# Change Blindness

Ronald A. Rensink
University of British Columbia

In L. Itti, G. Rees, and J.K. Tsotsos (eds). *Neurobiology of Attention*. (pp. 76-81). 2005.

San Diego, CA: Elsevier.

## Abstract

Large changes that occur in clear view of an observer can become difficult to notice if made during an eye movement, blink, or other such disturbance. This *change blindness* is consistent with the proposal that focused visual attention is necessary to see change, with a change becoming difficult to notice whenever conditions prevent attention from being automatically drawn to it.

It is shown here how the phenomenon of change blindness can provide new results on the nature of visual attention, including estimates of its capacity and the extent to which it can bind visual properties into coherent descriptions. It is also shown how the resultant characterization of attention can in turn provide new insights into the role that it plays in the perception of scenes and events.

# I. INTRODUCTION

As observers, we have a strong impression that our visual system produces a coherent and detailed description of the world in front of us; a description, moreover, that is always stable and complete. However, various studies have shown that our ability to perceive objects and events in our visual field is far more limited than subjective experience indicates. Among the more striking phenomena in this regard is *change blindness*, the inability to notice changes that occur in clear view of the observer, even when these changes are large and the observer knows they will occur. (For a general review, see e.g., Rensink, 2002).

Change blindness has turned out to be a powerful and robust effect that can be induced in a variety of ways, such as making the change during an eye movement, an eye blink, or a brief flash in the image. The generality of this effect indicates the involvement of mechanisms central to the way that we perceive our surroundings. The determination of these mechanisms and the way they relate to each other is far from complete. But it is clear that visual attention is critical; in particular, results indicate that focused attention is needed for the perception of change (Rensink et al., 1997). Given the strength of its effects and its tight connection with attention, change blindness appears to be a powerful way of exploring the nature of visual attention and the role it plays in our perception of the world.

## A. Basic Distinctions

To avoid the confusions that often hinder investigations into the perception of change, it is useful to first make a few basic distinctions (see Rensink, 2002). One of these is the distinction between *change* and *motion*. As used here, *change* refers to the transformation of an enduring (coherent) structure over time. In contrast, *motion* refers to the temporal variation of some quantity (such as intensity or color) at a fixed point in space. Motion does not involve structure, and motion detectors do not require attention for their operation. As such, the key characteristic of focused visual attention is the creation (and perhaps maintenance) of representations capable of describing coherent spatiotemporal patterns.

Another important distinction is that between the perception of *dynamic* change (i.e., seeing a change as a dynamic visual event) and the inference of *completed* change (i.e., noticing that something has changed at some time in the past, without a phenomenological experience of anything dynamic). During the perception of dynamic change, the spatiotemporal continuity of the internal representation is maintained. In contrast, the perception of completed change does not require such continuity; in principle, it could be carried out simply by a comparison of the currently-visible structure with memory, requiring at most only an intermittent application of attention.

Finally, it is also worth distinguishing between *change* and *difference*. Perception of difference is based on the lack of similarity in properties of two distinct structures. In contrast to change, difference involves no notion of temporal transformation; instead, similarity is defined via atemporal comparison. The question of whether attention is involved in the perception of difference (and perhaps of completed change) then reduces to the question of whether attention is needed for comparison.

## B. Methodological Considerations

The design of any change-detection experiment must provide a way to decouple change, motion, and difference. To decouple change from motion, at least two strategies are possible. First, the change can be made gradually enough that the accompanying motion signal does not draw attention (e.g., Simons et al., 2000). Second, the change can be made contingent on an event (such as a brief flash, eye movement, or occlusion) that creates a global motion signal that can swamp the localized signal associated with the change (e.g., Rensink et al., 1997).

Decoupling change from difference requires separating the effects of visual attention from the effects of long-term memory. One strategy is to have observers detect changes as soon as possible, thereby minimizing the contribution of memory. Another possibility is to have the observer respond differentially to the perception of dynamic change (which presumably relies on attention) and the inference of completed change (which presumably relies on a longer-term visual memory).

Techniques have been developed that incorporate most of these considerations into their design. Two examples are shown in Figure 1. Figure 1a shows the *one-shot* paradigm, in which an image is briefly presented, followed by a brief blank or mask, and then followed by a second display, possibly containing a changed version of the first. Performance here is measured by the accuracy of change detection. Figure 1b shows the *flicker* paradigm, where the two displays continually alternate until the observer reports the presence or absence of the change. The measure here is the time taken to detect the change. Note that these variants correspond to the use of brief and extended displays in visual search experiments on static stimuli [see VISUAL SEARCH], with the target being a spatiotemporal pattern rather than a purely spatial one. (For a more extensive review, see Rensink, 2002.)

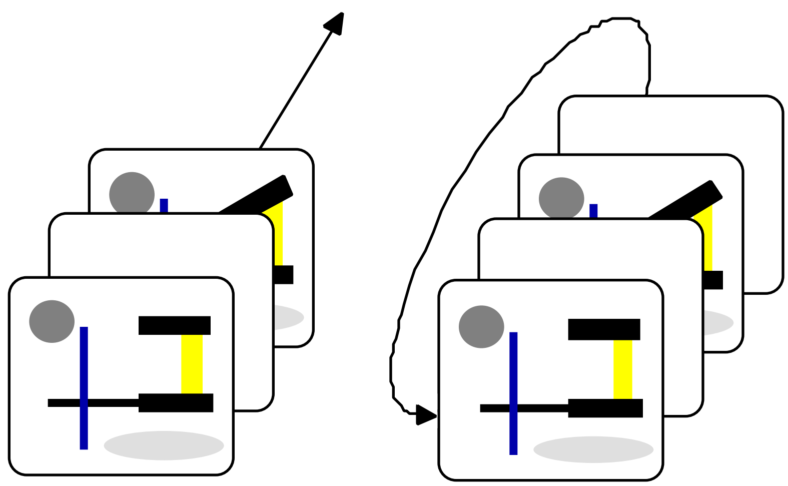


**Figure 1.** Examples of techniques used to induce change blindness. (a) One-shot paradigm. Here, the observer views a single alternation of displays, with a brief blank or mask between them. The task of the observer is to detect (or identify) the change; performance is measured via accuracy of response. (b) Flicker paradigm. Here, the observer views a continual cycling of displays, with a brief blank or mask after each display. The task of the observer is to detect (or identify) the change; performance is measured via response time. Both approaches can also be applied to other kinds of change, such as those made during eye movements or blinks.

# II. THE NATURE OF VISUAL ATTENTION

Before examining how studies of change blindness (and its flip side, change detection) cast light on our understanding of visual attention, it is important to specify what is meant by "attention". This term can refer to several rather different things, and it is not *a priori* evident which would be relevant here.

As it turns out, however, all results point towards the involvement of the *focused attention* believed to bind together properties in a static item (e.g., Kahneman et al, 1992). For example, many characteristics of change detection (such as speed, capacity, and selectivity) are similar to—or at least compatible with—those of focused attention (Rensink, 2002). Also, change blindness is attenuated both for "interesting" items and for cued items (Rensink et al., 1997), both effects being consistent with what is known about the control of focused attention. Furthermore, attentional priming occurs at the location of an item seen to be changing, but not when there is no visual experience of change (Fernandez-Duque & Thornton, 2000).

## A. Capacity Limits and Bottlenecks

Studies based on both one-shot and flicker paradigms show that when observers attempt to detect the *presence* of change, about 4 items can be attended at a time (e.g., Luck & Vogel, 1997; Rensink, 2002). This is similar to the limit found for other kinds of attentional task (e.g., Pylyshyn & Storm, 1988). The extent to which separate components of attention and visual short-term memory are involved is not clear. [See ATTENTIONAL BOTTLENECK AND WORKING MEMORY.]

Interestingly, detecting the *absence* of change among a set of changing items yields a limit of 1 item (Rensink, 2002). One explanation for this is that information from the attended items is in some way pooled into a single collection point, or *nexus* (Figure 2). If the nonchanging items do not contribute to the pooled signal, detecting a single change signal among 4 attended items could easily be done. But if 4 items were attended, it would be difficult to distinguish 4 changing items from 3 changing + one non-changing item; thus, to reliably detect a non-changing target, only one item can be attended at a time. (Note that this explanation is similar to that used to explain search asymmetry, where detecting the presence of a basic property is far easier than detecting its absence. [see VISUAL SEARCH].)

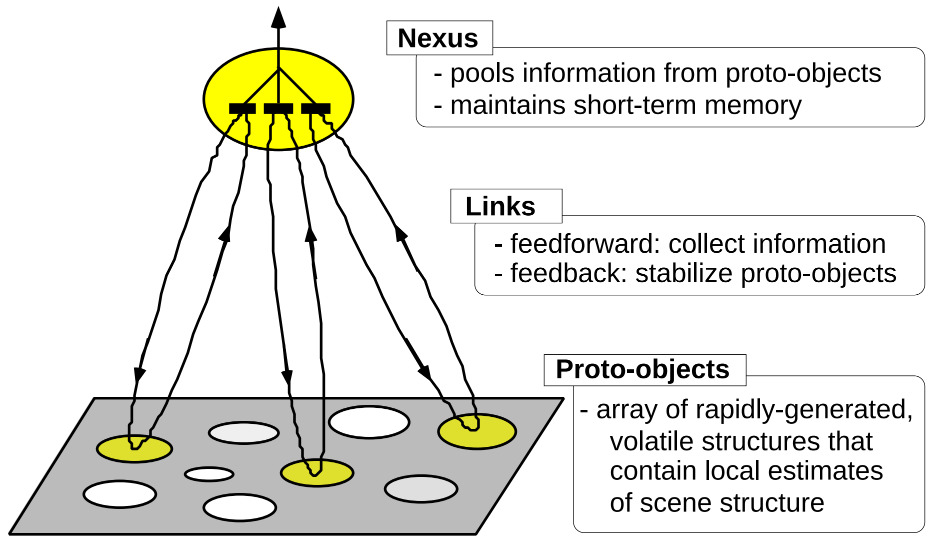


**Figure 2.** Pooling of attended information. Attended items are linked to a single *nexus*. (a) When searching for the presence of change, the nexus signal will either be 1 (target present) or 0 (target absent). A relatively strong signal therefore exists, even when information from several links is collected. (b) When searching for the absence of change, the nexus signal will either be n-1 (target present) or n (target absent). If all items are attended, n would be about 4, and the resulting signal would be quite weak; to obtain a strong signal, the nexus must collect information from only one link at a time.

These limits also cast light on the nature of the bottleneck involved. If 4 items can be held by attention, detecting the absence of change in any of them should be easy: Simply compare each with its counterpart in the image; even if only one comparison can be made at a time, all items could eventually be compared. The finding that this is not possible indicates that the bottleneck is not the number of comparisons that can be made, but rather, the number and nature of the representations constructed.

## B. Independence of Attentional Complexes

Given that several items can be held by attention at any one time, how are the corresponding complexes[1] related to each other? It may be that each is independent of the others (Pylyshyn & Storm, 1988); on the other hand, a higher-level structure may somehow link them, imposing constraints upon their operation (Rensink, 2002).

Results from change-detection studies support the latter view. For example, the constraint that only 1 non-change can be detected at a time would not exist if complexes were independent entities—each complex could simply be tested in turn, leading to a limit of at least 4 items. Additional evidence is the blindness found for switches of colors among tracked items (Saiki, 2003) and for switches of property assignments in static items containing multiple properties (Wheeler & Treisman, 2002), even when only 2 or 3 items were involved. If complexes were independent, and if properties were correctly bound to them, detection of such changes should be easy. The low level of performance actually found is compatible with some migration of properties among the attended items, a natural consequence of a pooled signal

## C. Contents of Attentional Complexes

Another issue of interest is the content of an attentional complex, i.e., the number of basic properties it contains, and the amount of detail for each property. Change-blindness studies indicate that this content is usually sparse, with only a small number of properties represented. For example, observers can miss large changes in an object even when it is attended, suggesting that the corresponding complex may be far from a complete representation of that object (Levin & Simons, 1997). Moreover, it appears that complexes are held in coherent form only as long as they are attended, falling apart when attention is withdrawn (Wolfe, 1999).

At least four simple properties—such as orientation, color, size, and curvature—can be simultaneously represented (Luck & Vogel, 1997), apparently via the concurrent coding of different kinds of properties (Wheeler & Treisman, 2002). Furthermore, such coding can captures not only the properties of each item, also their parts and the structural relations between them (Carlson-Radvansky, 1999). The relation of these complexes to the elements of visual short-term memory

### D. The Binding Problem

One of the more important concerns in the study of attention is the *binding problem*: how to prevent the properties (color, location, etc.) of one object representation from being erroneously assigned to another. [See ATTENTION AND BINDING.] The proposal that attended information is collected into a single nexus—i.e., that only one object is attended at a time—may provide a way out of this: If only one object is represented at a time, there can be no erroneous assignment of properties. Note that this solution would require the ability to construct a new complex for each object as it is needed. As such, the binding problem would be replaced by a *gating problem*: how to select the properties to be entered into the appropriate attentional complex at the appropriate moment in time.

## III. VISUAL ATTENTION AND SCENE PERCEPTION

Although change blindness has provided insights into the nature of visual attention, it has also provided insights into other aspects of visual perception, and the role that attention plays in them.

For example, it is believed that the initial stage of visual perception (*early vision*) involves simple visual elements created rapidly and in parallel across the visual field. Because these elements are believed to have a fleeting existence (existing only as long as light continues to enter the eye), the role of attention was sometimes seen as one of “welding” these elements into complexes that are more durable. These complexes then accumulate, providing a representation that is both dense (i.e. highly detailed) and coherent (i.e., all elements correctly bound together).

But change-blindness experiments show that only a few coherent complexes exist, with relatively little detail in each; as such, no representations exist that are both highly detailed *and* coherent. To reconcile this with the coherent, detailed picture of the world that we experience, it has been proposed that scene perception is based on a sparse, dynamic “just-in-time” system that creates object representations when (and only when) they are needed. If this co-ordination were done correctly, this *virtual representation* would appear to higher-level processes as if "real", i.e., as if all objects simultaneously have detailed, coherent representations (Rensink, 2002).

## A. Triadic Architecture

One possible implementation of a virtual representation is the *triadic architecture* (Rensink, 2002) shown in Figure 3. This is composed of three systems: (i) an early system that continually generates simple visual elements, (ii) an attentional system that enters a subset of these into a coherent representation of an object, and (iii) a nonattentional system that determines such things as the meaning (or *gist*) of the scene, and the spatial arrangement (or *layout*) of items in it.

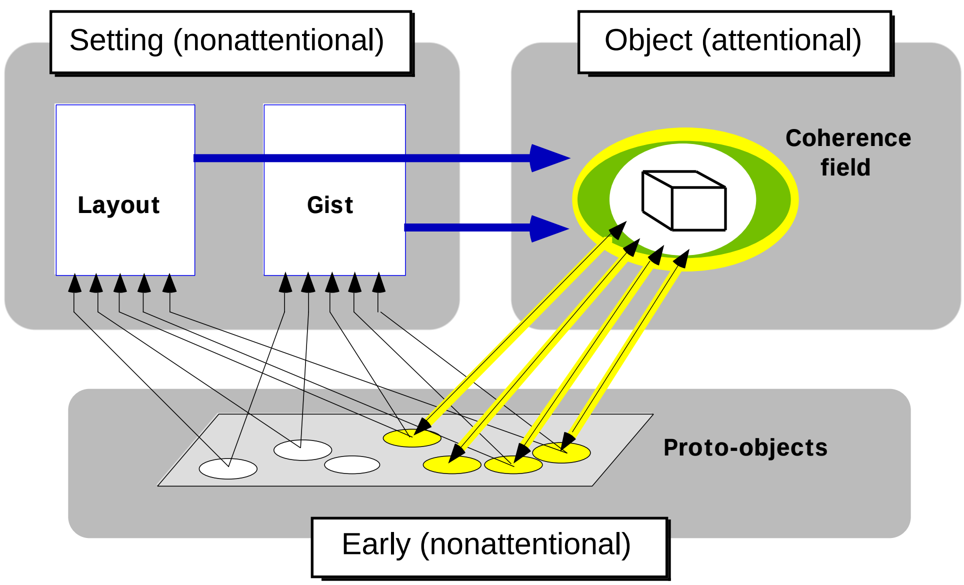


**Figure 3.** Triadic architecture. Thin lines indicate information flow; thick lines control. Here, visual perception is carried out by three largely independent systems: (i) an early system concerned with the formation of (unattended) elements rapidly and in parallel across the visual field, (ii) an object system concerned with the formation of coherent representations (complexes) via attention, and (iii) a nonattentional setting system that enables attentional guidance via high-level knowledge. These enable effective management of attention (and therefore conscious perception) via a combination of high- and low-level control.

Here, the constantly-regenerating elements in the early system provide a rapid estimate of scene gist and layout. Attention is controlled both by high-level considerations (knowledge) and by low-level considerations (salience of individual items) to create representations of the appropriate objects at the appropriate time. The objects so formed could then be used in turn as the basis of further attention guidance. As such, scene perception would involve a continually-circulating flow of information between low-level representations containing retinal input and higher-level representations containing knowledge about the scene [See ATTENTION AND SCENE UNDERSTANDING.]

## B. Observer Intention

Given the dynamic nature of scene representation, perception for a given task must rely on *attentional management*—i.e., deploying attention as effectively as possible. An important factor here is the degree to which the observer expects a change, and believes that reporting it is relevant. The degree of change blindness found is much higher when the observer does not expect a change (being asked to report it afterwards), although some ability to detect change still remains (Levin & Simons, 1997). This supports the view that the only properties put into coherent form (or at least compared) are those needed for the task at hand.

Another important factor is the type of change expected. Detection of orientation change is unaffected by irrelevant variations in contrast sign, again indicating that only those properties needed for the immediate task are encoded (Rensink, 2002). More generally, observers appear to be sensitive only to changes in those properties relevant to the task being carried out at the moment the change was made (see Rensink, 2002).

## C. High-level Knowledge

Attentional management is heavily dependent on the high-level knowledge of the observer. One way this can influence perception is via the particular *representations* available. For example, detection of change was better for objects learned at a specific rather than a general level (Archambault et al., 1999). This suggests that more detailed representations had been formed for the specific-level objects, with observers then taking advantage of these representations to improve their performance.

Another way that knowledge can influence perception is via a more effective *guidance* of attention. For example, a study that compared the performance of experts and non-experts in American football found that experts could spot changes in meaning more quickly, and could attentionally scan meaningful scenes more efficiently (Werner & Thies, 2000).

## D. Vision Outside the Focus of Attention

The proposal that attention is needed to see change implies that change cannot be seen outside the focus of attention. But this proposal was based on studies where observers made a volitional response; as such, the meaning of “see” must be restricted to conscious visual experience (Rensink, 2002). If other forms of response are considered, it may be that some other form of change perception is possible, perhaps mediated by the nonattentional streams proposed in the triadic architecture (section II.A).

Several change-blindness studies show interesting results in this regard. For example, even if an observer does not consciously experience a change, their visuomotor systems can still respond to it (e.g. Bridgeman et al 1979). Furthermore, observers without any visual experience of a change can guess its location with above-chance accuracy (Fernandez-Duque & Thornton, 2000). However, in such cases the underlying mechanisms are poorly understood, and it is not entirely certain that the involvement of focused attention can be ruled out. Moreover, there remains the possibility that while a considerable subset of properties may be attended (and so affect subjective experience), only a subset of these may be compared on a given task (and so affect objective performance). Further work will be needed to clarify these issues (see Rensink, 2002).

# IV. CONCLUSIONS

Visual attention appears to be critical for the creation (and perhaps maintenance) of internal representations with a spatiotemporal coherence that in some sense matches that of the external object(s) they describe. Change blindness reflects the ability of visual attention to create (and perhaps maintain) such representational structures. Results to date on the nature and role of attention are consistent with—and in places extend—results obtained using other approaches.

Looked at more broadly, the study of change blindness is the first stage of investigation into the more general issue of the perception of organized spatiotemporal patterns, such as movements and events. Based on the results obtained so far, it is likely that the perception of such patterns will critically depend upon visual attention.

## References

Archambault A, O'Donnell C, Schyns PG. 1999. Blind to object changes: When learning the same object at different levels of categorization modifies its perception. *Psychological Science*, **10**: 249-255.

Bridgeman B, Lewis S, Heit G, and Nagle M. 1979. Relation between Cognitive and Motor-Oriented Systems of Visual Position Perception. *Journal of Experimental Psychology: Human Perception and Performance*, **5**: 692-700.

Carlson-Radvansky LA. 1999. Memory for relational information across eye movements. *Perception & Psychophysics*, **61**: 919-934.

Fernandez-Duque D, Thornton IM. 2000. Change detection without awareness: Do explicit reports underestimate the representation of change in the visual system. *Visual Cognition*, **7**: 323-344.

Kahneman D, Treisman A, Gibbs B. 1992. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, **24**: 175-219.

Levin DT, Simons DJ. 1997. Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review*, **4:** 501-506.

Luck SJ, Vogel EK. 1997. The capacity of visual working memory for features and conjunctions. *Nature*, **390:** 279-280.

Pylyshyn ZW, Storm RW. 1988. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, **3**: 179-197.

Rensink RA. 2002. Change detection. *Annual Review of Psychology*, **53:** 245-277.

Rensink RA, O'Regan JK, Clark JJ. 1997. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, **8:** 368-373.

Saiki J. 2003. Feature binding in object-file representations of multiple moving items. *Journal of Vision*, **3**: 6-21.

Simons DJ, Franconeri SL, Reimer RL. 2000. Change blindness in the absence of visual disruption. *Perception*, **29:** 1143-1154.

Werner S, Thies B. 2000. Is "change blindness" attenuated by domain-specific expertise? An expert-novices comparison of change detection in football images. *Visual Cognition*, **7:** 163-173.

Wheeler ME, Treisman AM. 2002. Binding in short-term visual memory. *Journal of Experimental Psychology: General*, **131**: 48-64.

Wolfe JM. 1999. Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting Memories*. Cambridge, MA: MIT Press. (pp. 71-94.)

## Footnotes

1. A variety of terms have been used to describe the representational structures formed by focused attention, such as *object file* (Kahneman et al, 1992), FINST (Pylyshyn & Storm, 1988), and *coherence field* (Rensink, 2002). To discuss the results of various change-blindness studies without regard to a particular theoretical framework, the term *complex* will be used to denote the representational structure formed by attention for an item in a stimulus array, without regard to any particular theory.