

Change Detection

Ronald A. Rensink
Departments of Psychology and Computer Science
University of British Columbia
Vancouver BC
Canada

To appear in: Annual Review of Psychology, 53 (2002)

Running Head: Change Detection

Key words: change blindness, visual attention, scene perception,
eye movements, visual memory

Word count (abstract) : 140

Word count (body) : 11011

Author Notes

I would like to thank Dan Simons and Ian Thornton for their comments on an earlier draft of this paper. Correspondence may be addressed to R.A. Rensink, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver BC V6T 1Z4, Canada. (email: rensink@psych.ubc.ca).

Abstract

Five aspects of visual change detection are reviewed. The first concerns the concept of change itself, in particular the ways it differs from the related notions of motion and difference. The second involves the various methodological approaches that have been developed to study change detection; it is shown that under a variety of conditions observers are often unable to see large changes directly in front of them. Next, it is argued that this "change blindness" indicates that focused attention is needed to detect change, and that this can help map out the nature of visual attention. The fourth aspect concerns how these results affect our understanding of visual perception—for example, the proposal that a sparse, dynamic representation underlies much of our visual experience. Finally, a brief discussion is presented concerning the limits to our current understanding of change detection.

Change Detection

INTRODUCTION

Change detection is the apprehension of change in the world around us. The ability to detect change is important in many aspects of our everyday life—for example, noticing a person entering the room, coping with traffic, or watching a kitten as it runs under a table. But in spite of the pervasiveness of change detection in our lives, it has proven surprisingly difficult to study. Only recently have various approaches begun to converge in terms of what it is, and how it is carried out.

As used here, the term change detection will pertain primarily to the visual processes involved in first noticing a change. It will denote not only detection proper (i.e., the observer reporting on the existence of the change), but also identification (reporting what the change is) and localization (reporting where it is). The perception of dynamic patterns per se (e.g., the perception of movement) will not be discussed in detail here, since this involves a formidable set of issues in its own right (see, e.g., Jacobs et al 1988). Likewise, the focus will be on behavioral measures and their interpretation rather than investigations into underlying neural systems.

Restricted in this way, change detection might appear to be a fairly straightforward process. However, empirical studies have repeatedly proven otherwise. For example, we as observers tend to believe we could immediately detect any change in front of us if it were sufficiently large (Levin et al 2000). But this is not so—under a wide variety of conditions we can be amazingly "blind" to changes, failing to see them even when they are large, repeatedly made, and anticipated. (For reviews, see Rensink 2000a; Simons & Levin 1997). This change blindness (Rensink et al 1997) is a striking phenomenon, one that has often served as the "flip side" of change detection: just as our ability to detect change has cast light on some perceptual mechanisms, so has our inability to detect it cast light on others.

The study of change detection can be loosely divided into three phases. The first, occurring roughly between the mid 1950s and mid 1960s, included the work of French (1953) on changes in position in dot arrays and Hochberg (1968) on changes to faces; here, change usually occurred during a temporal gap of several seconds. Meanwhile, studies such as Ditchburn (1955) and Wallach & Lewis (1966) investigated displacements made during an eye movement (or saccade). All studies showed observers to be surprisingly poor at detecting changes made contingent on a gap or a saccade. However, no attempt was made to incorporate the entire set of findings into a systematic framework.¹

The second phase took place during the 1970s, with some work continuing into the 1980s. Examples here include Pollack (1972), Phillips & Singer (1974), and Pashler (1988), who systematically investigated the limits on the detection of gap-contingent changes in arrays of simple figures. This work—along with that on visual integration—formed the basis for the proposal of a limited-capacity visual short-term memory (vSTM). Concurrent with these developments, studies such as Mack (1970), Bridgeman et al (1975), and McConkie & Zola (1979) showed that observers were poor at detecting saccade-contingent change under a variety of conditions. Although a considerable body of knowledge about this phenomenon was eventually gathered (see e.g., Bridgeman et al 1994), no integration with the work on gap-contingent changes was achieved.

The third phase began in the early 1990s and has continued to the present day. Examples are Simons (1996), Rensink et al (1997), and Henderson & Hollingworth (1999a). Work here can be distinguished from previous approaches in at least one of three ways: (i) stimuli are more realistic—e.g., images of real-world scenes or dynamic events, (ii) repeated rather than single changes are used, allowing the use of time as a measure of performance, and (iii) an emphasis on integrating results obtained via different kinds of manipulations—e.g., the identification of transsaccadic memory with vSTM (Irwin 1991), which connected saccade-contingent and gap-contingent change blindness. These studies emphasize the idea that change detection is not a marginal process, but involves mechanisms central to the way that we perceive our world.

DISTINCTIONS AND DEFINITIONS

As with many ideas in psychology, the concept of change is something that initially appears unproblematic, but upon closer examination contains subtleties that may cause great confusion unless carefully handled. Indeed, the concept of change has had a particularly long and tortuous history, extending back at least to the time of the pre-Socratic philosophers, where it formed the basis for a host of problems. Some of these problems—such as the paradoxes of Zeno of Elea—were highly intricate, taking centuries to resolve² (see e.g., Kirk & Raven 1957); indeed, several persist to this day.

Consequently, this section will focus on some of the basic issues pertaining to the concept of change. In particular, it will attempt to clarify the distinctions (and relations) between change, motion, and difference. These distinctions will be examined both from the perspective of physical description and the perspective of perceptual mechanism.

Change vs. motion

The word change generally refers to a transformation or modification of something over time. As such, this notion presumes a non-changing substrate on which changes are imposed. More precisely, change will be defined here as the transformation over time of a well-defined, enduring structure. The complexity of the structure does not matter—it can range from an undifferentiated particle to a highly articulated object. All that is required is that the structure continue to exist over the course of its transformation. Although such continuity can be defined in a variety of ways, the most appropriate appears to be spatiotemporal continuity (see e.g., Smith 1998).

Stated this way, an important distinction can be drawn between change and motion. Motion is often taken to refer to change of position over time. But consider a situation such as a flowing stream. Here, the critical property is the velocity at each point in space; the complete array of these velocities forms a motion field. Referencing temporal variation to space in this way allows motion to be treated much like color or brightness, so that constructs such as borders can be defined on the basis of the motion pattern. More generally, motion can refer to the temporal variation at a point in space of any measurable quantity (see e.g., Adelson & Bergen 1991).

Thus, motion is most usefully defined as variation referenced to location, whereas change is referenced to structure.³ This distinction has important consequences for the perceptual processes involved. For example, motion can generally be described in terms of local derivatives—no other structure is needed.⁴ Motion detectors can therefore be located at the initial stages of visual processing, where spatial representations have minimal complexity (see e.g., Hildreth & Koch 1987; Nakayama 1985). In contrast, change is referenced to a particular structure that must maintain spatiotemporal continuity, and so more sophisticated processes are needed.

In this view, then, the transformation of any external entity is picked up by two concurrent perceptual systems: one describing motion (variation referenced to location), and the other change (variation referenced to a structure).⁵ Although the outputs of these systems are often correlated, they can on occasion be decoupled, e.g., by their differential response to different stimuli (Seiffert & Cavanagh 1998), or by using transformations that occur beyond the temporal window of most motion detectors, a window of approximately 50-80 ms (e.g., Woodhouse & Barlow 1982; van der Grind et al 1986).

Dynamic vs. completed change

Another important distinction is that between the detection of dynamic change (i.e., seeing a change in progress) and the detection of completed change (i.e., seeing that something has changed). Loosely speaking, this distinction reflects the difference between present progressive and past perfect tense.

More precisely, the detection of dynamic change refers to the perception of the transformation itself—the change is perceived as a dynamic visual event. This suggests that the spatiotemporal continuity of the external entity may be reflected in the spatiotemporal continuity of the internal representation. Note that the entity need not be continually present, however—mechanisms at early levels may allow its representation to be sustained over brief intervals.

In contrast, the detection of completed change refers to the perception that the structure changed at some point, such as might happen if the change took place while the entity was briefly occluded. Phenomenologically, there is no sense of a dynamic transformation—the change is simply noted to have taken place some time in the past. The mechanisms involved here would likely compare a property of a representation in memory against a representation of a currently-visible structure, with the continuity of the referents established by means other than the continuity of the visual representation itself.⁶

Change vs. difference

Finally, it is also important to distinguish between change and difference. As discussed above, change refers to the transformation over time of a single structure. In contrast, difference refers to a lack of similarity in the properties of two structures. The issue then is to clarify how these two notions differ. To the degree that they are not the same, trying to "spot the difference" between two side-by-side images will be a rather different activity than trying to detect the change in a pair of sequentially-presented images (cf. Gur & Hilgard 1975; Brunel & Ninio 1997; Scott-Brown et al 2000; Shore & Klein 2000).

It is important to note that the concepts of change and difference do have several elements in common. Both are referenced to structure, with the nature of the structures being unimportant. And both rely on the idea of similarity as applied to one or more of their properties.

But although they are somewhat alike, the two concepts are not the same. To begin with, change refers to a single structure; difference to two. Furthermore, change involves temporal transformation, the measures of similarity pertaining to the same structure at different points in time; this is especially pronounced in dynamic change. In contrast, difference involves no notion of transformation, with similarity defined instead via the atemporal comparison of structures that may or may not exist simultaneously. There consequently appears to be an ordering of sorts: (i) dynamic change, with dynamic transformation and spatiotemporal continuity, (ii) completed change, with inferred transformation (or comparison) and possibly a more abstract kind

of continuity, and (iii) difference, with no transformation (only comparison) and no continuity.

As in the case of motion, these distinctions have implications for the perceptual mechanisms involved. The detection of dynamic change involves spatiotemporal continuity, not only of the external entity, but likely the internal representations as well. As such, a memory of a fairly sophisticated sort is required, one that not only maintains continuity but perhaps also supports the perception of a dynamic transformation. In contrast, all that is needed to detect difference is to extract the relevant properties from each entity and compare them at some point; memory need not enable continuity to be determined. The mechanisms underlying the detection of completed change depend on the kind of continuity involved. If the external entity is somehow tracked, these mechanisms might be largely the same as those used for dynamic change. If the entity is not tracked, the mechanisms could be much the same as those used for detecting difference, along with an additional mechanism to identify the two structures as the same entity at different moments in time.

It is also important to maintain a careful distinction between structures defined as external entities and structures defined as internal representations. For example, a person may be encountered on two different occasions years apart. If the external entity is considered, there exists a set of transformations linking the current manifestation to the previous one; it is therefore the same person, who has simply changed. But suppose the observer previously knew the person, but does not now recognize them. In terms of internal representations, the old and the new manifestations will be seen as different people.

These distinctions not only affect the way that words are used—they also affect the types of experiments considered relevant to a discussion of change detection. For instance, "match-to-sample" experiments (e.g., Biederman & Gerhardstein 1993) do not pertain to change detection, for the test and sample items are not necessarily perceived as corresponding to the same entity. Meanwhile, some "same vs. different" studies (e.g., Carlson-Radvansky & Irwin 1999) do pertain to change detection, since successive stimuli are perceived as the same entity.

EMPIRICAL APPROACHES

As mentioned earlier, the study of change detection has evolved over many years, proceeding through phases that have emphasized different types of stimuli and different types of tasks. All studies, however, rely on the same basic design: an observer is initially shown a stimulus (e.g., a picture or array), a change of some kind is made to this stimulus (e.g., removal or alteration of an element), and the response of the observer is then measured. Furthermore, the wide variety of approaches that have been

developed around this design can be categorized via a relatively small number of dimensions.

Contingency of change

The design of any change-detection experiment must ensure that the results are not due to the detection of motion. Note that the goal is not to eliminate motion detection, for a changing stimulus is always accompanied by temporal variations in the incoming light. Rather, the goal is to decouple the outputs of the change- and motion-detection systems. A few studies (e.g., Brawn & Snowden 1999; Castiello & Jeannerod 1991) attempt this via the temporal pattern of responses to a sudden change. Others (e.g., Seiffert & Cavanagh 1998) look at how performance is affected by different types of stimuli. But for the most part, change and motion have been decoupled by making the change contingent on some event:

Gap-contingent techniques, for example, make the change during a temporal gap (or interstimulus interval) between the original and altered stimulus. A patterned mask is sometimes displayed during this gap, although often a simple blank field is used. Examples of this include Hochberg (1968), Phillips (1974), Pashler (1988), Simons (1996), and Rensink et al (1997). Observers are generally poor at detecting change if more than a few items are present.

Saccade-contingent approaches, meanwhile, make the change during a saccade of the eyes. Examples are Sperling & Speelman (1965), Bridgeman et al (1979), McConkie & Zola (1979), Carlson-Radvansky & Irwin (1995), Grimes (1996), and Henderson & Hollingworth (1999a). In all cases, observers are generally poor at detecting change. Indeed, this is true for position change if even just one item is present, provided it has no global frame of reference.

Shift-contingent techniques make the change during a sudden shift of the entire display; in a way this is like the saccade-contingent technique, but with a simulated saccade. Examples are Sperling & Speelman (1968; see Sperling 1990) and Blackmore et al (1995). A considerable amount of change blindness is found both when the eye does and does not move in response to the shift.

Blink-contingent procedures make a change during an eyeblink. An example of this is O'Regan et al (2000). Again, observers are generally poor at detecting such changes. (This effect was known to movie editors long ago, who would use a sharp noise to induce a blink—see Dmytryk 1984, p. 31). Interestingly, change blindness can occur even if the observer is fixating the item being changed.

Splat-contingent techniques make the change simultaneous with the appearance of brief distractors (or splats). An example of this is Rensink et al (2000b). The change blindness induced by this technique is less severe than that of others; nevertheless, it is

still found (O'Regan et al 1999; Rensink et al 2000b). This shows that change blindness can be induced even when the change itself is completely undisturbed.

Occlusion-contingent change occurs while the changing item is briefly occluded. Examples are Simons & Levin (1998), Vaughan & Yantis (1999), Scholl et al (1999), and Rich & Gillam (2000). In all cases, changes are much more difficult to detect than when the changing item is not occluded.

Cut-contingent methods involve items in movies, the change being made during a cut from one camera position to another. Examples of this are Levin & Simons (1997), and Levin & Simons (2000). Changes made this way are usually difficult to detect. An interesting exception is that location change can be easily detected if the left-right arrangement of the two leading characters is reversed (see e.g., Arijon 1976, p. 29).

Gradual change, meanwhile, has the transition between original and modified display made slowly, i.e., over the course of several seconds. (Note that the transition must still be fast enough that it can be easily seen once noticed.) An example is Simons et al (2000). Observers have great difficulty detecting this kind of change, even though no disruptions of any kind appear in the display.

Repetition of Change

Studies can also be characterized by the number of times the change is made; this is roughly analogous to the duration of a static presentation in a conventional detection experiment. As for visual experiments generally, brief and extended presentations are complementary approaches, with the weaknesses of the one largely compensated for by the strengths of the other.

In the one-shot approach, the change is made just once during each trial. Performance is primarily measured by accuracy, although response time (RT) is sometimes measured as well (Figure 1a). Examples of this are Hochberg (1968), Avons & Phillips (1980), Blackmore et al (1995), Levin & Simons (1998), and Wright et al (2000). This technique minimizes the involvement of eye movements and long-term memory. It can also distinguish between particular transformations (e.g., blue to yellow, present to absent), which would not be easily separated otherwise.

The repeated-change approach, in contrast, has the change made repeatedly until detected, or until the trial ends. Performance is primarily measured via RT, although accuracy is sometimes measured as well. Examples are Rensink et al (1997), Hollingworth & Henderson (2000), Aginsky & Tarr (2000), and Wallis & Bühlhoff (2000). In the case of gap-contingent change, the constantly-repeating gaps generate a visible flicker (Figure 1b); this technique is therefore known as the flicker paradigm (Rensink et al 1997). Because this technique allows sufficient time to process the input, the finding

of change blindness here rules out the possibility that it is due to a failure to consolidate the information in memory (Rensink et al 2000b).

----- Insert Figure 1 about here -----

Content of Display

Another dimension is the content of the displays used. As for visual stimuli generally, this can range from simple static figures on computer monitors to dynamic events in the world itself. The level of realism used reflects a particular choice of trade-offs: simpler displays afford more control and usually allow results to be more easily analyzed, whereas more realistic displays involve factors more difficult to compensate for, but more applicable to tasks in everyday life.

Simple figures form the simplest type of display; these are typically dots, lines, or letters arranged in circular, linear, or rectangular arrays. If configural effects can be neutralized, such stimuli are well-suited for investigating individual mechanisms, since the effects of other processes (e.g., scene organization) are minimal. Examples are Phillips (1974), Luck & Vogel (1997), Rensink (2000b), and Scott-Brown et al (2000). In all cases, a considerable amount of change blindness occurs whenever more than a few items are present.

Drawings of objects and scenes range from line-based sketches to full-color computer renderings; these are placed into simple arrays or form complete scenes. Such displays are a step towards greater realism, but without the full complexity of real-world stimuli. Examples can be found in Simons (1996), Henderson & Hollingworth (1999a), Scholl (2000), and Williams & Simons (2000). In general, a high degree of change blindness can be induced, although not always as dramatic as that obtained with more realistic stimuli (Hollingworth & Henderson 2000).

Images of objects and scenes provide even more realism. These are photographs of real-world objects, which—as in the case of drawings—are either placed into simple arrays or form complete scenes. Such stimuli have the advantage of avoiding an artificial parsing of objects into parts; in the case of complete scenes they also avoid an artificial parsing of the scene. Examples are Blackmore et al (1995), Grimes (1996), Rensink et al (1997), Zelinsky (1997) and Ro et al (2001). In all situations, a high degree of change blindness can be induced.

Dynamic displays such as movies also provide a greater degree of realism. Examples are Levin & Simons (1997), Gysen et al (2000), and Wallis & Bühlhoff (2000). Again, observers generally have great difficulty detecting change, especially in objects

irrelevant to the main events in the presentation. Changes to moving objects appear to be detected more easily than changes to stationary objects (Gysen et al 2000).

Real-life interactions provide the highest level of realism. An example of this is Simons & Levin (1998), where an experimenter requested directions from an unwitting observer, with an occlusion-contingent switch of experimenters occurring during the interaction. As in the case of other techniques, observers are poor at noticing the change. Another example is Wang & Simons (1999), where observers detected changes in the layout of a set of real-world objects. Changes were more difficult to detect if this set was rotated during a temporal gap; interestingly, performance was better if the observer was rotated instead, indicating possible involvement of the vestibular system.

Content of Change

Most studies to date have been careful to ensure that changes made to a display do not introduce a radical change in its overall appearance. For example, a change to a real-world image will be such that both the original and modified images are of the same kind of scene. Care is also taken that no anomalies are caused by the alteration (e.g., an automobile floating in the air), for otherwise performance could be influenced by the anomaly, and not the change *per se*.

But even with such constraints, changes can be made in many ways. It is difficult to compare different types of change with each other: performance depends on the magnitude of the change (Carlson-Radvansky & Irwin 1999; Smilek et al 2000; Williams & Simons 2000) and there is no simple way to equate the visibility of different kinds of changes. But a great deal can be learned by examining how change detection is affected by manipulations within each change type.

Of the various types of change, perhaps the simplest is that in the existence of an item, i.e., addition or deletion.⁷ Examples can be found in Rensink et al (1997), Henderson & Hollingworth (1999a), Aginsky & Tarr (2000), and Mondy & Coltheart (2000). Provided that the changed item is unique, some deletions appear to be detected more easily than additions. However, this asymmetry may be eliminated—or even reversed—depending on how the initial and modified images are to be used subsequently, suggesting that detection of existence change may be based on a feature-matching process (Agostinelli et al 1986; Mondy & Coltheart 2000).

Changes can also be made to various properties of an item; these are usually simple features such as orientation, size, shape, or color. Examples are Palmer (1988), Grimes (1996), Simons (1996), and Scott-Brown & Orbach (1998). Three variants are commonly used: change to an item with a property unique in the display, change to an item with a nonunique value, and a switch in properties among two or more items. Detection is best for changes to a unique property, and worst for switches (Saiki 1999;

Rich & Gillam 2000; Wheeler & Treisman 2001). Performance can differ for different properties (Aginsky & Tarr 2000; Rensink 2000b).

Property changes can also be more complex—e.g., disjunction, where one of two possible properties can change. This is as easy as detecting a single property change, indicating that both properties are concurrently encoded (Luck & Vogel 1997; Wheeler & Treisman 2001). Another variant is conjunction, where all items change one of two properties, with the target changing in both. Such changes are extremely difficult to detect, as is the absence of change among changing items (Rensink 1999a, 2001).

Another type of change is that of the semantic identity of an item—by rearranging its parts, for instance, or by substituting an entirely different item altogether. Examples of this can be found in Levin & Simons (1997), Zelinsky (1997), Archambault et al (1999), and Williams & Simons (2000). If the change is such that the overall appearance of the item is maintained (e.g., same overall size or shape), detection of the change is usually quite poor. Provided that visual similarity can be taken into account, such an approach may provide a powerful way to map out the various categories used in visual perception. Related to this is the issue of how the type of change connects to the item being changed—e.g., detecting the displacement of a car along its direction of travel (i.e. back and forth), versus a similar displacement sideways.

Finally, changes can also be made to the spatial arrangement (or layout) of the items in the display. Care must be taken to keep the number of items—and their properties—constant to avoid confounding factors. (For instance, removing an item would result in a change in layout, but would also result in a change in existence.) Examples are French (1953), Pollack (1972), Irwin (1991), Wang & Simons (1999), and Jiang et al (2000). Change blindness is generally found for layout changes involving more than a few items. In some cases, layout change appears to be easier to detect than feature change (e.g., Simons 1996), whereas in others it is more difficult (e.g., Mondy & Coltheart 2000). This divergence may be due to different encoding strategies: in some situations, a group of items might be seen as a single item, with layout change then corresponding to a change in the configuration of its parts.

Observer Intention

Another important dimension is the intention of the observer. Intentions affect the degree to which an observer will expect a change, which in turn can affect the mechanisms used (see Simons & Mitroff 2001). This is important for investigation into the mechanisms involved in everyday vision, which are not usually devoted to the detection of an anticipated change.

At one end of this spectrum is the intentional approach. Here, the observer fully expects a change and devotes all available resources to detecting it; as such, this is a good way to examine perceptual capacities. Examples are Pollack (1972), Pashler (1988), Jiang et al (2000), and Wright et al (2000). Change blindness often occurs under these conditions, even though all resources have been allocated to the task.

A less-extreme variant is the divided-attention approach, where some other task is primary—e.g., memorization of an image for a subsequent memory task. Meanwhile, observers are told that changes may "occasionally" occur, and to report any changes that they notice. Examples of this are Grimes (1996) and McConkie & Currie (1996). Change blindness is again found, although detection still occurs for changes to certain items, such as saccade targets (McConkie & Currie 1996).

At the other end of this spectrum is the incidental approach. Here, there is no mention at all of a possible change—observers are typically given some other task as their primary responsibility, and questioned only afterwards about whether they noticed a change. Examples can be found in Levin & Simons (1997) and Rich & Gillam (2000). The degree of blindness encountered is generally much higher than that found using intentional approaches. However, some ability to detect change remains.

Type of Task

Tasks can also be classified according to which aspect of the change is involved. As for perception generally, detection is not necessarily identical to localization, or to identification (see e.g. Henderson 1992; Pashler & Badgio 1987). Comparing performance for different types of task may therefore cast light on the various mechanisms at play.

Detection is perhaps the more widely-used type of task; here, the observer simply responds to the presence a change in the display. Examples are Phillips (1974), Verfaillie et al (1994), McConkie & Currie (1996), Luck & Vogel (1997), Austen & Enns (2000), and Scott-Brown et al (2000). A high degree of change blindness is generally found.

Localization, meanwhile, requires that the observer respond to the location of the change. Examples can be found in Fernandez-Duque & Thornton (2000), Scott-Brown & Orbach (1998), and Smilek et al (2000). In most cases, observers are relatively poor at determining the location of the change.

Identification is potentially more complex, requiring the observer to respond to the identity of the changing item. Two variants exist: identity of the changing item, and identity of the change itself (i.e., the type of change). Most work to date has been based on identity of the change, although a few studies have also looked at identity of the item. Examples of this type of task are Palmer (1988), Brawn & Snowden (1999), Wilken

et al (1999) and Mondy & Coltheart (2000). Results all indicate that identification is more difficult than detection of change.

Type of Response

An important aspect of any psychophysical experiment is the set of mechanisms generating the observer's response. These are determined both by the instructions to the observer (i.e., what aspect of their experience to respond to) and by the choice of system to regard as "responding". Different mechanisms may tap into different perceptual systems; consequently, a great deal may be learned from the various types of response to a change.

Explicit responses are triggered by the conscious visual experience of the observer; these most closely match the intuitive idea of "seeing" a change. Examples are French (1953), Simons (1996), Rensink et al (1997), and Jiang et al (2000). Two variants exist: "yes/no", in which the observer must answer either "yes" or "no" (or more generally, select from a fixed set of alternatives), and "go/no-go", in which the observer only responds "yes" if the change is seen (or if one particular type or location of change is seen). The go/no-go response has the advantage of guaranteeing that the observer truly did experience the change, but the response bias differs from that of yes/no responses (Wilken et al 1999). A high degree of change blindness is still found in all cases.

Semi-explicit responses are similar, except triggered by a "feeling" that a change is occurring—no visual experience is involved. This trigger is therefore explicit in some ways (i.e., an awareness that a change is occurring) but not others (i.e., what the change looks like). An example of this is Rensink (1998); see also Rensink (2000a). Although detection of change via this mindsight is not immediate, it can precede the visual experience of change by several seconds in approximately one-third of observers.

Implicit responses, in contrast, are measured by the extent to which a change not consciously perceived can influence a consciously-initiated decision—e.g., how the occurrence of an unseen change affects forced-choice guessing about its possible location. Examples of this are Fernandez-Duque & Thornton (2000), Thornton & Fernandez-Duque (2000), and Williams & Simons (2000). Although implicit detection of change is generally poor, it is above chance levels, suggesting that the underlying mechanisms may provide some information about the location and nature of the item changed.

Visuomotor responses are based on the reaction of a visually-guided motor system—usually manual pointing or eye fixation—to a change in the display. (Note that conscious perception of the change must not be able to influence this response.) Examples are Bridgeman et al (1979), Goodale et al (1986), Castiello et al (1991), and

Hayhoe et al (1998). Visuomotor responses are generally faster and more accurate than consciously-mediated ones, especially when the display contains only a few items; this is particularly true for location change, although an ability to detect property change also exists (Hayhoe et al 1998).

ATTENTION AND CHANGE DETECTION

As the previous section shows, a striking blindness to change can be induced under a wide variety of conditions. The sheer range of these conditions—together with the strength and robustness of the effect itself—indicates that the mechanisms involved are central to the way we perceive the world around us. But what might these mechanisms be?

Most results to date can be explained by the thesis that focused attention is needed to see change (Rensink et al 1997). A change in the world is always accompanied by a motion signal in the input; under normal circumstances, this signal will be unique—or at least larger than the background noise—and so attract attention to its location (see, e.g., Klein et al 1992). This in turn will enable the change to be seen. But if this signal is too weak (e.g., is made too slowly, or is swamped by transients associated with a saccade, flicker, or splat) it will not draw attention, and change blindness will result.

If this view is correct, the possibility arises of reversing things—instead of using focused attention to help clarify what change blindness is, change blindness could be used to help clarify what focused attention is. Much has been learned of attention using approaches such as visual search on static displays, and priming (see e.g., Pashler 1998). However, the high degree of change blindness found under many conditions might be harnessed to provide results with a high signal-to-noise ratio, which could allow the mechanisms of focused attention to be explored in great detail.

Involvement of Focused Attention

Given the thesis that change detection is mediated by "attention", it is important to specify exactly what is meant by this term, since several different meanings can be ascribed to it (see e.g. Allport 1992). In particular, it is important to determine whether the detection of change is mediated by the focused attention believed to bind together features in the perception of static displays (e.g., Treisman & Gormican 1988).

Evidence is accumulating in favor of this possibility. To begin with, many characteristics of the change-detection process (e.g., speed, capacity, selectivity) are similar to—or at least compatible with—what is known of focused attention (Rensink

2000b). For example, for orientation change no more than 4-5 items can be monitored simultaneously (Rensink 2000b; Rensink et al 2000a), a limit similar to that encountered in other types of attentional task (e.g., Pashler 1988; Pylyshyn & Storm 1988).

In addition, change blindness is reduced for items considered to be "interesting" (Rensink et al 1997), and by exogenous cues at the location of the change (Scholl 2000). In both types of situations, then, performance is consistent with the drawing of focused attention.

Finally, it appears that attentional priming occurs at the location of an item seen to be changing, and that such priming does not occur when there is no visual experience of change (Fernandez-Duque & Thornton 2000). Again, this supports the view that the relevant quantity is the focused attention involved in the perception of static displays.

Comparison vs. Construction

If attention is needed for change detection, how does it operate? At least two possibilities exist. First, attention could construct a limited number of relatively complex structures (e.g., the object files of Kahneman et al 1992, or the coherence fields of Rensink 2000c), with these complexes⁸ then being the basis for change detection. Alternatively, attention may simply enable a limited amount of comparison on an effectively unlimited amount of information (Scott-Brown et al 2000).

Several results argue against this latter proposal. First, it cannot explain the failure to combine the detailed contents of successive fixations (Irwin 1991) or why visual search for a changing item should be difficult (Rensink 2000b): if the detailed contents of successive fixations or displays could be accumulated, a distinctive pattern formed from these should be easy to detect. Second, when both the initial and changed displays are presented for increasingly long durations, a limit is reached in the number of items that can be seen to change orientation at any one time (Rensink 2000b). This would not occur if storage were unlimited, since all stored items could eventually be compared, even by a limited-capacity mechanism. Finally, the comparison of items already in (short-term) memory requires about 20 ms/item (Rensink 1999a); if limits on comparison dominated change detection, search for change would also proceed at this rate. However, it actually proceeds at about 100 ms/item (Rensink 2000b), showing that additional operations are involved. Relatively little of change blindness therefore appears due to a bottleneck in the comparison process. Rather, detection of change apparently involves both the construction and comparison of a limited number of structures.

Construction vs. Memory

If attention forms complexes capable of supporting change detection, an important issue is how these are stored in vSTM. A commonly-held view is that focused attention and vSTM are largely separate, with attention constructing complexes and vSTM then maintaining these. However, results on change detection are beginning to alter this picture.

For example, the absence of change is extremely difficult to detect—only one item can be compared at each alternation (Rensink 1999a, 2000a). If several items could be placed into vSTM, this difficulty should not exist, for each item could be examined in turn against the input. Conjunctions of changes are also extremely difficult to detect, also being limited to one per alternation (Rensink 2000a, 2001); again, this should not happen if several complexes could be held in vSTM and compared (one at a time, if need be) against the input. Finally, detection of change to a conjunction of features is improved when all items are removed except the test item, again indicating that items are not held in vSTM in bound form (Wheeler & Treisman 2001).

These results show that focused attention and vSTM overlap much more than previously believed, a conclusion also arrived at from other areas of study (see e.g., Cowan 1988). Indeed, the two may simply be different aspects of the same process, with items held in a coherent complex as long as they are attended, this complex then falling apart when attention is withdrawn (Wolfe 1999; Rensink 2000c).

Independence of Attentional Complexes

Several studies (e.g., Luck & Vogel 1997; Rensink 2000b) show that several items can be held by attention at any one time. But how are the corresponding complexes related to each other? One possibility is that each is completely independent of the others (Pylyshyn & Storm 1988). Alternatively, a higher-level structure may constrain what can be done with them (Yantis 1992; Rensink 2000a, 2001).

Results from change-detection studies support the latter view. In particular, it appears that although attention can hold on to 4-5 items at a time, there is some pooling of their properties into a single collection point, or nexus (Rensink 1999a, 2000c). The number of items attended therefore will depend on the nature of the task. Detecting the presence of change, for example, can be done by attending to several items at a time: all that is needed is a change in the pooled signal. In contrast, determining the absence of a change involves detecting the presence of a non-changing signal in a signal collected from several changing items; to do this reliably, items must be attended one at a time (see Rensink 2000a, 2001). Note that this constraint would not exist if complexes were independent entities—in such a case, each complex could simply be tested in turn.

Other studies also indicate that complexes are not independent. For example, a high degree of change blindness exists for switches of colors in tracked items (Saiki 1999; Scholl et al 1999) and for conjunctions of features in stationary items, even when only three items are in the display (Wheeler & Treisman 2001). In both situations, the relevant items are almost certainly attended; if the corresponding complexes were independent, detection should be near-perfect. The poor performance actually found indicates a migration of features among the attended items, something easily explained by a pooled signal. Furthermore, perceiving the identity of a change (e.g., bigger vs. smaller) is more accurate when attention is given to one item than to four, again indicating a pooling of attentional resources (Palmer 1990).

Contents of Attentional Complexes

Another important issue is the content of an attentional complex—the number of features included, the amount of detail for each feature, etc. Previous work (e.g., Kahneman et al 1992) indicated that this content is relatively sparse, with only a handful of features represented. Results from change-detection studies have reinforced this conclusion. For example, observers in an incidental change-detection task can miss relatively large changes in an item even when it is attended, suggesting that the corresponding complex may be far from a complete representation of the object (Levin & Simons 1997).⁹

Thus, attention may not be concerned with the construction of general-purpose representations, but rather, with the construction of more specialized representations suitable for the task at hand. It appears that at least four properties—such as orientation, color, size, and curvature—can be simultaneously represented in a complex (Luck & Vogel 1997), presumably via concurrent systems (Wheeler & Treisman 2001). Furthermore, items can be described not only in terms of their properties, but also their parts and the structural relations between these parts (Carlson-Radvansky & Irwin 1995). Such relations appear to be encoded via a concurrent set of spatial representations, the simplest of these describing the spatial configuration (or layout) of the attended items (Jiang et al 2000).

Observers are generally good at selecting the properties to be entered into a complex while screening out others. For example, the detection of orientation change is almost entirely unaffected by irrelevant variations in contrast sign (Rensink 2000b). Selection is possible for structural relations as well—for instance, observers can focus on a particular level (local or global) of a compound item, although this appears to be limited to one level at a time (Austen & Enns 2000).

Coherence Theory

The characteristics described above provide a necessary set of constraints on any model of visual attention. One model consistent with all these characteristics is coherence theory (Rensink 2000c). This proposal (Figure 2) has three parts:

- i) Prior to focused attention is a stage of early processing, i.e., processing that is low-level, rapid, and carried out in parallel across the visual field. The resultant structures (proto-objects) can be quite sophisticated, describing several aspects of scene structure. However, they have limited spatial coherence. They also have limited temporal coherence—they are volatile, and so are simply replaced by any new stimuli at their location.
- ii) Focused attention acts as a hand that "grasps" several proto-objects from this constantly-regenerating flux. While held, they are part of a coherence field representing an individuated object. This field is formed via feedback between low-level proto-objects and a mid-level nexus. The coherence formed this way allows the object to maintain continuity across brief interruptions; as such, it is perceived as being transformed when new stimuli arrive at its location.
- iii) To release attention, the feedback loop is broken. The field then loses its coherence, with the object representation dissolving back into its constituent set of volatile proto-objects.

----- Insert Figure 2 about here -----

Note that in contrast to previous models of attended complexes (e.g., the object files of Kahneman et al 1992), coherence theory posits that these structures (the coherence fields) collapse as soon as attention is withdrawn. Furthermore, these complexes are not independent, but are parts of a more integrated structure which loosely corresponds to a concrete instance (or token) of an articulated object.

IMPLICATIONS FOR VISUAL PERCEPTION

According to the arguments put forward in the previous section, attention is needed to see change. But although this thesis—and the developments stemming from it—can explain how observers (do or don't) see change under various conditions, the adoption of this thesis has a number of implications for various aspects of visual perception. For example, if only a small number of objects have a coherent representation at any one time, why do we experience all the objects in our surroundings as existing simultaneously? If we can implicitly perceive change in the

absence of attention, how might this be carried out? And how might it contribute to our visual experience? In general, then, work on change detection is forcing a reconsideration of what it means to see, and what might be the mechanisms involved.

Visual Buffer

Among the first casualties of the thesis that attention is needed to see change was the hypothesis of a visual buffer, a spatiotopic memory store believed to accumulate the contents of successive fixations (e.g., Feldman 1985). This was thought to provide a detailed representation of the incoming light independent of eye movements; among other things, it was taken as the basis for the richly-detailed experience we have of our surroundings. But the widespread existence of change blindness suggests that no such buffer exists (see Comparison vs. Construction, above), a conclusion consistent with complementary work on the visual integration of information (e.g., Irwin 1996; Henderson 1997). More generally, it appears that the detailed contents of successive presentations—including successive fixations—can never be added, compared, or otherwise combined in their entirety, thereby ruling out any large-scale accumulation of information.

It should be pointed out, however, that there still exists a detailed retinotopic representation of the incoming light at any moment in time—indeed, it is exactly this that constitutes the output of early vision (see, e.g., Rensink 2000c). But according to coherence theory this representation does not endure: unattended proto-objects are quickly replaced whenever the eye (or the visual stimulus) moves, and rapidly decay when the eye closes (or the stimulus disappears).

The memory trace of the decaying proto-objects may correspond to the informational persistence (or iconic memory) found using partial-report techniques (e.g., Sperling 1960). Although such memory persists for only about 300 ms, it may nevertheless enable a change in a proto-object to be represented without attention. According to coherence theory such an unattended change will not be consciously perceived. But it could have other effects—for example, it could be salient, drawing attention to its location sooner than otherwise would have been the case (Rensink et al 2000a; Smilek et al 2000). As such, informational persistence may play a larger role in visual perception than previously believed (cf. Haber 1983). One interesting possibility is that it enables the dynamic transformation that is experienced in dynamic change.

Coherent Structure

Various studies of gap-contingent change (e.g., Pashler 1988; Luck & Vogel 1997; Rensink et al 2000a) indicate that only about 4 items can be monitored at a time. This limit is similar to that obtained from work on saccade-contingent change, and indicates that transsaccadic memory may be largely—if not entirely—identical to vSTM (Irwin 1991). If vSTM and visual attention are also much the same, there may exist only one system concerned with the formation and maintenance of coherent visual structure. According to coherence theory (Rensink 2000c), this system would be primarily concerned with the perception of objects.

A key issue in determining the viability of this view is the fate of an attentional complex once attention has been withdrawn. Results suggest that the binding of features into coherent complexes does fall apart in the absence of attention (see Construction vs. Memory, above), supporting the idea that vSTM and visual attention are largely the same, and thus supporting the proposal of a single system for the formation and maintenance of coherent structure.

Scene Representation

The results of various studies of change detection (and related phenomena such as visual integration) appear to converge on two main points concerning visual representation: (i) if it is detailed, it cannot be coherent to any great extent, and (ii) if it is coherent, it cannot be highly detailed. Thus, no visual representations are both coherent and detailed.

To reconcile this with the coherent, detailed scene we believe we experience, consider what needs to be represented for a task. There is usually little need for a detailed representation of all the objects present; instead, all that is really needed is a representation of only those objects—and those particular properties—involved in the task at hand. If attention can form a coherent representation of any aspect of any object whenever requested, the result will be a virtual representation of the scene that will appear to higher levels as if "real", i.e., as if all objects simultaneously have detailed, coherent representations (Rensink 2000c).¹⁰ In this view, the notion of a static, all-purpose representation is replaced by that of a dynamic representation highly sensitive to the demands of the task and the expectations of the observer.

It is not yet known whether such a system is actually implemented in the human perceptual system. However, one possibility that is consistent with what is known of human vision is the triadic architecture (Rensink 2000c) shown in Figure 3. This is composed of three largely independent systems:

- i) Early processing. A low-level system that continually generates highly-detailed, volatile structures. (This could be e.g., the proto-objects posited by coherence theory.)
- ii) Object system. A limited-capacity attentional system that stabilizes these and forms them into a coherent object representation. (This could be e.g., the coherence field posited by coherence theory.)
- ii) Setting system. A limited-capacity nonattentional system that helps guide attention. This may be based on the meaning (or gist) of the scene, and the arrangement (or layout) of items in it. These attributes are based on properties obtained via early vision, largely without the involvement of attention.

----- Insert Figure 3 about here -----

Here, the constantly-regenerating set of proto-objects provides a rapid estimate of the visible properties of the scene. These then form the basis of a rapid determination of scene gist. This—together with layout information—then allows the observer to verify whether the initial impression was correct, and to gather additional detail and construct coherent representations whenever required. Note that such a process of sparse construction and verification is believed to underlie much of the conscious perception of scenes (Henderson 1992; Henderson & Hollingworth 1999b)

Dependence on Task

If the formation of coherent structure requires attention, successful perception for everyday tasks must depend on attentional management, i.e., deploying attention such that the limited amount of information stabilized is used as effectively as possible. An important part of this is to allocate attention in a way suitable for the task at hand.

Evidence of this kind of management appeared in a study which examined search for change in an array of figures (Rensink 1999b). When observers searched for an orientation change, speed was strongly influenced by the shape of the items. But when the same figures were used in a search for a contrast polarity change, no such influence was found. Apparently, some geometric information did not enter into the representations formed when only nongeometric (i.e., color) information needed to be compared.

Similar effects appeared in a study comparing incidental and intentional detection of changes made during movie cuts (Levin & Simons 1997). Intentional detection was found to be much more accurate than incidental detection, indicating that

the only properties put into coherent form (or at least compared) were the minimal ones sufficient for the immediate needs of the observer.

High-level Knowledge

Another important aspect of attentional management involves the high-level knowledge of the observer. Such knowledge often takes the form of a chunk, a structure in long-term memory containing a cluster of properties that can be attended and held in vSTM in much the same way as a simple element (e.g., Miller 1956). The contents of a chunk can affect what is attended, and thus, what is perceived. Evidence of this was found by Archambault et al, (1999), who had observers learn a set of objects either at a general level of categorization (e.g., "a mug"), or at a specific level (e.g., "the mug owned by Bill"). Detection of change was better for objects learned at a specific level, suggesting that more detailed representations had been formed.

High-level knowledge can also have a more direct influence via the selection of the content of attentional complexes. For example, Simons & Levin (1998) investigated detection of changes to people asking for directions. Observers were more likely to notice the change when the experimenter was considered a member of their social group, suggesting a high-level influence on the amount of information contained in the corresponding complexes.

Another aspect is the control of attention by various objects in a scene. Known objects may be rapidly identified via the setting system, and so preferentially draw (or at least hold) attention. For example, changes to real-world scenes are noticed most quickly for central interests, i.e., objects or regions mentioned most often in brief verbal descriptions of the scene (Rensink et al 1997). Although central interests may be preferentially attended in part because of the salience of their features (Shore & Klein 2000), semantic considerations are also involved (Hollingworth & Henderson 2000; Ro et al 2001).

Finally, if the observer has knowledge of which aspects of a scene are important, this can help guide attention to the appropriate items at the appropriate time. For instance, Werner & Thies (2000) examined the ability of observers to notice changes in scenes of American football games. Comparing the performance of experts and non-experts, it was found that experts could spot changes in meaning more quickly, and could attentionally scan meaningful scenes more efficiently.

Implicit Perception

The thesis that attention is needed to see change is somewhat ambiguous about the implicit detection of change that may occur in the absence of conscious awareness (see Type of Response, above). If "see" denotes any use of light to affect behavior (cf. Rensink, 2000a), this thesis makes the strong claim that no detection of change can occur without attention; if attention is sufficient for conscious experience, this would then rule out the very possibility of implicit change detection (a result in accord with some views on implicit perception, e.g., Dulany 1997). On the other hand, if "see" is restricted to conscious visual experience, then the thesis says nothing about how the implicit perception of change may be carried out.

Owing to various methodological difficulties, relatively little is known about this type of perception generally (see, e.g., Milner & Goodale 1995; Merikle & Reingold 1992). Several studies have examined the response of the visuomotor system to various types of change. Results show that even if the observer does not consciously experience a change, the visuomotor system may still respond to it (e.g., Bridgeman et al 1975; Goodale et al 1986; Hayhoe et al 1998). This has provided much of the support for the proposal that vision is composed of two largely independent streams: an on-line stream concerned with immediate visuomotor action, and a slower off-line stream concerned with object recognition (Milner & Goodale 1995). But the particular mechanisms underlying visuomotor change detection (including any possible involvement of attention) remain unknown.

Another approach is the forced-choice guessing of unseen changes. For example, observers may be able to guess the location of a change more often than chance, even if they have no awareness that it occurred (Fernandez-Duque & Thornton 2000). Note that the key issue here is not whether detection is "really" implicit, but whether performance accompanied by conscious awareness differs from performance unaccompanied by awareness (Merikle & Reingold 1992). Some evidence (Thornton & Fernandez-Duque 2000) suggests that this is the case, but there is some disagreement on this point (Mitroff & Simons 2000). If implicit detection of change does exist, the neural substrate might be the on-line stream believed to support visuomotor action; alternatively, it might be the setting system of the triadic architecture (Rensink, 2000c), a nonattentive system not concerned with conscious perception per se.

It may also be possible to obtain information about these mechanisms via explicit detection of change. For example, Deubel et al (1996) found that detection of saccade-contingent change improved greatly when the saccade target was blanked and did not reappear until a few hundred milliseconds after the eye had landed. This indicates that some information about position (and perhaps shape) is retained in a latent form not available to conscious perception. An interesting possibility is that this may be the information supporting implicit detection, with the reappearance of the target somehow triggering its transfer to a more explicit system.

OPEN ISSUES

As the previous sections have shown, much has been learned about the nature of change detection, both what it is and what mechanisms may be involved. And this in turn has cast considerable light on how the visual system operates. However, many issues concerning the nature of change detection still remain open. Some of these involve the properties of the particular mechanisms used—e.g., the capacity of attention/vSTM, the types of units used, etc. (see e.g., Rensink 2000b). But others are more fundamental and require resolution before change detection itself can be said to be well understood. These include:

1. How does identification of change differ from detection of change? Several studies have shown that identifying a change is more difficult than detecting it (see Type of Task, above). This may be in part because identification is simply a less sensitive process, requiring a higher signal-to-noise ratio. However, different mechanisms also appear to be involved (Wilken et al 1999). If this does turn out to be the case, it is not clear which ability—if either—should be considered as "basic"; for example, detection may occur via simple comparison, whereas identification may occur via mechanisms underlying the perception of dynamic events. More generally, different mechanisms may be used for various aspects of change perception, including not only detection and identification, but other tasks as well (e.g., localization).
2. How does implicit detection of change differ from explicit detection of change? Given the possibility that observers may be able to implicitly detect a change without any awareness of it (see Implicit Perception, above), an important issue is to determine whether this form of perception does exist outside of visuomotor responses. If so, a key concern would be the mechanisms involved—not just in their details of their operation, but also how they relate to each other (i.e., whether there exist several separate systems), and how they relate to the mechanisms underlying the explicit detection of change.
3. Does explicit detection of change involve only a subset of attended information? Although attention may be needed to explicitly experience change, it is unclear whether such detection is based on all attended information or just a selected subset. Some studies suggest that observers may hold only one property at a time in coherent form (e.g., Ballard et al 1995; Hayhoe et al 1998), a result at odds with our impression that we see many properties simultaneously. It could be

that our impressions result from a virtual representation that efficiently switches between properties as well as between objects. But it could also be that our visual experience reflects the many dimensions being attended, with only a subset of these being compared in any given task.

It is important to resolve this issue, for otherwise characteristics of the comparison mechanism could be ascribed to visual attention, or vice versa. For example, the relatively poor detection of incidental change (e.g., Levin & Simons 1997) cannot be used as conclusive evidence that fewer properties are being attended—it could simply be that fewer properties are being compared. However, careful analysis can often provide information about the attentional structures involved, even if the details of the comparison process are unknown (e.g., Rensink 2000a, 2001).

4. How might detection of dynamic change differ from detection of completed change? At a conceptual level, there is a considerable difference between these two types of detection, suggesting that separate mechanisms may be involved (see Dynamic vs. Completed Change, above). But whether this division really exists in the human visual system is unknown. Interestingly, a phenomenological dissociation appears to exist in the detection of gap-contingent change, with detection of dynamic change (seeing the dynamic transformation) found for ISIs of about 300 ms or less, and detection of completed change (seeing that something has changed) for longer ISIs (e.g., Phillips 1974; Rensink et al 2000a; see also Bridgeman et al 1975). To establish this difference firmly would require behavioral correlates (cf. Merikle & Reingold 1992). A possible candidate in this regard might be the finding of two separate mechanisms for the detection of displacement: one for ISIs less than 200 ms, the other for ISIs greater than 500 ms (Palmer 1986).

If the two types of detection are separate, the continued application of attention would likely be needed only for the detection of dynamic change; the detection of completed change might then be carried out via a referral to remembered information, with attention perhaps needed only for the comparison operation. This latter possibility may explain the finding that a change can sometimes be explicitly detected when the eye returns to a previously-fixated item (Henderson & Hollingworth 1999a); if it can be ascertained that attention is not continually applied to such items, this would indicate that the relevant quantity is indeed remembered information. But note that such information would not be a complete description of the scene—rather, it would still be a relatively sparse representation, such as the schematic map proposed by Hochberg (1968) for the guidance of eye movements, or the setting system proposed by Rensink (2000c) for the guidance of attention.

5. How might detection of dynamic change differ from detection of motion? If there is a separate system for the detection of dynamic change, its maximum temporal range would be approximately 300 ms. This is within the largest temporal range for motion detectors (van der Grind et al 1986), raising the possibility that detection of dynamic change is simply a form of motion perception. However, 300 ms is also a duration typical of informational persistence, suggesting that detection of dynamic change may be based on this instead. An interesting possibility in this regard is that changes may be picked up at the level of informational persistence, with these changes becoming explicitly detected only when attention is applied to the relevant structures (see Visual Buffer, above).

One way to settle this issue might be to examine how change detection is affected by a sudden displacement of the stimulus: performance would decline for a motion-based mechanism whereas it might remain unaffected for a mechanism referenced to stimulus structure. Evidence for the indifference of dynamic change detection to displacement has been found (Rensink et al 2000a), providing preliminary support for the existence of separate systems.

6. How might detection of completed change differ from detection of difference? If a distinct system exists for the detection of completed change, it might not need to have items continually attended—instead, a currently-visible item would simply be compared to a structure in memory. As such, detection of completed change may differ from perception of difference only in the use of an additional process to establish that the structures being compared refer to same external entity (see Change vs. Difference, above). But the detection of completed change can occur after ISIs of only 300 ms, something difficult to reconcile with the long-term memories that can be involved in the perception of difference. And preliminary work (e.g., Shore & Klein, 2000) suggests behavioral differences between the two forms of detection. It may be that several systems are at play, each involving a different time scale, and perhaps different ways of making comparisons.

7. What are the limits to which change-detection mechanisms can access visual representations? Although change blindness can arise from a failure to maintain relevant information, it can also arise from a failure to access the relevant representations (see e.g., Simons, 2000). For example, the explicit detection of saccade-contingent change is greatly improved when a brief blank follows the landing of the eye (Deubel et al 1996), showing the existence of preserved information normally inaccessible to the mechanisms underlying the conscious detection of change. Again, such information is unlikely to be a complete and

detailed description of the visible scene, being instead a sparser description containing information about particular aspects of it.

More generally, it may be that the group of mechanisms underlying each kind of change detection (e.g., explicit or implicit, dynamic or completed) can access only a subset of the representations used in vision, and that each group accesses a different subset. If so, it becomes important to determine the various types of information that are preserved, the subsets that can be accessed by each group of change-detection mechanisms, and the reasons for these limits.

Resolving these issues will not be easy. But as the discussion above has indicated, ways exist to investigate the various problems involved. If history is any guide, these investigations will likely lead to effects again at odds with our intuitions. These will then lead to further revisions in our understanding of what change detection is and how it is carried out. And this in turn will improve our ability to explore other aspects of visual perception. Thus, it appears that the concept of change, a concept that caused so many difficulties for so many years, is well on its way to becoming the basis for many important new insights into the way that we experience the dynamic world around us.

Literature Cited

- Adelson EH, Bergen JR. 1991. The plenoptic function and the elements of early vision. In Computational Models of Visual Processing, ed. M Landy, JA Movshon, pp. 3-20. Cambridge, MA: MIT Press.
- Aginsky V, Tarr MJ. 2000. How are different properties of a scene encoded in visual memory? Visual Cognition, 7: 147-162.
- Agostinelli G, Sherman SJ, Fazio RH, Hearst ES. 1986. Detecting and identifying change: Additions versus deletions. Journal of Experimental Psychology: Human Perception and Performance, 12: 445-454.
- 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.
- Arijon D. 1976. Grammar of the Film Language. Los Angeles: Silman-James Press.
- Austen E, Enns JT. 2000. Change detection: Paying attention to detail. Psyche 6(11). <http://psyche.cs.monash.edu.au/v6/psyche-6-11-austen.html>
- Avons SE, Phillips WA. 1980. Visualization and memorization as a function of display time and poststimulus processing time. Journal of Experimental Psychology: Human Learning and Memory, 6: 407-420.
- Ballard DH, Hayhoe MM, Pelz JB. 1995. Memory representations in natural tasks. Journal of Cognitive Neuroscience, 7: 66-80.
- Batchelor GK. 1967. An Introduction to Fluid Dynamics (pp. 71-73). Cambridge: Cambridge University Press.
- Biederman I, Gerhardstein PC. 1993. Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. Journal of Experimental Psychology: Human Perception and Performance, 19: 1162-1182.
- Blackmore SJ, 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.
- Brawn P, Snowden RJ. 1999. Can one pay attention to a particular color? Perception & Psychophysics, 61: 860-873.

- Bridgeman B, Hendry D, Stark L. 1975. Failure to detect displacement of the visual world during saccadic eye movements. Vision Research, **15**: 719-722.
- Bridgeman B, Lewis S, Heit G, 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.
- Bridgeman B, van der Heijden AH, Velichkovsky BM. 1994. A theory of visual stability across saccadic eye movements. Behavioral and Brain Sciences, **17**: 247-292.
- Brunel N, Ninio J. 1997. Time to detect the difference between two images presented side by side. Cognitive Brain Research, **5**: 273-282.
- Carlson-Radvansky LA, Irwin DE. 1995. Memory for structural information across eye movements. Journal of Experimental Psychology: Learning, Memory, and Cognition, **21**: 1441-1458.
- Carlson-Radvansky LA, Irwin DE. 1999. Memory for relational information across eye movements. Perception & Psychophysics, **61**: 919-934.
- Castiello U, Jeannerod M. 1991. Measuring time to awareness. Neuroreport, **2**: 797-800.
- Castiello U, Paulignan Y, Jeannerod M. 1991. Temporal dissociation of motor responses and subjective awareness. Brain, **114**: 2639-2655.
- Cowan N. 1988. Evolving conceptions of memory storage, selective attention, and their mutual constraints with the human information-processing system. Psychological Review, **104**: 163-191.
- Deubel H, Schneider WX, Bridgeman B. 1996. Post-saccadic target blanking prevents saccadic suppression of image displacement. Vision Research, **36**: 985-996.
- Ditchburn R. 1955. Eye-movements in relation to retinal action. Optica Acta, **1**: 171-176.
- Dmytryk E. 1984. On Film Editing. Boston, MA: Focal Press.
- Dulany DE. 1997. Consciousness in the explicit (deliberative) and implicit (evocative). In Scientific Approaches to Consciousness, ed. JD Cohen & JW Schooler, pp. 179-211. Mahwah, NJ: Erlbaum.

- Feldman JA. 1985. Four frames suffice: A provisional model of vision and space. Behavioral and Brain Sciences, **8**: 265-289.
- 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.
- French RS. 1953. The discrimination of dot patterns as a function of number and average separation of dots. Journal of Experimental Psychology, **46**: 1-9.
- Gelfand IM, Fomin SV. 1963. Calculus of Variations. (R.A. Silverman, translator.) Englewood Cliffs, NJ: Prentice-Hall.
- Goodale MA, Pelisson D, Prablanc C. 1986. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. Nature, **320**: 748-750.
- Grimes J. 1996. On the failure to detect changes in scenes across saccades. In Perception (Vancouver Studies in Cognitive Science, vol. 5), ed. K Akins, pp. 89-109. New York : Oxford University Press.
- Gur RC, Hilgard ER. 1975. Visual imagery and the discrimination of differences between altered pictures simultaneously and successively presented. British Journal of Psychology, **66**: 341-345.
- Gysen V, De Graef P, Verfaillie K. 2000. Detection of intrasaccadic changes in stationary and moving objects. Investigative Ophthalmology & Visual Science, **41**: 45. (Abstr.)
- Haber RN. 1983. The impending demise of the icon: A critique of the concept of iconic storage in visual information processing. Behavioral and Brain Sciences, **6**: 1-54.
- Hayhoe MM, Bensinger DG, Ballard DH. 1998. Task constraints in visual working memory. Vision Research, **38**: 125-137.
- Henderson JM. 1992. Object identification in context: The visual processing of natural scene. Canadian Journal of Psychology, **46**: 319-341.
- Henderson JM. 1997. Transsaccadic memory and integration during real-world object perception. Psychological Science, **8**: 51-55.

- Henderson JM, Hollingworth A. 1999a. The role of fixation position in detecting scene changes across saccades. Psychological Science, **10**: 438-443
- Henderson JM, Hollingworth A. 1999b. High-level scene perception. Annual Review of Psychology, **50**: 243-271.
- Hildreth EC, Koch C. 1987. The analysis of visual motion: From computational theories to neuronal mechanisms. Annual Review of Neuroscience, **10**: 477-533.
- Hochberg J. 1968. In the mind's eye. In Contemporary Theory and Research in Visual Perception, ed. RN Haber, pp. 309-331. New York: Holt, Rinehart, & Winston.
- Hollingworth A, Henderson JM. 2000. Semantic informativeness mediates the detection of changes in natural scenes. Visual Cognition, **7**: 213-235.
- Irwin DE. 1991. Information integration across saccadic eye movements. Cognitive Psychology, **23**: 420-456.
- Irwin DE. 1996. Integrating information across saccadic eye movements. Current Directions in Psychological Science, **5**: 94-100.
- Jacobs WJ, Blackburn JR, Buttrick M, Harpur TJ, Kennedy D, Mana MJ, MacDonald MA, McPherson LM, Paul D, Pfaus JG. 1988. Observations. Psychobiology, **16**: 3-19.
- Jiang Y, Olson IR, Chun MM. 2000. Organization of visual short-term memory. Journal of Experimental Psychology: Learning, Memory, and Cognition, **26**: 683-702.
- Kahneman D, Treisman A, Gibbs B. 1992. The reviewing of object files: Object-specific integration of information. Cognitive Psychology, **24**: 175-219.
- Kanizsa G. 1979. Organization in Vision. New York: Praeger.
- Kanizsa G, Gerbino W. 1982. Amodal completion: Seeing or thinking. In Organization and Representation in Perception, ed. J Beck, pp. 167-189. Hillsdale, NJ: Erlbaum.
- Kirk GS, Raven JE. 1957. The Presocratic Philosophers: A Critical History with a Selection of Texts. Cambridge: Cambridge University Press.
- Klein R, Kingstone A, Pontefract A. 1992. Orienting of visual attention. In Eye Movements and Visual Cognition: Scene Perception and Reading, ed. K Rayner, pp. 46-65. New York: Springer.

- Levin DT, Momen N, Drivdahl SB, Simons DJ. 2000. Change blindness blindness: The metacognitive error of overestimating change-detection ability. Visual Cognition, **7**: 397-412.
- Levin DT, Simons DJ. 1997. Failure to detect changes to attended objects in motion pictures. Psychonomic Bulletin & Review, **4**: 501-506.
- Levin DT, Simons DJ. 2000. Perceiving stability in a changing world: Combining shots and integrating views in motion pictures and the real world. Media Psychology, **2**: 357-380.
- Luck SJ, Vogel EK. 1997. The capacity of visual working memory for features and conjunctions. Nature, **390**: 279-280.
- Mack A. 1970. An investigation of the relationship between eye and retinal image movement in the perception of movement. Perception & Psychophysics, **8**: 291-298.
- Mack A, Rock I. 1998. Inattentional Blindness. Cambridge MA: MIT Press.
- Marr D. 1982. Vision. San Francisco: Freeman.
- McConkie GW, Currie CB. 1996. Visual stability across saccades while viewing complex pictures. Journal of Experimental Psychology: Human Perception and Performance, **22**: 563-581.
- McConkie GW, Zola D. 1979. Is visual information integrated across successive fixations in reading? Perception & Psychophysics, **25**: 221-224.
- Merikle PM, Reingold EM. 1992. Measuring unconscious perceptual processes. In RF Bernstein and TS Pittman (Eds.), Perception Without Awareness (pp. 55-80). New York: Guilford Press.
- Miller GA. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, **63**: 81-97.
- Milner AD, Goodale MA. 1995. The Visual Brain in Action. Oxford: Oxford University Press.
- Mitroff SR, Simons DJ. 2000. Changes are not localized before they are explicitly detected. Poster presented at ARVO, Ft. Lauderdale.

- Mondy S, Coltheart V. 2000. Detection and identification of change in naturalistic scenes. Visual Cognition, 7: 281-296.
- Nakayama K. 1985. Biological motion processing: A review. Vision Research, 25: 625-660.
- O'Regan JK, Deubel H, Clark JJ, Rensink RA. 2000. Picture changes during blinks: Looking without seeing and seeing without looking. Visual Cognition, 7: 191-211.
- O'Regan JK, Rensink RA, Clark JJ. 1999. Change Blindness as a Result of 'Mudsplashes'. Nature, 398: 34.
- Palmer J. 1986. Mechanisms of displacement discrimination with and without perceived movement. Journal of Experimental Psychology: Human Perception & Performance, 12: 411-421.
- Palmer J. 1988. Very short-term visual memory for size and shape. Perception & Psychophysics, 43: 278-286.
- Palmer J. 1990. Attentional limits on the perception and memory of visual information. Journal of Experimental Psychology: Human Perception & Performance, 16: 332-350.
- Pashler HE. 1988. Familiarity and visual change detection. Perception & Psychophysics, 44: 369-378.
- Pashler HE. 1998. The Psychology of Attention. Cambridge, MA: MIT Press.
- Pashler HE, Badgio PC. 1987. Attentional issues in the identification of alphanumeric characters. In Attention and Performance XII, ed. M Coltheart, pp. 63-81. Hillsdale, NJ: Erlbaum.
- Phillips WA. 1974. On the distinction between sensory storage and short-term visual memory. Perception & Psychophysics, 16: 283-290.
- Phillips WA, Singer W. 1974. Function and interaction of On and Off transients in vision. I. Psychophysics. Exp. Brain Research, 19: 493-506.
- Pollack I. 1972. Detection of changes in spatial position: III. Dot number or dot density. Perception & Psychophysics, 12: 487-491.
- Pylyshyn ZW, Storm RW. 1988. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. Spatial Vision, 3 179-197.

- Rensink RA 1998. Mindsight: Visual sensing without seeing. Investigative Ophthalmology & Visual Science, **39**: 631. (Abstr.)
- Rensink RA. 1999a. The Magical Number One, Plus or Minus Zero. Investigative Ophthalmology & Visual Science, **40**: 52. (Abstr.)
- Rensink RA. 1999b. Attentional Processing of Geometric Figures. Perception, **28** (suppl.): 55-56. (Abstr.)
- Rensink RA. 2000a. Seeing, sensing, and scrutinizing. Vision Research, **40**: 1469-1487.
- Rensink RA. 2000b. Visual search for change: A probe into the nature of attentional processing. Visual Cognition, **7**: 345-376
- Rensink RA. 2000c. The dynamic representation of scenes. Visual Cognition, **7**: 17-42.
- Rensink RA. 2001. Change blindness: Implications for the nature of visual attention. In M Jenkin and LR Harris (Eds.), Vision and Attention. New York: Springer.
- Rensink RA, Deubel H, Schneider WX. 2000a. The incredible shrinking span: Estimates of memory capacity depend on interstimulus interval. Investigative Ophthalmology & Visual Science, **41**:425. (Abstr.)
- 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.
- Rensink RA, O'Regan JK, Clark JJ. (2000b). On the failure to detect changes in scenes across brief interruptions. Visual Cognition, **7**: 127-145.
- Rich A, Gillam B. 2000. Failure to detect changes in color for lines rotating in depth: the effects of grouping and type of color change. Vision Research, **40**: 1469-1487.
- Ro T, Russell C, Lavie N. 2001. Changing faces: A detection advantage in the flicker paradigm. Psychological Science, **12**: 94-99.
- Saiki J. 1999. Relation blindness: Difficulty in detecting a violation of spatial relations in regular rotations of a triangular pattern. Investigative Ophthalmology & Visual Science, **40**: 49. (Abstr.)

- Scholl BJ. 2000. Attenuated change blindness for exogenously attended items in a flicker paradigm. Visual Cognition, **7**: 377-396.
- Scholl BJ, Pylyshyn ZW, Franconeri SL. 1999. When are featural and spatiotemporal properties encoded as a result of attentional allocation? Investigative Ophthalmology & Visual Science, **40**: 797. (Abstr.)
- Scott-Brown KC, Baker MR, Orbach HS. 2000. Comparison blindness. Visual Cognition, **7**: 253-267.
- Scott-Brown KC, Orbach HS. 1998. Contrast discrimination, non-uniform patterns, and change blindness. Proc. Royal Society London, **B 265**: 2159-2166.
- Seiffert AE, Cavanagh P. 1998. Position displacement, not velocity, is the cue to motion detection of second-order patterns. Vision Research, **38**: 3569-3582.
- Shore DI, Klein RM. 2000. The effects of scene inversion on change-blindness. Journal of General Psychology, **127**: 27-44.
- Simons DJ. 1996. In sight, out of mind: When object representations fail. Psychological Science, **7**: 301-305.
- Simons DJ. 2000. Current approaches to change blindness. Visual Cognition, **7**: 1-15.
- Simons DJ, Franconeri SL, Reimer RL. 2000. Change blindness in the absence of visual disruption. Perception, **29**: 1143-1154.
- Simons DJ, Levin DT. 1997. Change blindness. Trends in Cognitive Sciences, **1**: 261-267.
- Simons DJ, Levin DT. 1998. Failure to detect changes to people during a real-world interaction. Psychonomic Bulletin & Review, **5**: 644-649.
- Simons DJ, Mitroff SR. 2001. The role of expectations in change detection and attentional capture. In M Jenkin and LR Harris (Eds.), Vision and Attention. New York: Springer.
- Smilek D, Eastwood JD, Merikle PM. 2000. Does unattended information facilitate change detection? Journal of Experimental Psychology: Human Perception and Performance, **26**: 480-487.
- Smith BC. 1998. On the Origin of Objects. Cambridge, MA: MIT Press.

- Sperling G. 1960. The information available in brief visual presentations. Psychological Monographs, **74**: 1-29.
- Sperling G, Hoff CE. 2000. Third-order versus first-order and second-order motion in ambiguous stimuli: Competition reveals temporal tuning functions, monocular/binocular, and the role of attention. Perception, **29** (suppl.): 83. (Abstr.)
- Sperling G, Speelman RG. 1965. Visual spatial localization during object motion, apparent object motion, and image motion produced by eye movements. Journal of the Optical Society of America, **55**: 1576-1577.
- Sperling G. 1990. Comparison of perception in the moving and stationary eye. In Eye Movements and Their Role in Visual and Cognitive Processes, ed. E Kowler, pp. 307-351. New York: Elsevier.
- Thornton IM, Fernandez-Duque D. 2000. An implicit measure of undetected change. Spatial Vision, **14**: 21-44.
- Treisman A, Gormican S. 1988. Feature analysis in early vision: Evidence from search asymmetries. Psychological Review, **95**: 15-48.
- van der Grind WA, Koenderink JJ, van Doorn AJ. 1986. The distribution of human motion detector properties in the monocular visual field. Vision Research, **26**: 797-810.
- Vaughan B, Yantis S. 1999. Information preservation during occlusion. Investigative Ophthalmology & Visual Science, **40**: 415. (Abstr.)
- Verfaillie K, de Troy A, van Rensbergen J. 1994. Transsaccadic integration of biological motion. Journal of Experimental Psychology: Learning, Memory, and Cognition, **20**: 649-670.
- Wallach H, Lewis C. 1966. The effect of abnormal displacement on the retinal image during eye movements. Perception & Psychophysics, **1**: 25-29.
- Wang RF, Simons DJ. 1999. Active and passive scene recognition across views. Cognition, **70**: 191-210.
- 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. 2001. Binding in short-term visual memory. Journal of Experimental Psychology: General. In press.
- Wilken P, Mattingley JB, Korb KB, Webster WR, Conway D. 1999. Capacity limits for detection versus reportability of change in visual scenes. Paper presented at the 26th Annual Australian Experimental Psychology Conference. 9-11 April 1999, Sydney, Australia.
- Williams P, Simons DJ. 2000. Detecting changes in novel, complex, three-dimensional objects. Visual Cognition, 7: 297-322.
- Wolfe JM. 1999. Inattentional amnesia. In Fleeting Memories, ed. V Coltheart, pp. 71-94. Cambridge, MA: MIT Press.
- Woodhouse JM, Barlow HB. 1982. Spatial and temporal resolution and analysis. In The Senses, ed. HB Barlow, pp. 152-162. Cambridge: Cambridge University Press.
- Wright M, Green A, Baker S. 2000. Limitations for change detection in multiple Gabor targets. Visual Cognition, 7: 237-252.
- Yantis S. 1992. Multielement visual tracking: Attention and perceptual organization. Cognitive Psychology, 24: 295-340.
- Zelinsky G. 1997. Eye movements during change detection search task. Investigative Ophthalmology & Visual Science, 38: 373. (Abstr.)

Footnotes

1. The work on saccade-contingent change was intended to address the question of how visual experience could remain stable in the face of constant eye movements. The results of these studies—together with those of the following decades—did help to develop a framework for this problem (see e.g., Bridgeman et al 1994). But there does not appear to have been a serious attempt to relate the work on saccade-contingent change to the work on gap-contingent change, or to the issue of visual memory in general.

2. Perhaps the most famous of these is the paradox of Achilles and the tortoise, which involves change in position (i.e., movement). Here, a tortoise is given a head start in a race with Achilles. It is then argued that Achilles can never overtake the tortoise, for by the time he arrives at the previous location of the tortoise, the tortoise will have moved yet further ahead. Based on apparent paradoxes such as this, some early thinkers (e.g., Parmenides of Elea) argued that change was impossible, and that the perception of change was therefore illusory (see e.g., Kirk & Raven 1957).

3. This distinction can always be maintained, even when the variations are smooth and pertain to position. For the simple case of smoothly-moving particles, for example, there exist two distinct ways of describing the situation (Batchelor 1967). The first is an Eulerian specification, where derivatives are defined with respect to a fixed point in space. The second is a Lagrangian specification, where derivatives are defined with respect to the path of a particular (arbitrarily small) element. The two forms are generally different, and describe different quantities: motion and (position) change, respectively.

4. In general, simple continuity is not enough: additional conditions are generally required to ensure that the derivatives of that particular order and type can exist (see, e.g., Gelfand & Fomin 1963). However, the initial stages of visual processing usually involve only simple first- and second-order derivatives, which are calculated after smoothing by low-pass filters. As such, these quantities exist for virtually all conditions encountered when viewing the world (see, e.g., Hildreth & Koch 1987; Marr 1982).

5. The term motion is reserved here for variation measured with respect to a fixed location, and movement for a change in position of any part of a structure. As such, a change in the use of terms is to be advocated. For example, "first-order motion" and

"second-order motion" pertain to variation referenced to retinotopic space, and so do describe types of motion; "third-order motion", however, refers to the tracking of a moving structure, and is therefore a misnomer, describing instead a type of movement. This distinction is not just nominal—the mechanisms underlying the perception of first- and second-order motion are much alike, whereas they differ fundamentally from those underlying the perception of movement (e.g., Seiffert & Cavanagh, 1998; Sperling & Hoff, 2000). More generally, then, the displacement of an object gives rise to the perception of motion (temporal variation of intensity, color, etc.) at various points in the field of view, while the object itself is perceived as moving.

6. In some sense, these two types of perception are analogous to modal completion (i.e., the perception of visible properties not actually in the image to account for gaps between aligned fragments; see, e.g., Kanizsa, 1979) and amodal completion (i.e., the perception that aligned fragments in contact with a visible occluder are parts of the same structure; see, e.g., Kanizsa & Gerbino, 1982). Detection of dynamic change involves seeing a dynamic visual event that extends over a temporal gap; this is akin to seeing a modally completed surface that extends over a spatial gap. Meanwhile, detecting completed change involves a more abstract linking of items believed to refer to the same spatiotemporal structure; this is like the abstract linking of amodally-completed fragments.

7. This simplicity is only in terms of the pattern of light on the retina or the image on a visual display. In terms of entities in the external world, changes in existence are rather special. First, they pertain to the "well-defined" nature of the entity itself rather than to any particular properties it may have. Second, they are relatively uncommon, being limited to creation and destruction. This latter factor may account for the tendency to interpret a change of existence in a stimulus element as some other kind of change, e.g., a displacement of the corresponding entity to a distant location.

8. The term complex is used here to denote the representational structure formed by attention for an item in the stimulus array, without regard to any particular theory. Other terms (such as object file or coherence field) are associated with particular ideas about the formation and content of complexes, associations which are avoided by use of a more neutral word. The more theory-laden terms will be used only in the context of particular theories of attention.

9. Care must be taken in distinguishing between object defined as a structure in the external world (i.e., a concrete spatiotemporal entity) and as a structure internal to the

observer (i.e., an attentional complex). If "object" is taken to be an external structure, then attention to it will encode only some of its properties. As such, attention is necessary but not sufficient to perceive change in an object. If "object" is defined as the contents of the complex, the issue is whether there exists additional selectivity in the comparison process; if not, attention would be both necessary and sufficient to perceive change in this internal structure.

10. Note that this applies only to the visual experience of coherent structure—the structure manifested by, e.g., the ability to see change. Nothing is implied about how volatile structure is experienced. It may be, for example, that volatile structures support the visual experience of a dense "sea" of simple features, although changes per se will still not be seen. If so, the visual impression an observer has of their surroundings will be based on a just few coherent structures embedded in this "sea".

Figure Captions

Figure 1. Two variants of gap-contingent change. (a) One-shot paradigm. Here, the observer views a single alternation of displays, and determines if a change occurred. Performance is measured via accuracy of response. (b) Flicker paradigm. Here, the observer views a continual cycling of displays, and determines if a change is occurring. Performance is measured via response time. Both approaches can also be applied to other kinds of contingent change, such as those made during eye movements or blinks.

Figure 2. Coherence theory (Rensink 2000c). In the absence of focused attention, low-level structures (proto-objects) are volatile, and so are simply replaced by the representations of new stimuli at their location. Attention acts by establishing feedback links between a selected set of proto-objects that provides visual properties and a mid-level nexus that stabilizes them; the interacting set of proto-objects, links, and nexus is a coherence field. This interaction allows attended proto-object properties to be held in coherent form, both in time and in space. When the feedback loop is broken, coherence dissolves, with the previously-attended proto-objects reverting to a volatile state.

Figure 3. Triadic architecture (Rensink 2000c). Here, visual processing is split into three largely independent systems: (i) an early system concerned with the rapid formation of sophisticated—but volatile—proto-objects, (ii) an attentional system concerned with the formation of coherent objects, and (iii) a nonattentional setting system that guides attention. According to coherence theory, attention is needed for the visual experience of change, and so unattended structures do not support such experience. The contribution of unattended structures to the visual experience of static stimuli is unknown.




Figure 1

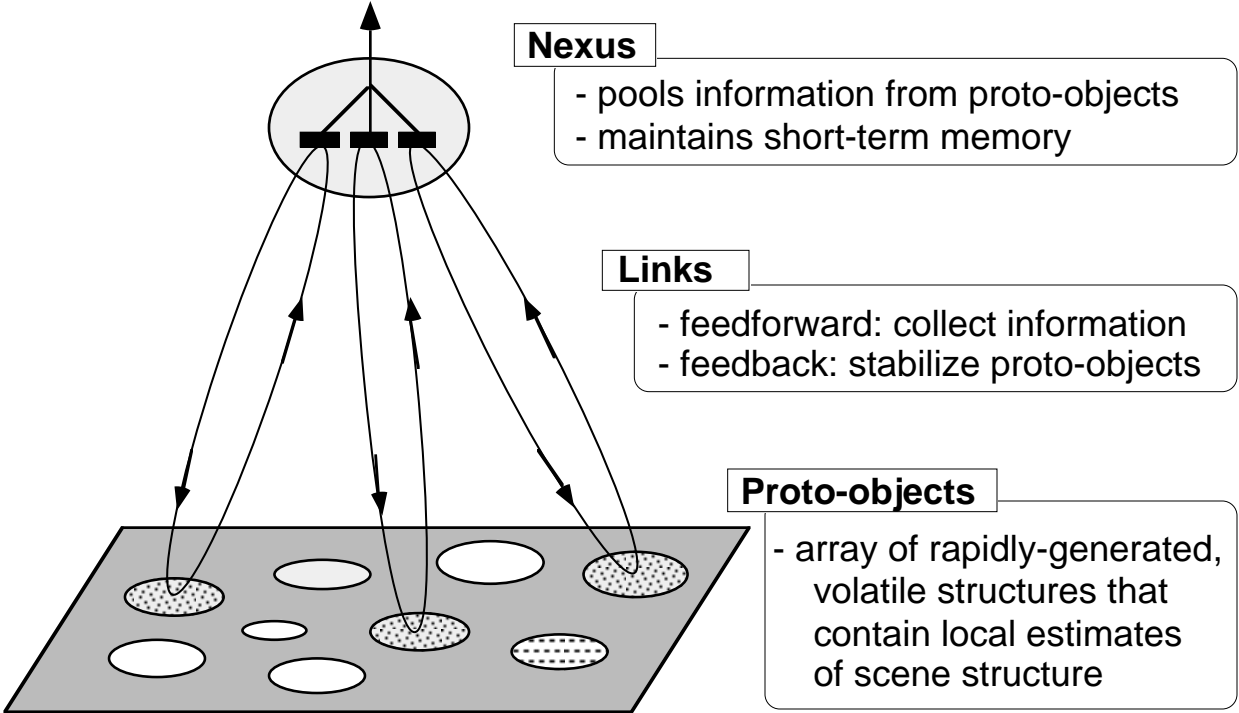


Figure 2

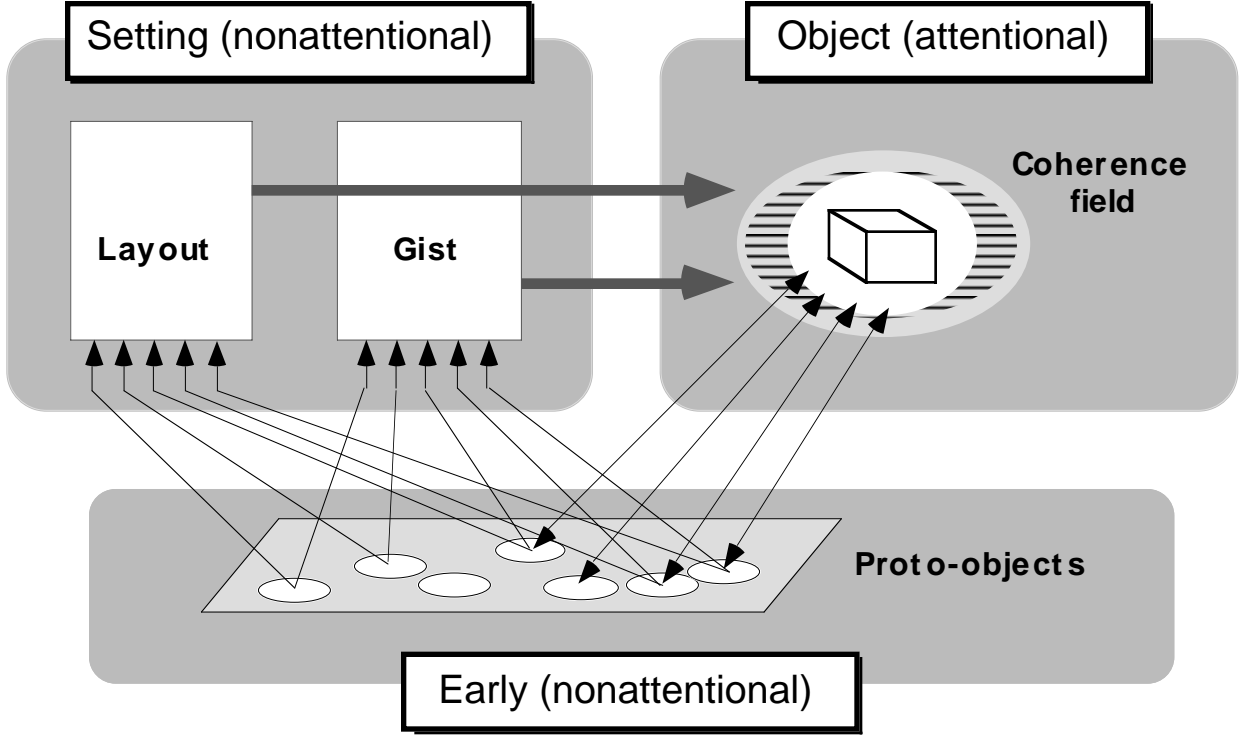


Figure 3