more natural behavior.

### *Influence of Prior Trials on Implicit Behavior*

Subtle changes in subjects' behavior across trials revealed sensitivity to cue repetition. Repetition of the same put-down cue reduced hand movement duration during the sorting decision. Repetition of the same pick-up cue, and thus repetitions of picking up a brick with the same feature value, reduced the probability of refixating the brick before the sorting decision. Repetition of both the pick-up and put-down cue reduced both the hand movement duration and the probability of refixations. These results may be related to the time cost of task-switching (Birnbom, 2003; Fagot, 1994; Pashler, Johnston, & Ruthruff, 2001; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977) where repeated trials with the same instruction, but not necessarily the same movement, result in progressively more efficient movement. It is not clear if this change in behavior is due to memory for the task or memory for the object features. Other paradigms have found that visual search within repeated spatial context (Chun & Jiang, 1998), or search for pop-out features (Maljkovic & Nakayama, 1994, 1996) is facilitated with recent repeated exposure to similar stimuli. Note that in those paradigms, the structure of the trial minimizes the number of cognitive processes between repetitions of stimuli. In contrast, trial structure in the present paradigm demands the use of many operations and subtasks between exposure to repeated cues. This suggests that memory for the cue in previous trials may be fairly robust. The fact that an effect should arise after a single trial suggests that processes forming memory of task instruction are also very flexible. Others have suggested that these fluctuations in performance, influenced by recent behavioral and reward history, may be necessary for reinforcement learning algorithms (Dayan & Balleine, 2002).

This result also suggests that the information in working memory may not be isolated to visual or perceptual information. Instead, knowledge of brick features may be bound to expectations of how this information relates to the behavioral task. Such an action-stimulus binding may comprise more complex *event files* that link together perceptual information (object features), current task context, and appropriate behaviors (Hommel, 2004). Recent experiments on PFC have shed light on how memory for behaviorally relevant features might be represented in the brain. Experiments recording activity in neurons in monkey PFC during visuo-motor tasks suggest a fundamental connection between memory for visual features and their learned behavioral relevance. PFC neurons have long been known to have sustained response during and following the presentation of a visual stimulus that the monkey was motivated to selectively act upon some time later (Fuster & Alexander, 1971). PFC connections to motor and limbic areas may also allow PFC neurons to modulate their activity with respect to the learned behavioral response (Miller & Cohen, 2001). Asaad et al. (1998) demonstrated that some individual PFC neurons are tuned not only to the visual stimulus, or to just the planned behavior, but to the conjunction between stimulus and upcoming behavior. Thus, these cells are neither exclusively stimulus or motor driven. Their response is modulated by the learned contin-

gencies between stimulus and reward-oriented behavior. It is important to note that this selective stimulus-response mapping can be relearned in only a few trials (Asaad et al., 1998; Miller & Cohen, 2001). Such a neural mechanism may underlie the modifications in hand and eye behavior during cue repetition across trials in the present experiment.

### *What Is Attention?*

Classical models of cognition often imply discrete serial stages of sensory analysis, planning, and response. The implicit assumption underlying most experimental paradigms is that the job of vision is to construct internal representations from which the brain generates a phenomenological percept, which is then used to select, plan, and execute a movement. Within this scheme, studies of visual attention have addressed the degree to which the mechanisms of feed-forward visual analysis are facilitated by the application of attention, acting either as an early- or late-stage filter or as a gain control (Luck, Woodman, & Vogel, 2000). This form of attention is thought to be manifest as a spatial spotlight capable of sweeping across the visual field, binding object features into object representations (Cave & Bichot, 1999). Other researchers have suggested that attention is more object based (Scholl, 2002; Yantis & Serences, 2003). More recently, Luck et al. (2000) have suggested that attention can operate in different cognitive subsystems under different conditions. For example, attention may bind features together to facilitate visual form recognition, facilitate the updating of working memory, or speed response selection. Similar to Luck et al.'s (2000) suggestion, the results of the present experiment suggest that a blanket label of *attention* disguises the many possible underlying processes that translate sensory input into the information used for guiding action. When possible, we have attempted to operationalize the notion of attention in the present article by referring to these fleeting functional processes in relation to the immediate demands of a task. *Encode*, *acquire*, and *storage* are all terms that have intuitive meaning but often elude formal definition. Our use of these terms may reflect computations in high-level cortical areas that moderate decisions guiding behavior. A more sophisticated understanding of these computations will be necessary to better understand the more complex challenge of how these operations are coordinated during extended behavior.

For example, consider again the variety of subtasks necessary in only the first part of a trial. Subjects must acquire the pick-up rule, search for a brick that fits the criterion for pick-up, grasp the selected brick, acquire the put-down rule, decide where the brick belongs (either using working memory for the relevant object feature or reacquiring it with an eye movement), and guide put-down. The central challenge to successfully coordinate behavior during brick sorting is to perform the right subtask at the right time. Performance of the sequence of subtasks described above requires qualitatively different kinds of operations. For example, a single fixation may reflect the acquisition of state information in the world (e.g., put-down rule). However, the operation used to deploy the next eye movement may rely on a combination of factors, including working memory and task context. For example, once the put-down rule is acquired, the subject may either next fixate the brick in hand or the conveyor belt. This decision process is arbitrated using the subjects' working memory of the brick feature and familiarity with the put-down rule. The rules of arbi-

tration, the degree to which different information influences the outcome of the response selection, are established by task context. For example, subjects reprioritize change detection following a noticed change, resulting in a higher probability of refixating the brick before put-down. Acquiring state information, choosing the next target for a fixation, and establishing task priorities are all computationally distinct operations that might all be considered as involving attention. The existence of these different computational levels becomes clearer in the context of a larger task, such as the one used in this experiment, than in more traditional paradigms that aim to isolate specific operations such as visual search or response selection. Note that this schema is radically different than the classical models, in which visual attention is thought of as a filter or as a gain control influencing a particular operation. These issues are discussed in more detail in Sprague, Ballard, & Robinson's (in press) article.

### *Change Detection and Task Performance*

A major challenge of visual cognition research is to identify what operations the brain performs throughout the course of everyday tasks. A contribution of the present experiment is to explore the possible use of the change blindness phenomenon to paradigms that demand immediate acquisition of relevant information, as found in ordinary behavior (Hayhoe et al., 2002; Land et al., 1999). If acquisition of information is constrained by task demands, providing different information within a scene during different stages of a task can reveal what information the subject is selecting and using to guide their behavior. Altering information within a scene can be construed as a change detection paradigm if detecting changes is the primary goal. However, altering information within a scene can also be broadly construed as a task use paradigm, in which experimenters can observe what information subjects chose to update and what information subjects chose to store. The present change detection experiment did not solely rely on performance at change detection per se but also emphasized how behavioral strategies could be used to acquire and store task-relevant visual information.

### References

- 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*, 723–767.
- Beck, D. M., Rees, G., Frith, C. D., & Lavie, N. (2001). Neural correlates of change detection and change blindness. *Nature Neuroscience*, *4*, 645–650.
- Birnboim, S. (2003). The automatic and controlled information-processing dissociation: Is it still relevant? *Neuropsychology Review*, *13*, 19–31.
- Blackmore, S. J., Brelstaff, G., Nelson, K., & Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, *24*, 1075–1081.
- Cave, K. R., & Bichot, N. P. (1999). Visuospatial attention: Beyond a spotlight model. *Psychonomic Bulletin & Review*, *6*, 204–223.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and



- memory of visual context guides spatial attention. *Cognitive Psychology*, 36, 28–71.
- Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 109–127.
- Cole, G. G., Kentridge, R. W., & Heywood, C. A. (2004). Visual salience in the change detection paradigm: The special role of object onset. *Journal of Experimental Psychology: Human Perception and Performance*, 30, 464–477.
- Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, 4, 519–525.
- Dayan, P., & Balleine, B. W. (2002). Reward, motivation, and reinforcement learning. *Neuron*, 36, 285–298.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113, 501–517.
- Duncan, J. (1993). Similarity between concurrent visual discriminations: Dimensions and objects. *Perception & Psychophysics*, 54, 425–430.
- Fagot, C. (1994). *Chronometric investigations on task switching*. Doctoral dissertation, University of California, San Diego.
- Fernandez-Duque, D., & Thornton, I. M. (2003). Explicit mechanisms do not account for implicit localization and identification of change: An empirical reply to Mitroff et al. (2002). *Journal of Experimental Psychology: Human Perception and Performance*, 29, 846–858.
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 1030–1044.
- Fuster, J. M., & Alexander, G. E. (1971, August 13). Neuron activity related to short-term memory. *Science*, 173, 652–654.
- Gordon, R. D., & Irwin, D. E. (1996). What's in an object file? Evidence from priming studies. *Perception & Psychophysics*, 58, 1260–1277.
- Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998, January 29). The representation of visual salience in monkey parietal cortex. *Nature*, 391, 481–484.
- Grimes, J. (1996). *On the failure to detect changes in scenes across saccades* (Vol. 2). New York: Oxford University Press.
- Hayhoe, M. M. (2000). Vision using routines: A functional account of vision. *Visual Cognition*, 7, 43–64.
- 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.
- Hoffman, J. E., Landau, B., & Pagani, B. (2003). Spatial breakdown in spatial construction: Evidence from eye fixations in children with Williams syndrome. *Cognitive Psychology*, 46, 260–301.
- Hollingworth, A. (2004). Constructing visual representations of natural scenes: The roles of short- and long-term visual memory. *Journal of Experimental Psychology: Human Perception and 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 and Performance*, 28, 113–136.
- Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences*, 8, 494–500.
- Huk, A. C., & Heeger, D. J. (2000). Task-related modulation of visual cortex. *Journal of Neurophysiology*, 83, 3525–3536.
- Irwin, D. E. (1991). Information integration across saccadic eye movements. *Cognitive Psychology*, 23, 420–456.
- Irwin, D. E., & Andrews, R. V. (Eds.). (1996). *Integration and accumulation of information across saccadic eye movements*. Cambridge, MA: MIT Press.
- Irwin, D. E., & Zelinsky, G. J. (2002). Eye movements and scene perception: Memory for things observed. *Perception & Psychophysics*, 64, 882–895.
- Ito, M., & Gilbert, C. D. (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron*, 22, 593–604.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40(10–12), 1489–1506.
- Itti, L., & Koch, C. (2001). Computational modeling of visual attention. *Nature Reviews Neuroscience*, 2, 194–203.
- Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman & 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., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, 28, 1311–1328.
- Levin, D. T., Simons, D. J., Angelone, B. L., & Chabris, C. F. (2002). Memory for centrally attended changing objects in an incidental real-world change detection paradigm. *British Journal of Psychology*, 93(Pt. 3), 289–302.
- Li, W., Piech, V., & Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience*, 7, 651–657.
- Luck, S. J., & Vogel, E. K. (1997, November 20). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Luck, S. J., Woodman, G. F., & Vogel, E. K. (2000). Event-related potential studies of attention. *Trends in Cognitive Sciences*, 4, 432–440.
- Mack, A., & Rock, I. (1996). *Inattention blindness*. Cambridge, MA: MIT Press.
- Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, 22, 657–672.
- Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Perception & Psychophysics*, 58, 977–991.
- 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.
- 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 of Neuroscience*, 24, 167–202.
- Miller, E. K., & Desimone, R. (1994, January 28). Parallel neuronal mechanisms for short-term memory. *Science*, 263, 520–522.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, 16, 5154–5167.
- Mitroff, S. R., Simons, D. J., & Franconeri, S. L. (2002). The siren song of implicit change detection. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 798–815.
- Most, S. B., Scholl, B. J., Clifford, C. W., & Simons, D. J. (in press). What you see is what you set: Sustained inattention blindness and the capture of awareness. *Psychological Review*.
- 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, 9–17.
- Motter, B. C. (1994). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *Journal of Neuroscience*, 14, 2190–2199.

- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46, 461–488.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–1031.
- O'Regan, J. K., Rensink, R. R., & Clark, J. J. (1999, March 4). Change-blindness as a result of "mudsplashes". *Nature*, 398, 34.
- Parkhurst, D. J., & Niebur, E. (2003). Scene content selected by active vision. *Spatial Vision*, 16, 125–154.
- Parkhurst, D. J., & Niebur, E. (2004). Texture contrast attracts overt visual attention in natural scenes. *European Journal of Neuroscience*, 19, 783–789.
- Pashler, H., Johnston, J. C., & Ruthruff, E. (2001). Attention and performance. *Annual Review of Psychology*, 52, 629–651.
- Pelz, J. B., & Canosa, R. (2001). Oculomotor behavior and perceptual strategies in complex tasks. *Vision Research*, 41(25–26), 3587–3596.
- Pelz, J., Hayhoe, M., & Loeber, R. (2001). The coordination of eye, head, and hand movements in a natural task. *Experimental Brain Research*, 139, 266–277.
- Rensink, R. A. (2000a). Change blindness: Implications for the nature of visual attention. In L. R. J. Harris, M. (Ed.), *Vision and Attention* (pp. 169–188). New York: Springer.
- Rensink, R. A. (2000b). The dynamic representation of scenes. *Visual Cognition*, 7, 17–42.
- Rensink, R. A. (2000c). Seeing, sensing, and scrutinizing. *Vision Research*, 40(10–12), 1469–1487.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8, 368–373.
- Ress, D., & Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nature Neuroscience*, 6, 414–420.
- Roelfsema, P. R., Khayat, P. S., & Spekreijse, H. (2003). Subtask sequencing in the primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 100, 5467–5472.
- Roelfsema, P. R., Lamme, V. A., & Spekreijse, H. (1998, September 24). Object-based attention in the primary visual cortex of the macaque monkey. *Nature*, 395, 376–381.
- Roelfsema, P. R., Lamme, V. A., & Spekreijse, H. (2000). The implementation of visual routines. *Vision Research*, 40, 1385–1411.
- Saiki, J. (2003). Feature binding in object-file representations of multiple moving items. *Journal of Vision*, 3, 6–21.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1–66.
- Scholl, B. J. (Ed.). (2002). *Objects and attention*. Amsterdam: Elsevier Science.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127–190.
- Shinoda, H., Hayhoe, M. M., & Shrivastava, A. (2001). What controls attention in natural environments? *Vision Research*, 41(25–26), 3535–3545.
- Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, 7, 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 & Review*, 5, 644–649.
- Sprague, N., Ballard, D., & Robinson, A. (in press). Modeling attention with embodied visual behaviors. *ACM Transactions on Action and Perception*.
- Thornton, I. M., & Fernandez-Duque, D. (2000). An implicit measure of undetected change. *Spatial Vision*, 14, 21–44.
- Thornton, I. M., & Fernandez-Duque, D. (2002). Converging evidence for the detection of change without awareness. *Progress in Brain Research*, 140, 99–118.
- Treisman, A. (1988). Features and objects: The fourteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 40(A), 201–237.
- Treue, S., & Maunsell, J. H. R. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *Journal of Neuroscience*, 19, 7591–7602.
- 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 symposium on Eye Tracking Research & Applications* (pp. 95–102). New York: ACM Press.
- Ullman, S. (1984). Visual routines. *Cognition*, 18(1–3), 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.
- Wheeler, M. E., & Treisman, A. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131, 48–64.
- Yantis, S., & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 661–676.
- Yantis, S., & Hillstrom, A. P. (1994). Stimulus-driven attentional capture: Evidence from equiluminant visual objects. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 95–107.
- Yantis, S., & Serences, J. T. (2003). Cortical mechanisms of space-based and object-based attentional control. *Current Opinion in Neurobiology*, 13, 187–193.
- Yarbus, A. L. (1967). *Eye movements and vision*. New York: Plenum Press.

(Appendixes follow)

## Appendix A

### Criteria for Fixation Sequence Analysis

For this analysis, three new criteria were applied to screen trials and to categorize fixation targets. The first criterion involved further screening of eye position data. Of the 43 subjects with adequate tracking, occasional track loss was inevitable. Track loss may result in the failure to detect a fixation on a specific area. For a participant to perform the task, a fixation was necessary in each of four areas: the pick-up cue, a brick in the array, the put-down cue, and the belt area. (It was technically possible, though unusual, to perform a trial without fixating the brick in hand.) Because of eyeblinks, track loss, or imprecise fixation calibration, the automated analysis sometimes failed to record fixations to one or more of the necessary four areas. Although these trials were included in the analysis in Figure 4, it is more critical for analysis of fixation sequences. Therefore trials without a fixation to one of these four areas were excluded. This criterion excluded an average of 32% of trials for each subject. Such a high percentage were excluded because of the requirement for tracking several specific fixations. Analysis of the video record suggested that these excluded trials often had fixation close, if not on, each required target. The pattern of fixations during excluded trials were subjectively similar to those included for subse-

quent analysis. For example, the majority of the excluded trials did not include fixations to the pick-up cue or the brick array, either because fixations were directed in between the cue and the array or because the fixation was not detected on the brick until it was grasped during pick-up. The second criterion concerned the categorization of fixation area. Fixations to the brick in hand sometimes occurred while the brick was being picked up, sometimes immediately after fixating the sorting cue while the brick was still held in between the belts, and sometimes as the brick was being placed on a conveyor belt. Fixations to the brick while it occluded the belts invariably occurred when subjects were placing the brick down. Thus, these fixations were classified as belt fixations and not as fixations to the brick in hand in Table 1, because this analysis was designed to better reflect the functional stage of the task. The third criterion was that no trial ended in a fixation to either the pick-up cue or the array. After releasing a brick on the conveyor belt at the completion of a trial, subjects often fixated on the pick-up cue or brick array to prepare for the subsequent trial (note bumps in Figures 4A and 4B). Thus, fixations to the pick-up cue or the brick array at the end of the trial were not included in the analysis of fixation sequence.

## Appendix B

### Comparing Rates of Change Detection Between *One Feature* and *Two Feature* Trials

The three irrelevant features in the *One Feature* task were subdivided into two categories: (a) pre-put-down, the one feature that would later be used for put-down in the *Two Feature* task and (b) never relevant, the two other features that remained irrelevant throughout the experiment. During the *One Feature* task the pre-put-down feature differs in both the quality of the feature and the long-term task relevance, because subjects used this feature for sorting during the practice trials and may have anticipated its future use for the upcoming *Two Feature* condition. We then analyzed rates of change detection in each task depending on the relevance of the feature throughout the duration of the experiment. There was no significant difference in change detection rates between pre-put-down and never-relevant

features during the *One Feature* task. Nor was there a difference in change detection between the *One Feature* and *Two Feature* condition for features that remained irrelevant throughout the duration of the experiment. Lastly, when a previously irrelevant feature was made relevant for sorting (pre-put-down feature in *One Feature* and *Two Feature* tasks), subjects increased their rates of detection for that feature,  $F(1, 50) = 14.33, p < .01$ .

Received October 5, 2004

Revision received May 13, 2005

Accepted May 23, 2005 ■