

- 63 Sporns, O., Tononi, G. and Edelman, G.M. (1991) Modeling perceptual grouping and figure-ground segregation by means of active reentrant connections *Proc. Natl. Acad. Sci. U. S. A.* 88, 129-133
- 64 Tononi, G., Sporns, O. and Edelman, G.M. (1992) Reentry and the problem of integrating multiple cortical areas - simulation of dynamic integration in the visual system *Cereb. Cortex* 2, 310-335
- 65 Leonards, U., Singer, W. and Fahle, M. (1996) The influence of temporal phase differences on texture segmentation *Vision Res.* 36, 2689-2697
- 66 Leonards, U. and Singer, W. (1997) Selective temporal interactions between processing streams with differential sensitivity for colour and luminance contrast *Vision Res.* 37, 1129-1140
- 67 Leonards, U. and Singer, W. Two segmentation mechanisms with differential sensitivity for colour and luminance contrast *Vision Res.* (in press)
- 68 Kiper, D.C., Gegenfurtner, K.R. and Movshon, J.A. (1996) Cortical oscillatory responses do not affect visual segmentation *Vision Res.* 36, 539-544
- 69 Freiwald, W.A., Kreiter, A.K. and Singer, W. (1997) in *Göttingen Neurobiology Report 1997* (Elsner, N. and Wässle, H., eds), p. 550, Thieme

Change blindness

Daniel J. Simons and Daniel T. Levin

Although at any instant we experience a rich, detailed visual world, we do not use such visual details to form a stable representation across views. Over the past five years, researchers have focused increasingly on 'change blindness' (the inability to detect changes to an object or scene) as a means to examine the nature of our representations. Experiments using a diverse range of methods and displays have produced strikingly similar results: unless a change to a visual scene produces a localizable change or transient at a specific position on the retina, generally, people will not detect it. We review theory and research motivating work on change blindness and discuss recent evidence that people are blind to changes occurring in photographs, in motion pictures and even in real-world interactions. These findings suggest that relatively little visual information is preserved from one view to the next, and question a fundamental assumption that has underlain perception research for centuries: namely, that we need to store a detailed visual representation in the mind/brain from one view to the next.

Since antiquity, scholars have assumed the need for precise, veridical representations of our visual world¹. Modern researchers recognize that the two-dimensional retinal image cannot fully and unambiguously represent a three-dimensional world, and since Descartes, they have posited adjustments to the retinal image to compensate for distortions and ambiguities. In order to form an accurate, stable representation, we must somehow extract the invariant structure of the world from our ever-changing sensory experience. For example, as we view scenes in the real world, we move our eyes (saccade) three to four times each second. Across fixations, objects in the world are projected onto different parts of the retina. Somehow, we integrate information across these fixations to achieve a stable representation; we must recognize that two consecutive views are, in fact, of the same scene even when the viewpoint or viewing angle differs. Recently, research on how we integrate visual information across fixations has spawned a new series of studies focusing on our ability to detect changes from one view of a scene to the next. These studies have produced a set of results that, consistent with earlier evidence for memory distortions, suggest a high degree of 'change blindness'; observers do not appear to retain many visual details from one view to the next. These

recent findings and their implications for how we represent our visual world are the primary focus of this review.

Evidence for change blindness

Although change detection has only recently become a topic of intense inquiry, research spanning many areas of cognitive psychology has hinted at current findings. At times, research on visual integration of information across eye movements has revealed striking examples of our inability to detect changes (see Box 1). Research on recognition memory for large numbers of photographs also suggested the possibility of change blindness. The primary purpose of such studies was to demonstrate an impressive capacity to remember photographs from a single presentation, but they also revealed a lack of specificity in representations (see Box 2). Perhaps the most intriguing precursor to contemporary studies of change blindness comes from the informal observations of film makers and editors (see Box 3; several recent studies are described in the text).

Over the past ten years, a number of studies have begun to address our inability to detect changes to objects and scenes from one view to the next (see Box 4). Some experiments changed images during saccades. Others made changes

D.J. Simons
Department of
Psychology,
University of
Kirkland S.
820, Cambridge
MA 02138

tel: +1 617 495 5300
fax: +1 617 495 3728
e-mail: dsimons@wjh.harvard.edu

D.T. Levin is in the
Department of
Psychology, Kent
State University, PO
Box 5190, Kent, OH
44242-0001, USA

tel: +1 330 672 3785
fax: +1 330 672 3786
e-mail: dlevin@kent.edu

Box 1. Visual integration across eye movements

During the 1970s, evidence for visual masking and for the integration of visual information at a single retinal location^{a,b}, together with the general acceptance of the construct of iconic memory (a short-term sensory memory that retains a detailed picture-like representation of a scene)^c, inspired a model for achieving a continuous experience^d. This model suggested that visual images from consecutive views are combined in a visual buffer, much as two overhead transparencies can be superimposed.

Although this model seems plausible, it cannot account for continuity under natural viewing conditions. Somehow, visual integration in the real world must accommodate changes to our eye, head and body positions. In order for the visually integrative buffer model to work, stimuli presented on two different fixations or at two different retinal locations must be integrated visually. That is, our visual system must determine that an object is the same even when it stimulates different areas of the retina on consecutive fixations. In one test of this hypothesis, subjects fixated a point in the center of a display and a 12-dot pattern was presented briefly to parafoveal vision. Shortly thereafter, subjects moved their eyes to the parafoveal location and a second 12-dot pattern was presented. When the two patterns were combined, one dot was missing, and subjects were asked to determine the location of the missing dot. Although initial studies supported the notion of a visually integrative buffer^e, later studies controlling for methodological and display artifacts have failed to replicate the initial finding^{f-k} (for recent reviews, see Refs l and m). Therefore, this research fails to support the hypothesis that we form an accurate representation by storing and integrating precise visual information from one fixation to the next.

Additional evidence for the absence of integrated visual representation across eye movements comes from the study of preview effects in reading. One particularly dramatic example comes from a task in which observers read lines of text that alternated case with each letter (e.g. AlTeRnAtEd CaSe)ⁿ. During some saccades, every letter in the sentence changed case, so that the visual form of every word was different. Surprisingly, when the changes occurred during an eye movement, subjects almost never noticed. That is, subjects not only failed to integrate the visual form of the letters from one instant to the next, they could not even tell that the visual form was changing. Apparently, the information integrated across fixations during reading is not contingent on the precise visual form of the word. More recently, studies of reading have inspired a series of studies of integration of pictorial information across eye movements. These studies have focused on the benefits of a parafoveal preview on processing during a subsequent fixation^{m,o,p,q}. One recent study showed that when complementary sets of contours from an object were shown before and after an eye movement, in general, observers were unable to detect the change and the contour change had no effect on naming

latencies^r. The visual form was not sufficiently represented to allow detection of the change.

In a sense, studies of visual integration and studies of change detection address the same issues using complementary methodologies. Studies of visual integration focus on the ability to combine two distinct images, essentially adding their contents. Studies of change detection focus on the ability to subtract one image from another, thereby finding the difference. Both approaches allow an exploration of the specificity of scene representations.

References

- a Di Lollo, V. (1980) Temporal integration in visual memory *J. Exp. Psychol. Gen.* 109, 75-97
- b Kahneman, D. (1968) Method, findings, and theory in studies of visual masking *Psychol. Bull.* 70, 404-425
- c Neisser, U. (1967) *Cognitive Psychology*, Appleton-Century-Crofts
- d McConkie, G.W. and Rayner, K. (1976) Identifying the span of the effective stimulus in reading: Literature review and theories of reading, in *Theoretical Models and Processes of Reading* (2nd edn) (Singer, H. and Ruddell, R.B., eds), pp. 137-162, International Reading Association
- e Jonides, J., Irwin, D.E. and Yantis, S. (1982) Integrating visual information from successive fixations *Science*, 215, 192-194
- f Bridgeman, B. and Mayer, M. (1983) Failure to integrate visual information from successive fixations *Bull. Psychonomic Soc.* 21, 285-286
- g Irwin, D.E., Brown, J.S. and Sun, J-S. (1988) Visual masking and visual integration across saccadic eye movements *J. Exp. Psychol. Gen.* 117, 276-287
- h Irwin, D.E., Yantis, S. and Jonides, J. (1983) Evidence against visual integration across saccadic eye movements *Percept. Psychophysiol.* 34, 49-57
- i Jonides, J., Irwin, D.E. and Yantis, S. (1983) Failure to integrate information from successive fixations *Science* 222, 188
- j Rayner, K. and Pollatsek, A. (1983) Is visual information integrated across saccades? *Percept. Psychophysiol.* 34, 39-48
- k Sun, J-S. and Irwin, D.E. (1987) Retinal masking during pursuit eye movements: Implications for spatiotopic visual persistence *J. Exp. Psychol. Hum. Percept. Perform.* 13, 140-145
- l Irwin, D.E. (1991) Information integration across saccadic eye movements *Cognit. Psychol.* 23, 420-456
- m Pollatsek, A. and Rayner, K. (1992) What is integrated across fixations? in *Eye Movements and Visual Cognition: Scene Perception and Reading* (Rayner, K., ed.), pp. 166-191, Springer-Verlag
- n McConkie, G.W. and Zola, D. (1979) Is visual information integrated across successive fixations in reading? *Percept. Psychophysiol.* 25, 221-224
- o Henderson, J.M. (1997) Transsaccadic memory and integration during real-world object perception *Psychol. Sci.* 8, 51-55
- p Pollatsek, A., Rayner, K. and Collins, W.E. (1984) Integrating pictorial information across eye movements *J. Exp. Psychol. Gen.* 113, 426-442
- q Pollatsek, A., Rayner, K. and Henderson, J.M. (1990) Role of spatial location in integration of pictorial information across saccades *J. Exp. Psychol. Hum. Percept. Perform.* 16, 199-210

during blinks or during a blank interval between two pictures. Still others made changes to scenes while observers viewed a motion picture cut or a real-world occlusion event. What is striking about this diversity of approaches is the similarity of the results. In all of these experiments, observers fail to notice dramatic changes to displays. We will now turn to evidence for and mechanisms underlying change blindness across saccades.

Saccade-contingent changes

Imagine viewing a set of photographs for an upcoming recognition test. As you study the photographs, you shift your attention among the objects in the image and you scan the image with your eyes. Periodically, while your eyes are moving rapidly from one object to the next, something in the scene is changed. The experimenters mention that the scenes may change at times and that you should let them

know if you see something change. Sounds easy, right? You are studying the photographs intently for a test, so you should have a fairly complete representation.

Surprisingly, observers failed to notice when two men in a photograph exchanged different colored hats and only 50% noticed when two people exchanged heads². In all, subjects missed nearly 70% of the changes that occurred during an eye movement. Subsequent studies^{3,4} have confirmed the basic pattern described by Grimes. During an eye movement, we apparently lose, or at least lose access to, many of the visual details of the previous view. But is change blindness limited to cases in which information must be integrated across eye movements? These dramatic findings, accompanied by theoretical predictions of sparse visual representations⁵⁻⁷, spurred a flurry of investigations into the mechanisms underlying change blindness.

Change blindness across simulated saccades

One aspect of eye movements that might account for change blindness is the existence of motion transients across the retina. In a sense, the target change cannot be identified because the eye is processing change signals from every location. If global transients effectively mask the ability to localize an individual transient or change, then any display that creates global transients should make change detection difficult. In order to examine this possibility, several laboratories independently developed a technique designed to mimic eye movements without changing the fixation location^{8,9} (see Fig. 1).

Experiments using the flicker paradigm found that almost none of the changes were detected during the first cycle of alternation, and many changes were not detected even after nearly one minute of alternation⁹. When the blank screen was removed, eliminating the disruption caused by the global transient, the changes were detected easily^{9,10}. The change blindness caused by a flashed blank screen suggests the possibility that other forms of global transient should be equally successful in hiding changes. In fact, another recent study has demonstrated comparable results when changes are made contingent on blinks¹¹.

All of these findings illustrate the absence of a precise visual representation that survives global transients, yet they all suffer from one potential criticism. In all of these experiments, the global transient effectively covers the location of the change. Perhaps the blank screen, eye movement or blink actually serves as a mask, interrupting processing at that location. Another recent series of studies eliminates the masking explanation using the same alternating images used in the flicker studies. Rather than interspersing a blank screen, experimenters flashed a set of dot patterns at arbitrary positions on the image simultaneously with the image change¹². These dots created several additional local transients, giving the appearance of a mud splash hitting the windscreen of a car. However, they did not mask the image change. Even so, these additional local transients had an effect similar to an eye movement or a flashed blank screen: observers could not detect changes immediately, even though the changes produced local retinal transients. Although the degree of change blindness was somewhat attenuated relative to performance in the flicker paradigm (R.A. Rensink, pers. commun.), change detection still required substantial time and effort.

Box 2. Recognition memory for photographs

In typical studies of scene memory, observers viewed hundreds and sometimes thousands of photographs of natural scenes. Later, they tried to identify which photographs they had studied and which they had never viewed before^{a,b,c,d}. Although the larger conclusion of these studies was that observers can recognize previously viewed photographs at surprisingly high rates (sometimes exceeding 95% recognition after extended delays), several studies noted that memory for the images was not tied to the precise visual form of the image^{e,c}. When previously viewed photographs were mirror-reversed during a test, observers did not detect the change and reported them as previously viewed^c. These findings suggest not only that we fail to detect changes to the exact visual form of a scene but, also, that we can extract the gist or meaning of a scene and use it for recognition^a.

References

- a Nickerson, R.S. (1965) Short-term memory for complex meaningful visual configurations: A demonstration of capacity *Can. J. Psychol.* 19, 155-160
- b Shepard, R.N. (1967) Recognition memory for words, sentences, and pictures *J. Verb. Learn. Verb. Behav.* 6, 156-163
- c Standing, L., Conezio, J. and Haber, R.N. (1970) Perception and memory for pictures: Single-trial learning of 2500 visual stimuli *Psychonomic Sci.* 19, 73-74
- d Standing, L. (1973) Learning 10,000 pictures. *Q. J. Exp. Psychol.* 25, 207-222
- e Pedzek, K. et al. (1988) Picture memory: Recognizing added and deleted details *J. Exp. Psychol. Learn. Mem. Cognit.* 14, 468-476

The role of attention and expectations

The relatively long detection latencies in these studies of change blindness suggest that change detection is an active searching process in which individual objects are encoded and compared sequentially across views. Although the data do not speak directly to the nature of the search for changes, one particular result reinforces this possibility. Changes to objects in the 'center of interest' of a scene (according to independent ratings) are detected more readily, even when the physical magnitude of the change is comparable to that of a non-central change^{9,10}. This finding suggests that attention is focused on central objects either more rapidly or more often, thereby allowing faster change detection. The notion of a center of interest has important implications for how we encode our environment. Given the results of the mud splash and flicker experiments, we know that changes will not be encoded automatically and that some effort is needed to detect changes even with a localized transient. Apparently, the center of interest benefit derives not from the automatic representation of a precise visual image but from the abstraction of a scene's contents.

If abstraction plays a central role in our representation of scenes from one view to the next, then broad expectations about a scene may influence how we encode objects in that scene and, consequently, how we represent those objects. Interestingly, some of the early work on scene context and expectations found similar change blindness and center of interest effects^{13,14}. For example, in one experiment, observers

Box 3. Insights from film makers

In the movie *Ace Ventura: When Nature Calls*, the pieces on a chess board disappear completely from one shot to the next. In *Goodfellas*, a child is playing with blocks that appear and disappear across shots. One inevitable consequence of film production is the need to shoot scenes out of order, and often to shoot components of the same scene at different times. As a result, unintentionally, many details within a scene may change from one view to the next. Although film makers go to considerable effort to eliminate such errors, almost every movie – in fact, almost every cut – has some continuity mistake. Yet, most of the time, people are blind to these changes. (Film makers are, of course, justified in trying to eliminate glaring errors given the potential costs of some audience members noticing the change. If just one viewer notices such a change, the popular media and the Internet community will publicize the change and inspire people to look for the editing mistake rather than focusing on the movie.)

Film makers have long had the intuition that changes to the visual details of a scene across cuts are not detected by audiences, particularly when editing allows for smooth transitions⁴. For example, the film maker Lev Kuleshov⁵ notes that: 'convincing montage makes the audience overlook...minor defects (for example when the actor's costume changes between shots), though I repeat that this is only possible if the scene is edited correctly (in case of bad montage the blunder will leap to the eye).' Dmytryk⁶ notes that change blindness is evident when mistakes occur: 'far from the viewer's center of interest. If he is watching the actor's eyes, a mismatch of an arm or hand will be ignored nine times out of ten.' Such intuitions underlie Hochberg's more recent speculations about the 'sketchiness' of visual memory⁷ and clearly predict the center of interest effects described in the text.

The craft of film editing constitutes a rich body of knowledge about vision. Film makers must do explicitly what our visual system does automatically: they must combine a series of partial views (individual shots) into a coherent whole (a continuous scene) without audiences noticing the transitions. Some editors even suggest ways to cause

audiences to move their eyes or blink, thereby allowing a cut to go unnoticed⁸. This process of constructing a continuous visual scene has taught film makers that the visual details are not central to our understanding of a scene, but it has also given them intuitions about what is central. For example, film makers track the gaze direction of actors in each shot and are careful not to violate the relative spatial locations of gaze targets. Ongoing empirical studies of the importance of cues such as gaze direction and motion in motion pictures may provide a better understanding of how we perceive the layout of scenes in the real world.

References

- a Dmytryk, E. (1984) *On Film Editing: An Introduction to the Art of Film Construction*, Focal Press
- b Kuleshov, L. (1987) *Selected Works: Fifty Years in Films* (Agrachev, D. and Belenkaya, N., translators), Raduga Publishers
- c Hochberg, J. (1986) Representation of motion and space in video and cinematic displays, in *Handbook of Perception and Human Performance* (Vol. 1) (Boff, K.R., Kaufman, L.K. and Thomas, J.P., eds), pp. 22.1–22.64, John Wiley and Sons




Fig. Person change in a motion picture. This figure depicts four frames from a motion picture in which one actor is replaced by a different actor during a change in camera angle (B and C). Even though the actor was clearly the central object of the scene, many observers failed to detect this change.

viewed scenes that included both consistent and inconsistent objects in preparation for an upcoming recognition task⁴. During the testing phase, observers were asked to discriminate previously viewed images from similar images in which an object had been changed. Observers were more

likely to notice changes to the schema-inconsistent objects than the schema-consistent ones (also see Ref. 15). Unexpected objects are more likely to garner attentional resources, and attended objects are more likely to be retained from one view to the next. More recently, superior recognition of

Box 4. What is in a view

The word 'view' has a number of distinct meanings in perception research. A view can refer to a single fixation, a single viewing position or angle, or even a photograph of a scene. Here we take a view to be an unchanging image, essentially a snapshot of a scene. For studies of visual integration across eye movements (see Box 1), views are equivalent to fixations; observers fixate a single image and after they move their eyes, the scene has been changed, producing a different view. In studies of change detection across simulated saccades, each time the image changes observers experience a different view. In motion pictures, each change in camera angle produces a new view of the scene even though the content of the scene itself may be unchanged. Across a cut, a view change occurs between the final frame of the first shot and the initial frame of the second shot. Using this definition, all of the work we discuss involves changes across views, even when the changes occur across long delay intervals.

inconsistent objects has been shown in recognition of objects from real rooms¹⁶ (but see Ref. 13 for evidence that schema-consistent objects are better retained). Friedman anticipated many current claims of change blindness when she noted that for expected objects: 'local visual details of objects... are thus not generally encoded'¹⁴. She also made the prediction that any change not altering the abstract description of a scene substantially is unlikely to be detected.

If scene schemas help determine which changes will be noticed, then models proposing precise representations of the visual details of our environment may still prevail. Perhaps change blindness only applies to peripheral and unattended objects. Even if the visual details of peripheral objects are not represented precisely, the details of centrally attended objects may be. Several laboratories have tested this possibility.

Changes to attended objects

One test of the detectability of changes to attended objects comes from a task in which participants duplicate a pattern of colored blocks¹⁷ (see Fig. 2). Participants were observing the display carefully, in order to perform an action, so the model must be considered the center of interest of the scene. Yet, they failed to notice a change to the model pattern when it occurred during an eye movement.

In another series of studies¹⁸, observers viewed an array of five objects on a computer monitor. After two seconds, the array was replaced with a blank screen, followed shortly by another array of five objects. The second array was either identical to the first or was different in one of three ways: (a) one of the objects was moved to a previously empty location, (b) one object was replaced by an object that was not in the original array, or (c) two objects in the original array switched places. Observers were asked to determine whether or not any change had occurred (see Refs 19 and 20 for earlier work using a similar method). As in the flicker paradigm, observers missed changes when an object was replaced with a different object or when two objects switched places. As in the block-copying task, observers clearly focused attention on the objects in each display. The primary goal of the task was change detection, so they knew that they should encode the objects. In recent pilot studies using the same paradigm coupled with eye tracking (D.J. Simons and M. Spivey-Knowlton, unpublished), we found that, typically, observers look at all of the objects in the display. Thus, change blindness does not appear to result from a failure to focus attention on the target object during the trial.

Although these two studies using different methods converge on the conclusion that even attended objects may not be encoded sufficiently to allow change detection, both involved displays in which the observer's attention shifts from object to object during encoding. Perhaps our visual system can only tolerate one central object at a time. Successful change detection may only occur when the target object is the central object immediately before and after the change. To examine this possibility, we used motion pictures to change an object that remained central throughout the scene²¹ (see Box 3). In these films, a single character performed a simple action such as rising from a chair and

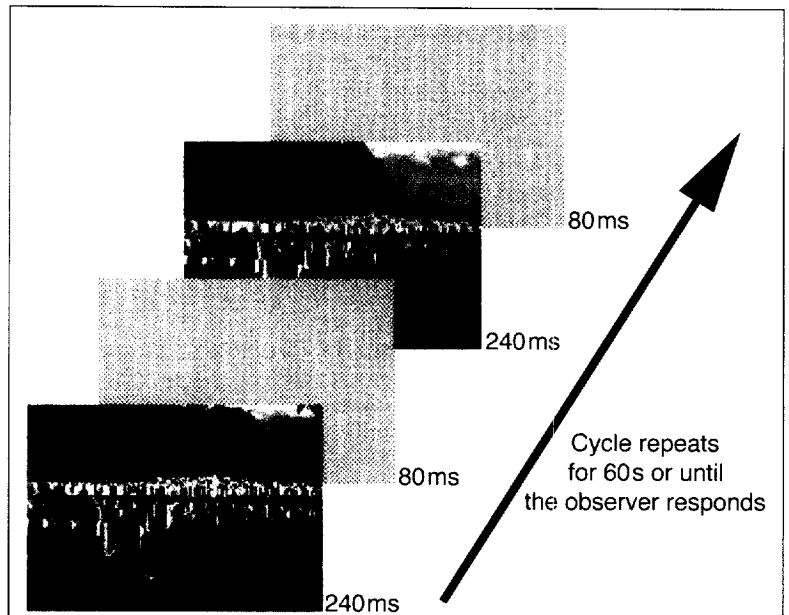


Fig. 1 The flicker technique. An original image and a modified image of a natural scene are presented in rapid alternation with a blank gray screen interposed between them, giving the display a flickering appearance (this method is often called the 'flicker paradigm'). The cycle of alternation repeats until observers report the change, and response latency is used as a dependent measure of change blindness. If change blindness is caused by masking local transients with global ones, then the flashed blank screen should decrease change detection. This figure was modeled on a figure appearing in a recent paper by Rensink *et al.*⁹.

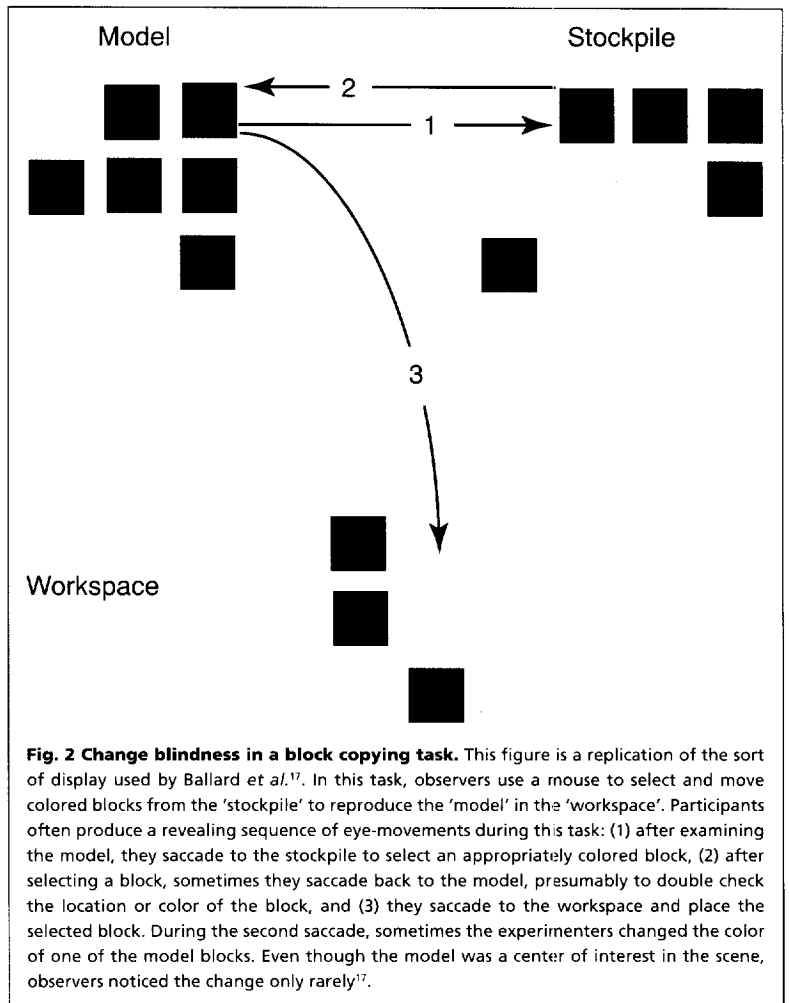


Fig. 2 Change blindness in a block copying task. This figure is a replication of the sort of display used by Ballard *et al.*¹⁷. In this task, observers use a mouse to select and move colored blocks from the 'stockpile' to reproduce the 'model' in the 'workspace'. Participants often produce a revealing sequence of eye-movements during this task: (1) after examining the model, they saccade to the stockpile to select an appropriately colored block, (2) after selecting a block, sometimes they saccade back to the model, presumably to double check the location or color of the block, and (3) they saccade to the workspace and place the selected block. During the second saccade, sometimes the experimenters changed the color of one of the model blocks. Even though the model was a center of interest in the scene, observers noticed the change only rarely¹⁷.

Outstanding questions

- Evidence reviewed in the text suggests that scene-inconsistent objects are coded more thoroughly than scene-consistent objects^{14,16,29}. How do we form expectations for a scene and to what extent do they influence our coding of objects and our ability to detect changes? If expectations influence what we attend to in a scene, then an instructional manipulation might lead to a radical difference in change blindness. Specifically, we would predict that a single change occurring in an ambiguous scene might be noticed for one interpretation but not another. More specifically, would changes only be detected when they violate expectations about a scene?
- When observers are not searching actively for a change, in general, they do not detect changes, even to the central object in a scene. When our behavioral goal is not change detection, we tend to encode the gist of a scene without explicitly coding the details that would allow change detection. In such cases, we are attending to objects and scenes, but we do not appear to encode the details necessary for change detection. What then does it mean to 'attend to' an object? Are there different kinds of attention that might influence change detection? How does the attention that leads to successful change detection differ from the kind of attention that makes an object the center of interest?
- What sorts of changes are likely to be noticed and what differentiates those from unnoticed ones? We have some evidence that changes to spatial layout are detected more readily than changes to object properties¹⁸, but a simple thought experiment shows that some property changes should be easy to notice. For example, if one of the actors in a video were replaced with a skunk, people would certainly notice the change. What aspect of this change would make it noticed more easily?
- Does change blindness indicate the absence of representations? Although this conclusion may be appealing, the possibility remains that more sensitive measures than change detection would reveal some underlying representation. Recent work on inattention blindness shows that background objects in the visual field can influence judgments even if observers are unaware of their existence³⁰. If so, what is the nature of these representations (for example, how precise and detailed are they) and in what ways do they influence our behavior? If such 'implicit' representations are precise and detailed, why do they not allow us to detect changes?
- How can we reconcile evidence for form-specific visual memory in priming studies^{31,32} with change blindness?

answering a telephone or entering through a doorway and sitting in a chair. During the action sequence, we cut from one camera angle to another, and during that cut, the original actor in the scene was replaced by a different person who then completed the action. The changes in camera angle followed conventional editing techniques by cutting in the middle of the action²². Even though the character was clearly the central object in the scene, 67% of observers failed to detect the change from one actor to another. Despite their change blindness, observers could describe the action sequence accurately and sometimes described properties of one or both actors. The actors were clearly discriminable; when given a set of films, half of which had changes to the actor and half of which did not, observers who had been instructed to look for changes had little trouble detecting them (see Box 3).

Although these findings demonstrate that changes to objects in the center of interest are not necessarily noticed, they support the claim that attention and abstract encoding are necessary for change detection. When observers search for and encode the features that individuate people, they can detect the change. Yet, under natural viewing conditions, they are unlikely to do so. Instead, they encode the

gist of the scene (in this case, the specific action and a few characteristics of the actor) and ignore the visual details. As long as the gist remains the same, change detection is unlikely because observers have not expended the effort to encode more details. This encoding strategy makes sense given the innumerable perceptual features in a natural scene or event, but it also illustrates the degree to which we lack a detailed representation of our world.

Although these findings of change blindness, taken together, support the conclusion that we lack a precise representation of our visual world from one view to the next, none of them focused on the representation of real objects in our environment. All of them presented scenes and objects on computer monitors and television displays which clearly lack many of the properties of real objects. One final series of studies examined the possibility that computer displays and motion pictures do not reflect how we process objects in the real world. In these studies, we extended the person change video studies²¹ to the real world. Imagine that a person approaches you and asks for directions. Kindly, you oblige and begin describing the route. While you are talking, two people interrupt you rudely by carrying a door right between you and the person you were talking to. Surely you would notice if the person you were talking to was then replaced by a completely different person. When we actually conducted this study, only 50% of observers noticed the change (D.J. Simons and D.T. Levin, unpublished). The two experimenters wore different clothing, were different heights and builds, had different haircuts and had noticeably different voices.

Interestingly, those who did notice the change were students of roughly the same age as the experimenters and those who failed to notice it were older than the experimenters. We theorized that this age difference may reflect a difference in how people abstracted the gist of the scene. Older participants would be more likely to encode the event as 'some student asking directions'; younger participants would be more likely to individuate the features of a person in their own social group and, thereby, would be more likely to encode those features that would discriminate the two experimenters²³. To examine this possibility, we replicated the door event with the same two experimenters dressed as construction workers (again with different clothing) under the assumption that these costumes would place the experimenters in a social group distinct from the students. Under these conditions, fewer than half of the students noticed the change. These studies demonstrate convincingly that paying attention to an object by no means guarantees change detection. The central object in a scene and the focus of a social interaction changed without observers noticing. As in studies of recognition memory for objects in photographs and rooms^{14,16}, observers were unlikely to notice changes that did not violate the gist of the scene. In this case, as long as the rough description of the scene was the same before and after the change, observers did not notice.

Summary and future directions

Taken as a whole, these findings provide a striking picture of our ability to perceive and represent scenes. Although the ability to discriminate and recognize photographs of scenes

can be exceptionally good²⁴⁻²⁸, memory for the properties and features of objects in scenes is surprisingly transitory. Findings from research on perception for action, change detection, motion picture perception and real world interactions all suggest that the visual details of object properties are not retained automatically from one view to the next. We fail to notice changes to scenes when they do not produce a motion on our retina that attracts attention. Although changes to central objects are more likely to be detected, they are not detected automatically. Therefore, attention is necessary, but not sufficient, for change detection.

Given failures of change detection, we must question the assumption that we have a detailed representation of our visual world. And, given our success in interacting and behaving in our environment, we must ask whether such detailed representations are even necessary. Although change blindness might appear to contradict our phenomenal experience of a stable, continuous world it may actually account for this impression. During any fixation we have a rich visual experience. From that visual experience, we abstract the meaning or gist of a scene. During the next visual fixation, we again have a rich visual experience, and if the gist is the same, our perceptual system assumes the details are the same. Consider, for example, a busy city street. In this kind of scene, a variety of property changes occur during the normal course of events. People are occluded by walking behind barriers, cars move, revealing previously hidden objects on the sidewalk, and people may shift a bag to the other arm or take out a handkerchief. All of these changes are rapid and might occur during saccades or between successive glances. A system that is too precise in tracking visual details would, in the words of William James, present a 'blooming, buzzing confusion'. In contrast, a system that gives a rich perceptual experience at any instant, but only integrates the gist (and perhaps the layout and movement direction) from one view to the next would give the impression of stability rather than chaos. This system would be successful at ignoring unreliable object property information, focusing instead on the information that the perceiver needs to know. Thus, change blindness supports the phenomenal experience of continuity by not preserving too much information from one view to the next.

Acknowledgements

The author thanks John Henderson and Ron Rensink for their comments and suggestions on an earlier draft of this manuscript.

References

- 1 Lindberg, D.C. (1976) *Theories of Vision from Al-Kindi to Kepler*, University of Chicago Press
- 2 Grimes, J. (1996) On the failure to detect changes in scenes across saccades, in *Perception (Vancouver Studies in Cognitive Science)* (Vol. 2) (Akins, K, ed.), pp. 89-110, Oxford University Press
- 3 Currie, C. et al. (1995) Maintaining visual stability across saccades: Role of the saccade target object. Technical Report UIUC-BI-HPP-95-01, Beckman Institute, University of Illinois
- 4 McConkie, G.W. and Currie, C.B. (1996) Visual stability across saccades while viewing complex pictures *J. Exp. Psychol. Hum. Percept. Perform.* 22, 563-581
- 5 Dennett, D.C. (1991) *Consciousness Explained*, Little, Brown and Company
- 6 O'Regan, J.K. (1992) Solving the 'Real' mysteries of visual perception: The world as an outside memory *Can. J. Psychol.* 46, 461-488
- 7 Stroud, J.M. (1955) The fine structure of psychological time, in *Information Theory in Psychology: Problems and Methods* (Quastler, H., ed.), pp. 174-207, Free Press
- 8 Blackmore, S.J. et al. (1995) Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes *Perception* 24, 1075-1081
- 9 Rensink, R.A., O'Regan, J.K. and Clark, J.J. (1997) To see or not to see: The need for attention to perceive changes in scenes *Psychol. Sci.* 8, 368-373
- 10 Tarr, M.J. and Aginsky, V. (1996) From objects to scenes: Speculations on similarities and differences. Paper presented at the Scene Recognition Workshop, Max-Planck-Institut für Biologische Kybernetik, Tübingen, Germany.
- 11 O'Regan, J.K. et al. (1997) Picture changes during blinks: Not seeing where you look and seeing where you don't look *Invest. Ophthalmol. Vis. Sci.* 38, 5707
- 12 O'Regan, J.K., Rensink, R.A. and Clark, J.J. (1996) "Mud splashes" render picture changes invisible *Invest. Ophthalmol. Vis. Sci.* 37, 5213
- 13 Brewer, W.F. and Treyns, J.C. (1981) Role of schemata in memory for places *Cognit. Psychol.* 13, 207-230
- 14 Friedman, A. (1979) Framing pictures: The role of knowledge in automatized encoding and memory for gist *J. Exp. Psychol. Gen.* 108, 316-355
- 15 Henderson, J.M., Hollingworth, A. and Weeks, P.A.J. (1996) The influence of scene context on object perception. Paper presented at the Scene Recognition Workshop, Max-Planck-Institut für Biologische Kybernetik, Tübingen, Germany
- 16 Pedzek, K. et al. (1989) Memory for real-world scenes: The role of consistency with schema expectation *J. Exp. Psychol. Learn. Mem. Cognit.* 15, 587-595
- 17 Ballard, D.H., Hayhoe, M.M. and Pelz, J.B. (1995) Memory representations in natural tasks *J. Cogn. Neurosci.* 7, 66-80
- 18 Simons, D.J. (1996) In sight, out of mind: When object representations fail *Psychol. Sci.* 7, 301-305
- 19 Pashler, H. (1988) Familiarity and visual change detection *Percept. Psychophysiol.* 44, 369-378
- 20 Phillips, W.A. (1974) On the distinction between sensory storage and short-term visual memory *Percept. Psychophysiol.* 16, 283-290
- 21 Levin, D.T. and Simons, D.J. Failure to detect changes to attended objects in motion pictures *Psychonomic Bull. Rev.* (in press)
- 22 Dmytryk, E. (1984) *On Film Editing: An Introduction to the Art of Film Construction*, Focal Press
- 23 Rothbart, M. and John, O.P. (1985) Social categorization and behavioral episodes: A cognitive analysis of the effects of intergroup contact *J. Social Issues* 41, 81-104
- 24 Intraub, H. (1981) Rapid conceptual identification of sequentially presented pictures *J. Exp. Psychol. Hum. Percept. Perform.* 7, 604-610
- 25 Nickerson, R.S. (1965) Short-term memory for complex meaningful visual configurations: A demonstration of capacity *Can. J. Psychol.* 19, 155-160
- 26 Potter, M.C. (1976) Short-term conceptual memory for pictures. *J. Exp. Psychol. Hum. Learn. Mem.* 2, 509-522
- 27 Shepard, R.N. (1967) Recognition memory for words, sentences, and pictures *J. Verb. Learn. Verb. Behav.* 6, 156-163
- 28 Standing, L., Conezio, J. and Haber, R.N. (1970) Perception and memory for pictures: Single-trial learning of 2500 visual stimuli *Psychonomic Sci.* 19, 73-74
- 29 Henderson, J.M. (1997) Transsaccadic memory and integration during real-world object perception *Psychol. Sci.* 8, 51-55
- 30 Moore, C.M. and Egeth, H. (1997) Perception without attention: Evidence of grouping under conditions of inattention *J. Exp. Psychol. Hum. Percept. Perform.* 23, 339-352
- 31 DeSchepper, B. and Treisman, A. (1996) Visual memory for novel shapes: Implicit coding without attention *J. Exp. Psychol. Learn. Mem. Cognit.* 22, 27-47
- 32 Graf, P. and Ryan, L. (1990) Transfer-appropriate processing for implicit and explicit memory *J. Exp. Psychol. Learn. Mem. Cognit.* 16, 978-992