VISUAL COGNITION, 2000, *7* (1/2/3), 345–376

Visual Search for Change: A Probe into the Nature of Attentional Processing

Ronald A. Rensink

Cambridge Basic Research, Nissan Research and Development, Inc., Cambridge, MA, USA

A set of visual search experiments tested the proposal that focused attention is needed to detect change. Displays were arrays of rectangles, with the target being the item that continually changed its orientation or contrast polarity. Five aspects of performance were examined: linearity of response, processing time, capacity, selectivity, and memory trace. Detection of change was found to be a self-terminating process requiring a time that increased linearly with the number of items in the display. Capacity for orientation was found to be about five items, a value comparable to estimates of attentional capacity. Observers were able to filter out both static and dynamic variations in irrelevant properties. Analysis also indicated a memory for previously attended locations.

These results support the hypothesis that the process needed to detect change is much the same as the attentional process needed to detect complex static patterns. Interestingly, the features of orientation and polarity were found to be handled in somewhat different ways. Taken together, these results not only provide evidence that focused attention is needed to see change, but also show that change detection itself can provide new insights into the nature of attentional processing.

Change blindness is a rather striking phenomenon: A change made to an image during a saccade, flicker, or other such interruption will often be difficult to detect, even when it is large and easily seen once noticed (Rensink, O'Regan, & Clark, 1997; Simons & Levin, 1997). Much has been made of the "negative" aspect of this phenomenon—the fact that detecting change can be

---

Please address all correspondence to R.A. Rensink, Cambridge Basic Research, Nissan Research and Development, Inc., 4 Cambridge Center, Cambridge, MA 02142, USA.

I would like to thank Ian Thornton, Carol Yin, Greg Zelinsky, and an anonymous reviewer for comments on an earlier draft of this paper. I would also like to thank Monica Strauss, Nayantara Santhi, Janis Breeze, and Carolyn Harp for their help in running the experiments. Parts of this paper were presented at the 19th European Conference on Visual Perception, Strassburg, France, September 9–13, 1996.

difficult, even after several seconds of viewing. For example, it has been used as evidence for the idea that observers never form a detailed, spatiotopic "picture" of their surroundings (Rensink et al., 1997; Simons, 1996). It has also been used to support the idea that withdrawing attention from a representational structure causes it to lose spatial and temporal coherence (Rensink, 1997, this issue).

But such blindness has its limits. Although some types of change require large amounts of time before being seen, others are detected more quickly (Rensink et al., 1997; Simons, 1996). And even those changes that are difficult to notice are eventually seen, and seen clearly. As such, this "positive" aspect of change blindness—the eventual perception of change—yields an interesting set of effects in its own right, one that also has the potential to provide new insights into the way that we see.

As an example of how this potential can be developed, this paper examines visual search for change in arrays of simple stimuli. It has been argued that focused attention is needed to see change (Rensink, 1997, this issue). It has also been argued that focused attention is needed to see static patterns not immediately evident to an observer (e.g. Treisman & Gormican, 1988; Wolfe, 1994). If both these views are correct, it may be possible to combine them, and so extend the framework developed for the attentional search of static spatial patterns to the case of dynamic spatiotemporal patterns. And visual search for these changing patterns may in turn shed new light on the nature of the attentional processes involved.

ATTENTION AND THE PERCEPTION OF CHANGE

In what follows, the connection between focused attention and visual perception is taken to be that given by *coherence theory* (see Rensink, this issue). This theory assumes that object-based attention is intimately involved with the formation of representational structures with spatiotemporal coherence. More precisely (see Figure 1):

1. Prior to focused attention, structures are formed rapidly and in parallel across the visual field. These preattentive structures (or *proto-objects*) can be quite complex, but have limited spatial coherence (see e.g. Rensink & Enns, 1995, 1998). Their temporal coherence is similarly limited—they are volatile, and so need to be constantly regenerated. As such, they are simply *replaced* when a new stimulus appears at their retinal location.
2. Focused attention acts like a hand to "grab" proto-objects from this flux and stabilize them. While held, these structures form a *coherence field* corresponding to an individuated object. This coherence allows the object to retain

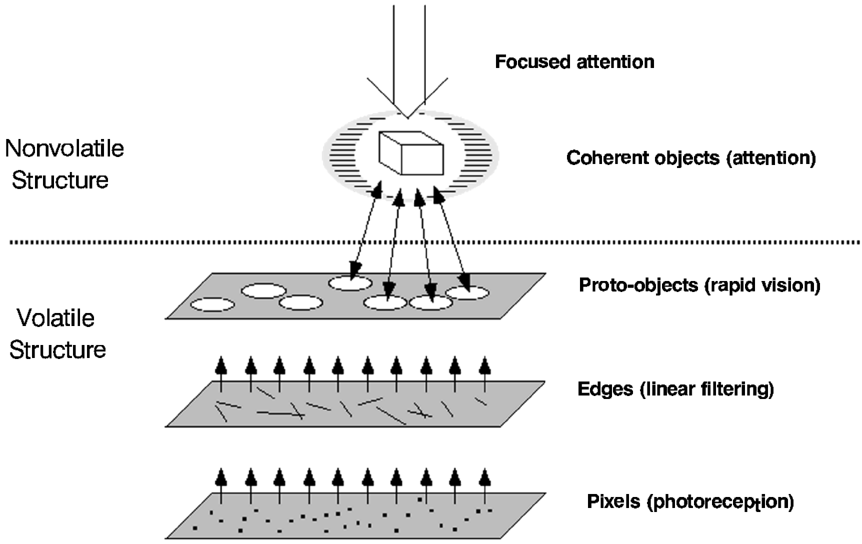

FIG. 1. Coherence theory. Early level processes produce volatile structures rapidly and in parallel across the visual field. Focused attention acts as a hand to "grab" these structures and stabilize them. As long as these structures are being held, they form an individuated object with both temporal and spatial coherence.

its continuity[1] across time, so that any new stimulus at that location is treated as the transformation of an existing structure rather than the appearance of a new one.

3. After focused attention is released, the coherence field is no longer maintained, and the object dissolves back into the original set of proto-objects. As such, there is little (if any) visual short-term memory of structure apart from what is currently being attended.

This concept of an object as a spatiotemporal coherence field is akin to the notion of an object file (Kahneman, Treisman, & Gibbs, 1992) but much more dynamic, with the field existing only as long as attention is being directed to the object. (For a more extensive discussion of these matters, see Rensink, this issue.) Among other things, coherence theory implies that there is little visual short-term memory (vSTM) apart from what is being attended; indeed, it

---

[1] The property of spatiotemporal continuity is sometimes described in terms of maintaining object identity. But "identity" often generates confusion with the rather different concept of semantic identity (cf. "identify") and so this term will not be used here. (For further discussion of this point, see Rensink, this issue.)

suggests that much—if not all—of vSTM may be the same as attentional hold.[2] This position is supported not only by results on change blindness, but also by results indicating a lack of attentional aftereffect even on static displays (Wolfe, 1996). In any event, the systematic investigation of change detection carried out here will provide a more thorough test of this proposal.

THE APPROACH

The approach taken here will be based on the detection of change under "flicker" conditions (Rensink et al., 1997). In these conditions, an original and a modified image continually alternate until the observer responds (Figure 2). A blank field briefly appears between each image; the transients generated by this field swamp the local motion signals that would normally draw attention to the location of the change.

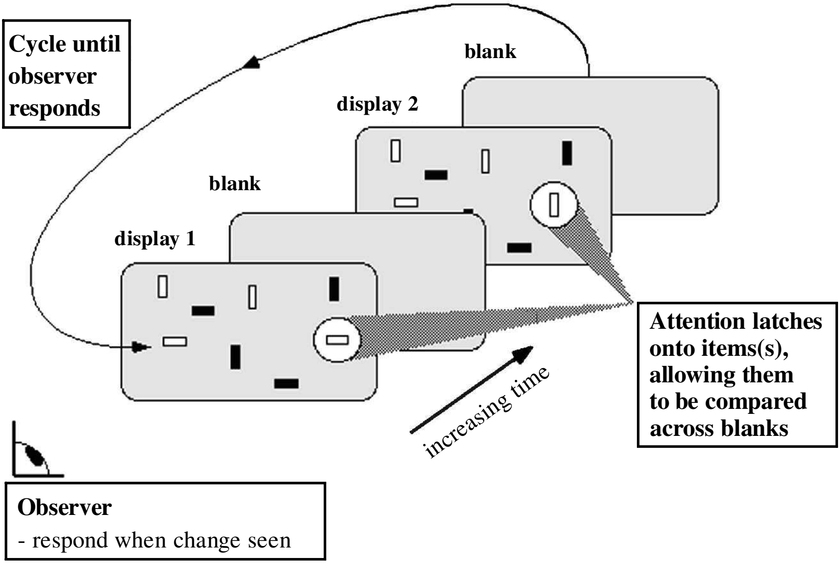

FIG. 2. Schematic of flicker paradigm. An original and modified image continually alternate, with a brief blank field presented after each image. The observer views the display until the change is seen. Since the target (a changing item) cannot be determined from any single display, attention must be applied to form stable structures that can be compared across the temporal gaps.

---

[2] There is of course a short-term memory for items that have been previously attended (see e.g. Cowan, 1988). However, in the view taken here this is not a *visual* memory supporting the formation object *tokens*. Rather, it is a more *abstract* memory containing information on object *types*. (For further discussion, see Rensink, this issue.)

All experiments will involve the *detection* of change, where observers report whether or not a change exists in the display. To avoid any influence of meaning (or *gist*) on performance, displays will be simple arrays of rectangles. Target items (when present) will change their properties, while distractor items retain theirs. This approach is somewhat akin to that of Phillips (1974) and Pashler (1988), except that (a) simpler displays are used, and (b) performance is not measured in terms of accuracy on briefly-presented displays, but rather in terms of the reaction time (RT) needed for displays of longer duration. As shown later, this approach can facilitate the disentanglement of several potentially confounding factors, such as the speed and capacity of the underlying mechanisms.

The issue to be investigated here is whether the mechanism used for change detection is the same as the attentional mechanism used in search for complex static patterns (e.g. Treisman & Gormican, 1988). This will be done by examining various aspects of visual search for change, and comparing these to what is believed true of focused attention. In particular, five aspects will be examined: Linearity of response, processing time, capacity limit, between-dimension selectivity, and memory for spatial location. Investigating the first of these involves determining if visual search for change is formally similar to search for complex static patterns. Investigating the next three involves showing that their characteristics are compatible with what is known about focused attention. Occam's razor can then be invoked to conclude that the processes involved are not just formally similar, but largely identical. Having established this identity, it will then be shown that search for change can also be used to shed light on other aspects of attentional processing, such as whether it has a memory of where it has been.

GENERAL METHOD

Each of the experimental conditions used well-known visual search methodology in which observers search as rapidly as possible for a target among a set of distractor items (e.g. Rensink & Enns, 1995; Treisman & Gormican, 1988). Each display was composed of an array of rectangles, with the two values of the changing property distributed roughly equally among the rectangles (Figure 2). The target did not have a unique static feature in any display—it was simply the item that continually changed its properties. Each distractor, in contrast, remained constant. In the interests of simplicity, only the properties of orientation and contrast polarity are examined here.

Displays were composed of 2, 6, or 10 items. Items were outlined rectangles of dimension 1.2° × 0.4°. These were positioned at random on an imaginary 5 × 4 grid of possible locations, with density controlled, that is, the average inter-item distance was approximately the same for all displays. The display area subtended approximately 15° × 12° of visual angle. The position of each item in the grid was jittered by ± 0.5°. The locations of all items were identical in the pair of displays used in each trial.

The temporal pattern (or *cadence*) used in all experiments is shown in Figure 2. After showing the first set of rectangles for a fixed display time (or *on-time*), the display was blanked for a brief interval (ISI, or *off-time*), during which the entire field was set to the colour of the display background. The next set of rectangles then appeared for a similar on-time, followed by a blank field for a similar off-time. The display then cycled back to the first set of rectangles, with the entire display sequence repeating until the observer responded. The colour of the blanks and the display backgrounds were always the same (medium grey), so that the items appeared to continually flicker on a motionless background.

Off-times were always 120 msec. With this value, differences in the motion signals created by the change are not sufficient to draw attention to the target (Pashler, 1988; Phillips, 1974; see also Appendix A for empirical verification). Since the target cannot be determined from any single display, attention must form the items into objects with sufficient spatiotemporal coherence to support the perception of change.

A Macintosh computer was used to generate the displays, control the experiments, and collect the data. Each condition tested 12 adult observers with normal or corrected-to-normal visual acuity. In all experiments, half the observers were naive to RT testing and visual search methodology, whereas the other half had extensive experience with search tasks. Observers completed three sets of 60 trials in each condition. Each condition was blocked, and conditions were counterbalanced in regards to presentation order.

The target was present on half the trials (chosen randomly) and absent in the other half. Observers were asked to determine the presence or absence of the target as quickly as possible, while maintaining an accuracy of at least 90%. Target presence or absence was reported by pressing one of two response keys. Visual feedback was given after each response.

In all experiments, the independent variable was the *set size*, defined as the number of items in the display. The primary dependent variable was the *search slope*, defined here as the slope of RT over display size; the inverse of this measure is the *search rate*. Another measure occasionally used was the *baseline*, defined as the (extrapolated) RT for a display of size of 1. Search slopes and baselines were obtained for each observer by an unweighted least-squares fit of the average RTs obtained for each set size. Although accuracy fell below 90% for some particularly difficult conditions, this did not appear to have much effect on search slopes (Appendix B), largely eliminating the possibility of speed–accuracy trade-offs influencing the results.

Data analyses usually focused on mean search slopes for the observers in each experiment. Unless otherwise indicated, differences were analysed using within-observer *t*-tests. For some conditions, between-observer analyses were also carried out. None of these latter analyses took advantage of the fact that some observers were common to both conditions; instead, the more conservative assumption was made that observers were sampled independently. All statistical tests were two-sided, and all reported differences were significant at the $p < .05$ level or better. Although target-absent measures will occasionally be used, consideration will usually be limited to target-present measures, since these are less susceptible to higher-level strategic influences (Chun & Wolfe, 1996).

ASPECT 1: LINEARITY OF REACTION TIMES

Issues

According to coherence theory, focused attention is needed to see change. To ascertain whether this is indeed the case, the first step is to determine if search for change is formally similar to search for static patterns. If so, the framework developed for spatial patterns can be legitimately applied to spatiotemporal patterns, allowing both kinds of search to be measured and analysed in much the same way.

At first glance, establishing formal similarity would appear to be a straightforward matter. Given that focused attention is needed to see change, detecting a target under flicker conditions will require an attentional scan of the display, with attention now being used to detect a spatiotemporal (and not just spatial) pattern. However, visual search for change is granular—responses are usually made after a display alternation. This is not fatal from the point of view of analysis, since performance can always be measured in terms of the number of alternations needed. But the framework developed for static search is based on RTs and search slopes rather than alternations, and it would be useful if these measures could be maintained.[3] Earlier change-detection studies (e.g. Rensink, 1996; Zelinsky, 1997) have shown that RTs and search slopes are useful measures of performance. But a careful investigation is still needed to establish their legitimacy.

The first issue in this regard is the extent to which the averaging of responses for a given set size provides a reliable estimate of RT. If the variance of responses is sufficiently high, averaging will allow the granular nature of the change-detection task to be ignored. Just as the average height of a population can be measured to millimetre precision using a ruler with only centimetre increments, so too can highly precise RTs be determined by averaging over

---

[3] Performance for search is sometimes measured in terms of detection accuracy on briefly presented displays (e.g. Sagi & Julesz, 1985). Accuracy measures are also natural for "one-shot" change-detection tasks (e.g. Pashler, 1988). But RT slopes are a more common performance measure for static patterns. The issue is therefore whether RT slopes can be reliably used for the flicker paradigm described here.

coarser measures (Ulrich & Giray, 1989). The question is whether the variance in responses here is large enough for this to occur.

The second issue is whether average RTs increase linearly with set size, so that search slopes can be used to describe performance. The detection of change has several components: loading information into vSTM, holding it across a blank interval, comparing the stored to the visible information in the new display, and—if search needs to be continued—unloading vSTM and shifting attention to a new location. Although some of these operations are used in search for static patterns, others are not. If any of these additional actions requires time that does not increase linearly with set size, RT linearity will fail, and slopes will not be legitimate measures of performance.

Thus, to establish a formal similarity between search for change and search for static patterns, it is critical that RT linearity be verified. Note that such linearity does not necessarily imply that search is carried out sequentially on an item-by-item basis. As for the case of static patterns, attention could be directed to several (or even all) items in parallel, with detection simply taking longer when more items are included (see e.g. Townsend, 1990).

Finally, if search does turn out to be linear, it becomes important to determine if it terminates after the target has been found. The critical measure here is the *slope ratio*, defined as the ratio of target-absent to target-present slopes. If the process is self-terminating—as is true for complex static patterns (Treisman & Gormican, 1988)—this ratio will be about 2. Note however, that its exact value may vary somewhat, since strategic factors can influence how long observers continue to search for absent targets (Chun & Wolfe, 1996).

Approach

Experiment 1 examined search for change at two different on-times: A short on-time of 80 msec and a long on-time of 800 msec. Observers were run on both the 80 msec and 800 msec conditions to allow within-observer comparisons. Experiment 1A examined the case of orientation change. Here, all items were black, with orientations either horizontal or vertical. Experiment 1B examined the corresponding case of polarity change. Items here were vertically oriented, and could be either black or white.

Results

Reaction times for the various conditions are shown in Figure 3. As is evident from the figure, approximately linear behaviour was found in all conditions, with target-absent slopes greater than target-present slopes.

Linearity was tested by comparing the RT increments from 2 to 6 items against those from 6 to 10 items: if RT is a linear function of set size, these increments should be identical. For orientation change at 80 msec, no

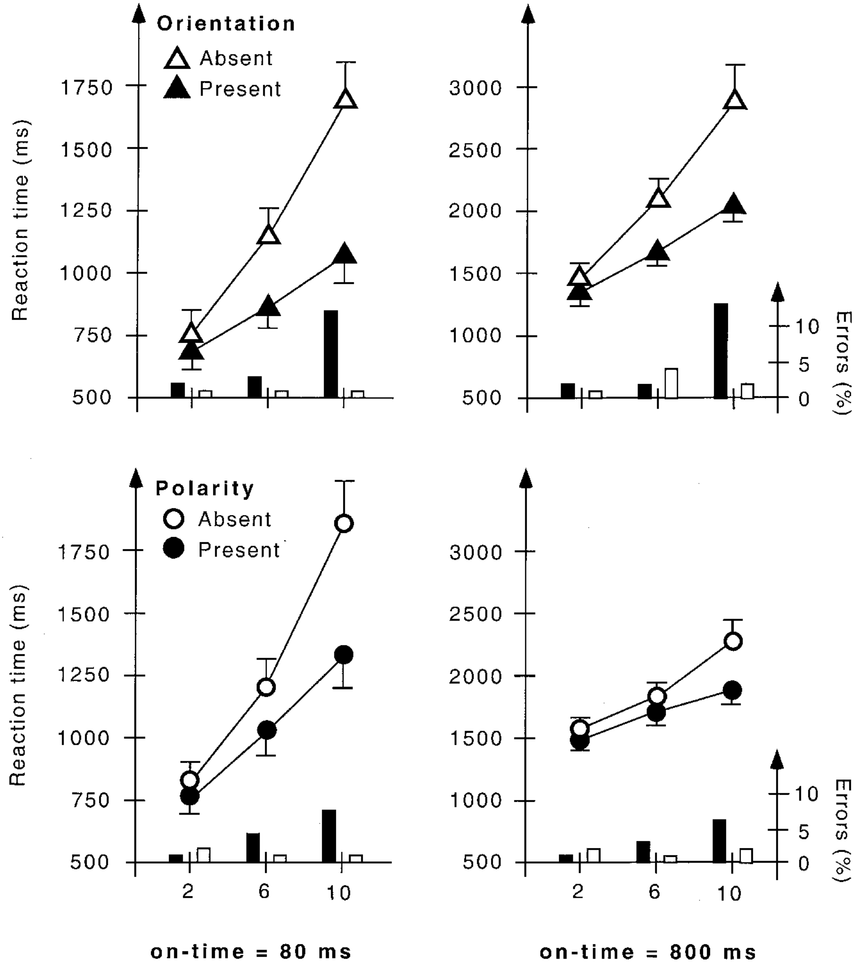

FIG. 3. Results of Experiment 1 (test for linearity). Upper graphs show results for changes in orientation; lower graphs show results for changes in polarity. [Note that scales on left-hand side differ by a factor of 2 from those on right-hand side.] In all cases, RTs have a strong linear dependence on set size, justifying the use of slope measures. For both types of change, target-absent slopes are about twice the value of target-present slopes, showing that the process is self-terminating.

significant difference was found for the target-present responses ($p > .15$). A slight positive acceleration appeared in target-absent responses, but the difference in increments just failed to reach significance ($p < .07$). For on-times of 800 msec, target-present increments were not significantly different ($p > .4$), and neither were target-absent increments ($p > .1$). In all cases, the fit of the RT data to a straight line was excellent, with r^2 measures exceeding .99.

A similar pattern was found for polarity change. For on-times of 80 msec, no significant difference was found for the target-present responses (*p* > .2), although a positive acceleration did appear in the target-absent case (*p* < .02). Likewise, for on-times of 800 msec, no reliable difference between increments was found for target-present responses (*p* > .5), and a positive acceleration was found for target-absent responses (*p* < .01). However, even where acceleration occurred, a strong linear trend still existed—r^2 values for all conditions were greater than 0.97.

Given that average RT is effectively a linear function of set size, it follows that slopes are a valid measure of performance, and so can be used to determine if search is self-terminating. For orientation change at 80 msec, target-present slopes were 48 msec/item, and target-absent slopes 115 msec/item, yielding a slope ratio of 2.4. Although target-absent slopes were reliably greater than twice the target-present slopes (*p* < .03), such a ratio is still consistent with a self-terminating search (Chun & Wolfe, 1996). For on-times of 800 msec, slopes were 85 msec/item and 174 msec/item, yielding a slope ratio of 2.04. Target-absent slopes differed significantly from target-present slopes (*p* < .0003), but not from twice the target-present slopes (*p* > .85).

Self-terminating search was also found for polarity change. For on-times of 80 msec, slopes were 69 msec/item and 127 msec/item, yielding a ratio of 1.9. Consistent with this, target-absent slopes differed from target-present slopes (*p* < .001), but not twice the target-present slopes (*p* > .5). For on-times of 800 msec, slopes were 51 msec/item and 92 msec/item, resulting in a ratio of 1.8. Again, target-absent slopes differed from target-absent slopes (*p* < .002), but not from twice the target-present slopes (*p* > .4).

Discussion

These results clearly show that search for change behaves much like search for complex static patterns. Target-present responses in all conditions were reliably linear, with the RT increment between set sizes of 2 and 6 being much the same as the RT increment between set sizes of 6 and 10. Target-absent responses did not have quite this degree of linearity. However, reaction times for target-absent trials reflect strategic considerations as well as perceptual processing (Chun & Wolfe, 1996), and so slight deviations from linearity are not unreasonable. In any event, a conservative approach will be taken here in subsequent experiments, with the analyses of slopes based only on target-present data.

Finally, the data not only show that the search process is linear, but also that it is self-terminating, that is, it ends as soon as the target is found. Consequently, visual search for change behaves much like visual search for complex static patterns, and can therefore be analysed using the same formal framework.

ASPECT 2: PROCESSING TIME

Issues

Although Experiment 1 showed that search slopes are meaningful measures, it did not consider what these slopes signify. Presumably, search slope reflects an overall processing time *P*: the time needed to load information into vSTM, hold it across the temporal gap, compare it, and—if necessary—unload vSTM and shift processing to the next candidate item(s). If coherence theory is correct and focused attention is involved, processing should never be faster than the speed of attentional search in complex static patterns.

Approach

Before proceeding with the analysis, note that the results of Experiment 1 indicate a potential complication: Search slopes are not always constant. Although the slopes for polarity change were roughly the same with both 80 and 800 msec on-times ($p > .05$), this was not true for orientation change ($p < .0002$). As such, it is important to first determine the extent to which search slope is affected by on-time.

Experiment 2 examined performance at several values of on-time: 80 msec, 160 msec, 320 msec, 480 msec, 640 msec, and 800 msec. Measures for the first and the last of these were taken directly from Experiment 1; measures for the others were obtained from separate sets of observers. Although less powerful than a complete within-observer comparison, this test should provide a rough determination of the extent to which slopes are influenced by on-time.

Experiment 2A examined the case of orientation. As in Experiment 1A, all items were black, with orientations either horizontal or vertical. Experiment 2B examined the corresponding case of polarity. As in Experiment 1B, items were vertically oriented, and could be either black or white.

Results

Slopes for the various conditions are shown in Figure 4. For orientation, there appears to be an overall trend for slope to increase with on-time; a one-way ANOVA confirms that this trend does reach significance, $F(5,66) = 2.38$; $p < .05$. However, this is largely due to the 800 msec condition: When this is removed from consideration, the trend effectively disappears, $F(4,55) = 1.43$; $p > .2$. Between-observer comparisons show that the 800 msec condition is not only slower than the 80 msec condition, but the 160 msec condition as well ($p < .03$). In contrast, no significant differences were found between any on-times of 640 msec or less ($p > .05$ for all pairwise comparisons). Thus, although slope does increase with on-time, this is significant only for on-times longer than 640 msec. For smaller values, no strong trend emerges: Even the best-fitting line has a relatively poor fit ($r^2 = 0.64$). To a first approximation, then, search slope

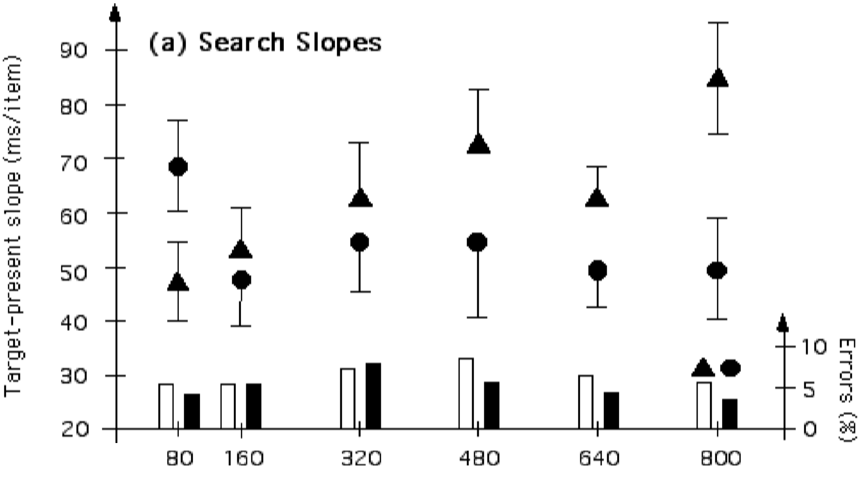

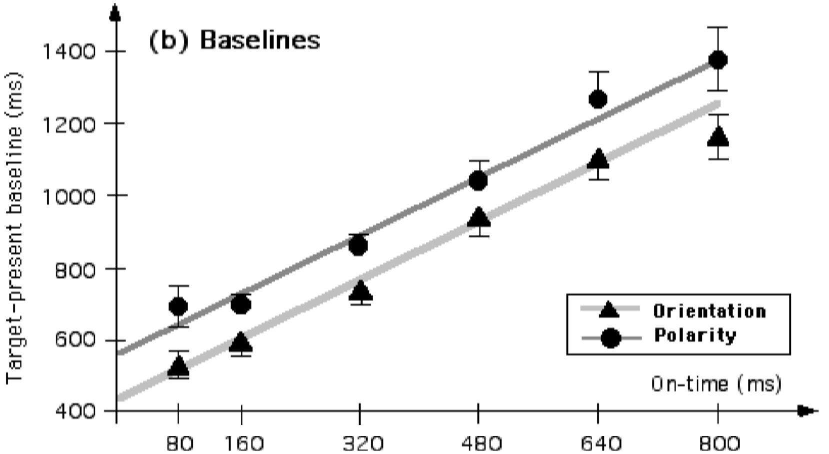

FIG. 4. Results of Experiment 2 (performance as a function of on-time). (a) Search slopes. For both orientation and polarity, there is no great tendency for slopes to increase with on-time, at least for values up to 640 msec. (b) Baselines. These increase in direct proportion to the increase in on-time. (For orientation, best-fitting line has slope 0.95 and intercept 561.8; $r^2 = .98$. For polarity, best-fitting line has slope 1.03 and intercept 568.8; $r^2 = .98$.) In the interests of clarity, orientation baselines have been lowered by 100 msec.

for on-times of 80-640 msec is constant, with an average value of about 60 msec/item.[4]

For polarity, there is no overall tendency for slope to vary with on-time, $F(5,66) = 0.62$; $p > .6$. Although search for the 80 msec condition appears to be slightly slower, between-observer comparisons fail to show that it differs from the speed of any other condition (all *p*s > .1). It therefore appears that search slope is largely unaffected by on-time, at least over the range of 80–800 ms. Within this range, the average slope is about 54 msec/item.

The baseline values for the various conditions are also shown in Figure 4. As is evident from the figure, baselines for both types of change are largely identical. (In the interests of clarity, baselines for orientation have been lowered by 100 msec.) Between-observer comparison showed no reliable differences between the two types of change for any condition (all *p*s > .15). Neither were there any reliable differences between target-present and target-absent baselines (all *p*s > .2). However, baselines did increase linearly with on-time. For orientation—excluding the 800 msec condition—the slope was 1.04 ($r^2 > .99$); for polarity it was 1.03 ($r^2 > .98$). The only exception to this pattern occurred for the 80 msec baselines: no significant differences were found between these and the 160 msec baselines, either for orientation ($p > .2$) or polarity ($p > .9$).

Discussion

These results show that over a fairly wide range of on-times, search speed does not greatly vary with on-time. Under these conditions, speed is evidently governed by intrinsic processing constraints rather than factors such as stimulus quality or memory limitations. The extent of this *processing range* depends on the property involved: for orientation, it extends from 80 to 600 msec; for polarity, it continues past 800 msec.

Assuming that target-present slopes represent half the actual search slope (Chun & Wolfe, 1996; Treisman & Gormican, 1988), and that search slope is due entirely to processing constraints, the processing time P can be readily calculated: for orientation, it is about 120 msec/item; for polarity, 108 msec/item. In accord with the prediction of coherence theory, both of these values are greater than the 30–50 msec/item thought to be needed for attentional shifts in static displays (Julesz, 1984; Wolfe, 1994).

---

[4] The fact that search for changing orientation is so difficult provides evidence that superposition cannot be carried out on unattended structures, for otherwise observers could search for the unique cross-shaped item formed from the horizontal and vertical rectangles. Instead, the old item is completely replaced by the new. Note that this replacement is at the level of proto-objects rather than pixels, for otherwise there would exist remnants of each of the two target items in those areas where they do not overlap, causing search to be easy.

The near-coincidence of the estimates for orientation and polarity suggests that processing time may be much the same for all basic properties (or "features"). If true, this would prove to be useful for determining which properties are features. However, it is unlikely that processing time is the same for all features: Between-observer comparison of the slopes for the 80 msec conditions shows that the difference between orientation and polarity is very close to significance ($p < .06$). More investigation is needed to settle this matter. Among other things, it is necessary to determine if similar processing times are found for other known features (such as colour and size), and if these values depend on the particular shapes of the items.

Although the baselines are not constant, they vary in a simple way, increasing directly with on-time. This behaviour is a direct consequence of the granular nature of the change detection task: An observer must wait for a display alternation before a change can be perceived. Thus, when on-time is increased, so is the average time of the nearest alternation, and consequently, the average time of each response. The failure to find a strong difference between the 80 and 160 msec on-times may stem from the relatively fast (< 300 msec) alternation time in those conditions. If this is smaller than the variance in the decision and the motor processes, it would allow an observer to effectively overcome the quantization effects that occur with slower cadences.

ASPECT 3: CAPACITY

Issues

Another aspect of performance is the capacity C, that is, the maximum number of items that an observer can see change at any one time (i.e. at any single alternation). If on-time is long enough to let the process pick up all the items it can, performance will not be limited by intrinsic processing speed, but by the maximum number of items that can be held across the temporal gap. If coherence theory is correct, this can be identified with the span of attention, or equivalently, the capacity of vSTM (believed to be about five items—see, for example, Luck & Vogel, 1997; Phillips, 1974; Pylyshyn & Storm, 1988). As such, no more than about five items ought to be seen to change at any single alternation.[5]

---

[5] Long term memory (LTM) may help observers hold on to presented material, and so inflate the estimates of short-term capacity. However, given that displays are on for relatively short amounts of time—typically alternating a few times per second—relatively little information is likely to be transferred into LTM. Empirical support for this position comes from the capacity analysis of the results of Experiment 2, which shows that the hold for orientation does not change when display on-time times increase from 640 to 800 msec. Such a plateau would not be expected if a significant amount of information was placed into LTM during the 800 msec of on-time.

Approach

To determine how many items are held across a gap, consider first the case where only one item is held. Starting from the onset of the first new display, one of the items in that display (i.e. the item at the attended location) can be compared with the item held in memory. Assuming that no change has been detected, the contents of memory must then be cleared, attention shifted to a new item, and the contents of the next item entered into memory, after which the process waits for the next alternation to begin. Consequently, only one item on average is examined per alternation, and so the search rate is identical to the alternation rate. Thus, if the display alternates every 100 msec, the search slope will also be 100 msec/item.

Following the same logic, if two items are held across each gap, search should be twice as fast. Thus, if the display alternates every 100 msec, the slope will now be 50 msec/item. More generally,

$$\text{search slope} = \text{alternation time} / \text{hold} \tag{1}$$

Transposing terms, this yields:

$$\text{hold} = \text{alternation time} / \text{search slope} \tag{2}$$

$$= (\text{on-time} + \text{off-time}) / \text{search slope} \tag{3}$$

Hold describes how many items on average are held across each temporal gap. Note that it does not assume memory to be limited to whatever can be extracted from a single display; its contents might be loaded in over several alternations.

Capacity C can be defined as the asymptotic hold that exists with increasing on-time; in this case, the items *are* gathered over a single display. To determine the value for C, equation (3) was used to transform the speeds obtained in Experiment 2 into holds. This was done both for changes in orientation and changes in polarity.

Note that hold provides an alternate way to analyse performance. If search is constrained by processing limitations, search slope s will have a constant value P (i.e. the processing time); equation (2) then indicates that hold will increase linearly with alternation time, the proportionality constant being $1/P$. Similarly, if search is constrained by memory limitations, hold h will have a constant value C (i.e. the capacity); equation (1) then implies that search slope increases linearly with alternation time, the proportionality constant being $1/C$. As such, s and h are *duals*, behaving in similar ways under the two different kinds of resource-limited conditions (Figure 5).

Among other things, this duality provides an alternate test of the extent to which processing time is constant. Via equation (2), an estimate of processing time can be made from the inverse of the slope of the best-fitting line through the estimated holds. (This is a least-squares fit of the data, with hold set to zero at zero alternation time.) This line can be tested for goodness of fit, and the

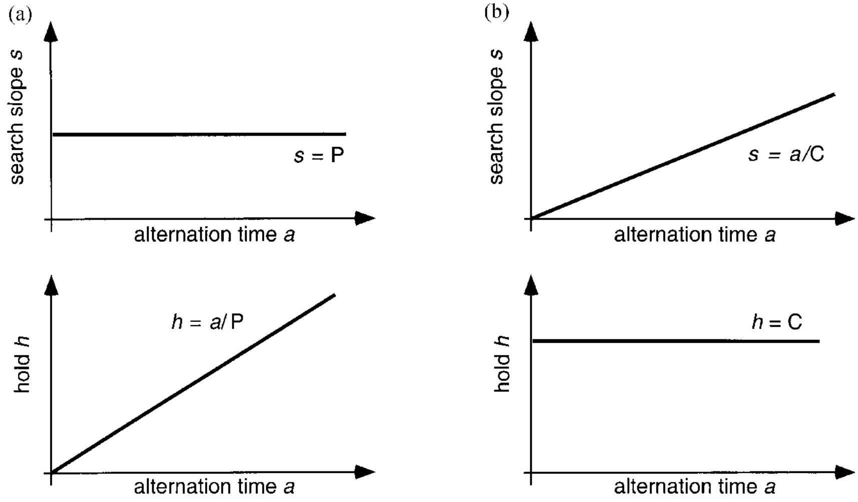

FIG. 5. Duality of hold and search slope. (a) For process-limited search, search slope *s* is a constant value P (the processing time); hold *h* increases in proportion to alternation time *a*, the proportionality constant being 1/P. (b) For memory-limited search, hold *h* is a constant value C (the capacity); search slope *s* increases in proportion to alternation time *a*, the proportionality constant being 1/C.

estimate of processing time made this way can be compared against the direct estimate obtained in Experiment 2.

Results

The dependence of hold on display on-time is shown in Figure 6. For orientation, hold increases linearly with on-time up to around 600 msec, after which it is constant at about 5.5 items. The best-fitting line for on-times in the processing range 80–640 msec has a slope of 7.64 item/sec, corresponding to a processing time of 131 msec/item. This value is slightly higher than the estimate of 120 msec/item obtained in Experiment 2, but the fit to the data is good ($r^2 = .92$), confirming that processing time is indeed approximately constant in this range.

A somewhat different pattern emerges for polarity. Here, hold is roughly the same as for orientation when on-times are short. This is hardly surprising, since according to equation (1) similar processing times correspond to similar holds. But rather than reaching an asymptote, hold continues to increase, reaching a value of about 9 items for on-times of 800 msec. The best-fitting line for on-times in the range 80–800 msec has a slope of 9.59 item/sec, corresponding to a processing time of 104 ms/item, a value very close to the direct estimate of 108 ms/item obtained in Experiment 2. As for the case of orientation, the fit of this line is good ($r^2 = .99$), indicating that processing time is constant.

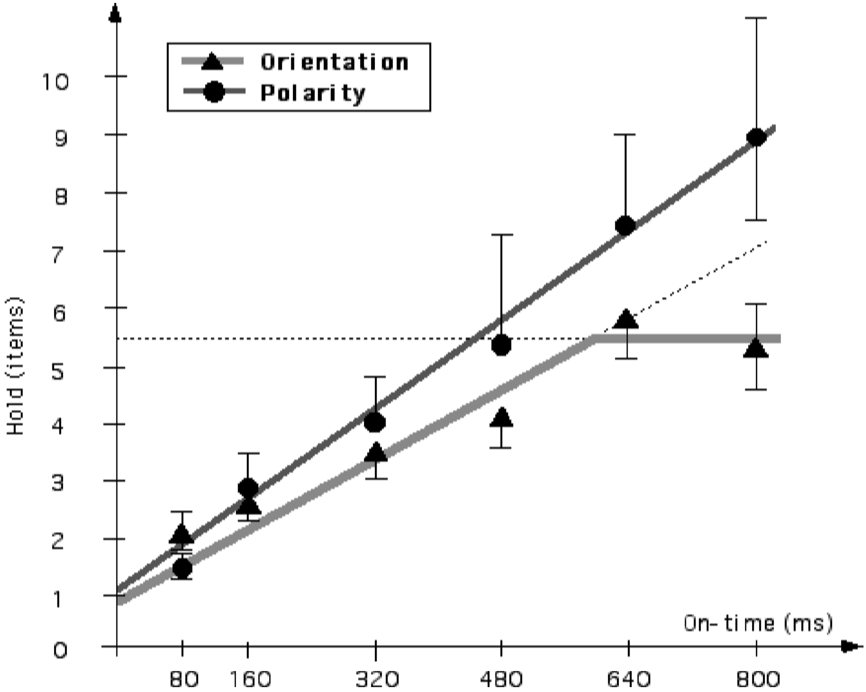

FIG. 6. Hold as a function of on-time. Data here are the search slopes of Figure 4, transformed via equation (3) into holds. For both orientation and polarity, hold increases linearly with on-time. But hold for orientation reaches an asymptote (i.e. capacity) of about 5, whereas polarity hold continues to increase, reaching a value of about 9 at 800 ms. (For orientation, best-fitting line, in the range 80–640 msec, has slope 7.64 items/sec; $r^2 = .92$. For polarity, best-fitting line has slope 9.59 items/sec; $r^2 = 0.99$. Note that because off-time is 120 msec, these lines intersect the x-axis at –120 msec.) Dashed lines are linear extrapolations.

Discussion

Taken together, these results show that orientation and polarity give rise to rather different behaviours. Although both lead to an estimate of hold that increases linearly over a range of on-times, the nature of this linearity is not the same: The hold for orientation asymptotes at about five items, whereas that for polarity continues to increase, reaching a value of about nine items at 800 msec.

Note that this difference is essentially a restatement of the result encountered in Experiment 2, where search for orientation change was found to be process-limited for on-times shorter than 600 msec. Here, it is seen that performance at longer on-times is governed by memory limitations, indicating a transition between process- and memory-limited modes of operation. In contrast, search for polarity change is entirely process-limited—there are no effects of memory limitations, at least for the on-times examined here. Evidently, the capacity for polarity is high, greatly exceeding that for orientation.

The question naturally arises as to why there should be such a divergence in behaviour. The limit of 5.5 items for orientation is readily understandable—it is close to estimates for the capacity of vSTM, and for the span of attention (e.g. Pashler, 1988; Pylyshyn & Storm, 1988). The mystery is therefore the existence of an apparent "supercapacity" for polarity. It is unlikely that eight or nine individual items are actually held in memory. Instead, it may be that there exists a type of grouping for polarity that does not exist (or at least, exist to the same degree) for orientation. This grouping may form larger, more complex "chunks", each of which collects polarity information from several items. In this view, the capacity of the attentional grasp would still be around five items, but the items that can be held are different for different kinds of visual properties. Resolving this issue is a matter for future work. But whatever the reason turns out to be, it is clear that it involves a mechanism that does not treat all basic visual features the same way.

ASPECT 4: BETWEEN-DIMENSION SELECTIVITY

Issues

The final aspect of performance tested here is the ability to selectively filter out variations in irrelevant properties. One of the most important characteristics of focused attention is its ability to selectively filter out irrelevant properties (or dimensions) of a stimulus. Indeed, this aspect is so important that focused attention is sometimes defined in terms of the ability to reject irrelevant messages (e.g. Schneider, Dumais, & Shiffrin, 1984).

As manifest in these experiments, this issue becomes the question of whether search can ignore variations in any property separable from the (changing) feature used to define the target. Two kinds of variation are worth examining here: static and dynamic. In both cases, items vary in their irrelevant properties, that is, the properties that do not define the target. For static variation, items are heterogeneous in the property being varied, but they do not change over time. A more demanding type of interference is dynamic variation, in which these properties are not only heterogeneous but are also changed at each display alternation. In both cases, the selectivity of this kind of filtering can be determined by measuring the extent to which search slope is affected by these variations. Orientation and polarity are known to be separable properties (Garner, 1974). If coherence theory is correct, and if variations in irrelevant properties cause no interference with the formation of lower-level proto-objects, performance should be unaffected by both types of variation.

Approach

Two slightly different experimental designs were used, investigating static and dynamic variation, respectively. Both involved two conditions: (a) orientation change defined the target, and the polarity of all items was varied, and (b) polarity change defined the target, and the orientation of all items was varied.

Experiment 3 examined the case of static variation. In Experiment 3A, the task was to detect change in orientation. In the homogenous displays, items were black rectangles; half were horizontal and half vertical. In the heterogeneous displays, similar rectangles were used, but half were black and half were white. Experiments were blocked, so that observers ran all trials of one condition before beginning the other. Experiment 3B examined the corresponding case of polarity change. Here, the items in the homogeneous displays were vertical rectangles; half were black and half were white. In the heterogeneous displays, half were vertical and half were horizontal. Observers were run (in counterbalanced order) on both homogeneous and heterogeneous displays. As for the case of orientation, experiments were blocked on the basis of heterogeneity. On-times were always 80 msec.

Experiment 4 examined the case of dynamic variation. For all displays, items were rectangles, half of which were black and half white, and half horizontal and half vertical. In non-varying displays, only the relevant property (orientation) changed. In varying displays, the polarity of all items—targets and distractors—changed with each display alternation. Experiment 4A examined orientation change; Experiment 4B examined polarity change. For all conditions, on-times were 80 msec. As in the static case, observers were run on both varying and non-varying displays. Experiments were blocked so that observers completed all trials of one condition before starting the other.

Results

Results for Experiment 3 (static variation) are shown in Figure 7. For orientation change, the slope for homogeneous displays was 47 msec/item, while that for heterogeneous displays was 54 msec/item, an insignificant difference ($p > .1$). As such, static polarity variations appeared to have little effect. The results for polarity change were similar. The rate for homogeneous displays was 61 msec/item, whereas that for heterogeneous displays was 67 msec/item. This difference was not significant ($p > .4$), indicating that static orientation variations had no great effect.

The results for Experiment 4 (dynamic variation) are shown in Figure 8. For orientation, search in non-varying displays was 53 msec/item, whereas search in varying displays was 50 msec/item, an insignificant difference ($p > .4$). The case of polarity change showed a small slowdown of 70 msec in the baselines of the varying displays ($p < .05$)—presumably the changing orientations made it

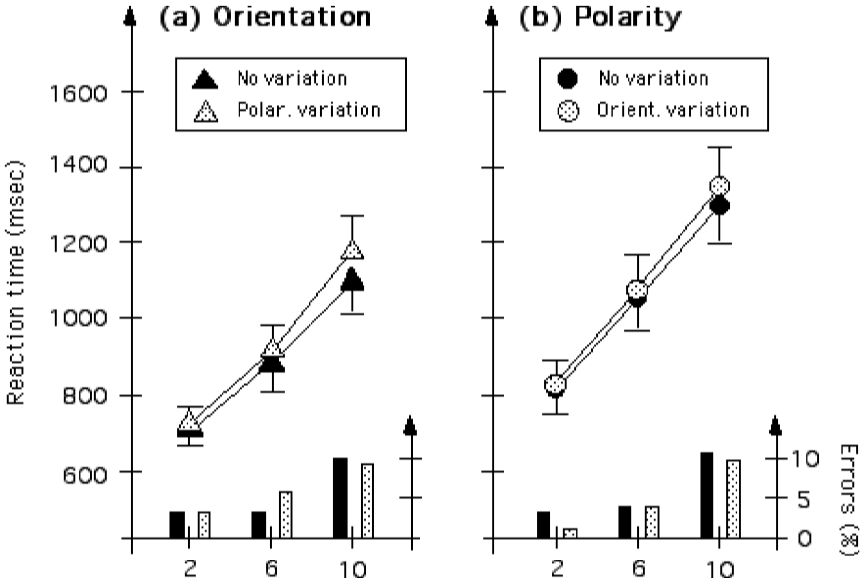

FIG. 7. Results of Experiment 3 (effect of static variation). (a) Effect of irrelevant static variation on search for orientation change. (b) Effect of irrelevant static variation on search for polarity change. As is evident, there is no effect of variation for either type of change.

slightly more difficult to verify targets. But more importantly, search slopes remained unaffected: The slope for non-varying displays was 49 msec/item, whereas that for varying displays was 56 msec/item ($p > .3$). Thus, the dynamic variation of irrelevant properties appeared not to interfere with search.

Discussion

The results of Experiments 3 and 4 are clear: Variation in separable properties—whether static or dynamic—can be filtered out quite well. As a methodological note, it is worth pointing out that (as for the case of processing time) the technique used here is quite general, and may provide a useful way to determine which pairs of visual properties are separable from each other.

But more importantly for present purposes, these results verify yet another prediction made from the assumption that attention is needed to see change. This verification—together with the compatibility found for the other aspects examined—indicates that visual search for change is not just something formally similar to search for complex static patterns, but, rather, involves attentional processes that are largely the same.

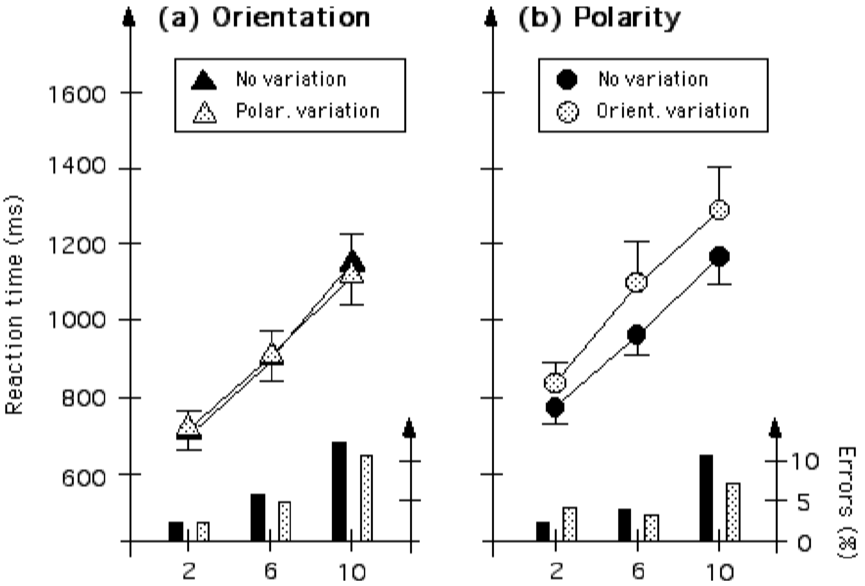

FIG. 8. Results of Experiment 4 (effect of dynamic variation). (a) Effect of irrelevant dynamic variation on search for orientation change. (b) Effect of irrelevant dynamic variation on search for polarity change. As is evident, only polarity change is affected by dynamic variation; this is limited to a slight (70 msec) increase in baselines.

ASPECT 5: SPATIAL MEMORY TRACE

Issues

Up to this point, the granular nature of the change-detection task has been largely ignored—indeed, Experiment 1 was expressly designed to show that granularity can be disregarded when measuring slopes and baselines. But the natural partitioning of reaction times in this task allows measurement not only of their averages, but also of their distribution after each alternation. Such a perspective is unique to change detection,[6] and may prove useful for analysing various aspects of attentional processing.

---

[6] A distribution of RTs can of course be obtained from search for static targets, as well as from "one-shot" change-detection tasks. However, the flicker paradigm described here is unique in that multiple distributions are obtained—one for each *epoch* (i.e. each interval corresponding to a particular display change). It is the distribution of responses *between* epochs that is of interest here.

As an example of this, consider the issue of whether attention leaves a memory trace that allows it to avoid items already examined.[7] Recent studies argue that no such trace exists; in such a case, attention can revisit previously attended items, and so the true speed of search will be about twice the measured speed (Horowitz & Wolfe, 1998). However, if speed is doubled, equation (3) forces the estimates of attentional capacity to be doubled as well. The estimate for orientation will then be 10–11 items, a value rather discrepant with other measures of attentional limits.

It is important to note here that capacity itself is not changed by the existence or non-existence of a memory trace. In visual processing, at least two different kinds of memory mechanisms exist. The first is the *hold* of the coherence field; this is effectively a memory of “what”, existing only as long as attention is focused on the constituent items. The second is the *trace* of locations that have been attended; this is effectively a memory of “where”, and being non-attentional, may last for a considerable time. Indeed, such a trace may be related to the representation of layout in a scene (Rensink, this issue).

Thus, the existence of a memory trace does not affect change-detection capacity *per se*. What it does affect is the *estimate* of this quantity. The key issue here is whether the original estimates of capacity are valid. If so, this will indicate that measured speed is similar to true speed, and consequently, that a memory trace does exist. If not, true speed will be twice measured speed, and the estimates obtained earlier will need to be doubled. In this latter case, the identification of change-detection capacity with attentional capacity will be compromised, and thus challenge the view that search for change is carried out by the same attentional mechanism as used for static patterns.

Approach

One way to determine whether the estimates of capacity are valid is to examine the distributions of reaction times within a trial. This can be done by considering the fraction of responses made in each *epoch*, that is, each interval between the appearance of a given display and the appearance of the following one. The responses analysed here will be those for the search for orientation change in Experiment 1, with on-times of 800 msec. Performance for this condition is

---

[7] The issue of memory effects is most easily described in terms of a serial model of attention, in which attention operates on an item-by-item basis. Assuming such a model, the question is whether the attention ever returns to an item that it has previously encountered. Because of its simplicity, this description will be used when discussing the issue of memory traces. However, this is not meant to imply that parallel models are to be dismissed in this regard. For a parallel model, attention acts concurrently on all items, and so the idea of a previously attended item is no longer relevant. But memory effects are still an issue. For example, one question is whether the search process has to restart if all items in the display change their location (Horowitz & Wolfe, 1998). Thus, this concern does not depend on whether the search process is serial or parallel.

memory-limited, with response distributions determined by change-detection capacity.

If there is no memory trace, the capacity estimate is about 11 items. In this case, it should be possible to load all items of a display into vSTM during the 800 msec of the initial display (epoch 0), and compare them during the 800 msec of the next display (epoch 1). More precisely, since the total display time is 1600 msec (even more if the intervening off-time is included), and the processing time is 120 msec/item, it should be possible to process at least 1600/120 = 13.3 items, more than on any display.[8] Consequently, responses ought to be found mostly within the first epoch, with perhaps a few stragglers in the succeeding ones. This will apply to both target-present and target-absent responses for all set sizes.

In contrast, if a memory trace does exist, the capacity estimate is about five items. In this case, target-present responses for set size 10 will be distributed across the first two or three epochs. Meanwhile, target-absent responses will be mostly gone from epoch 1—determining absence requires that all items be examined, and this will not be complete until at least the second epoch. For set size six, behaviour will depend on the capacity of the individual observer. For observers with a capacity of six or more, all responses can be made during epoch 1. For observers with a capacity of less than six, target-present responses will be distributed over at least two epochs, whereas target-absent responses will be largely absent from epoch 1. Finally, for set size two, all responses should occur during epoch 1.

Results

Distributions of target-present and target-absent responses were determined for each observer. Data were pooled for the six observers with a capacity of six or more (high-capacity observers; median = 8.0 items), and for the six observers with a capacity of less than six (low-capacity observers; median = 4.2 items).

Distributions for the high-capacity observers are shown in Figure 9. Most responses for set sizes two and six fall in epoch 1; this is expected from both views of attentional memory. More interesting is set size 10. Here, target-present responses still fall mostly in epoch 1 (66.0%), but there is now an appreciable presence in epoch 2 (29.0%). Importantly, target-absent responses are less common in epoch 1 (31.2%) than in epoch 2 (44.9%). Although the number of responses in epoch 1 is not small, two of the six observers had capacities greater than 10; if their responses are removed, the proportion of epoch 1 responses

---

[8] This conclusion still holds if the slower rate of 131 msec (obtained from the least-squares fit of hold estimates) is used, since there will be enough time to process at least 1600/131 = 12.2 items.

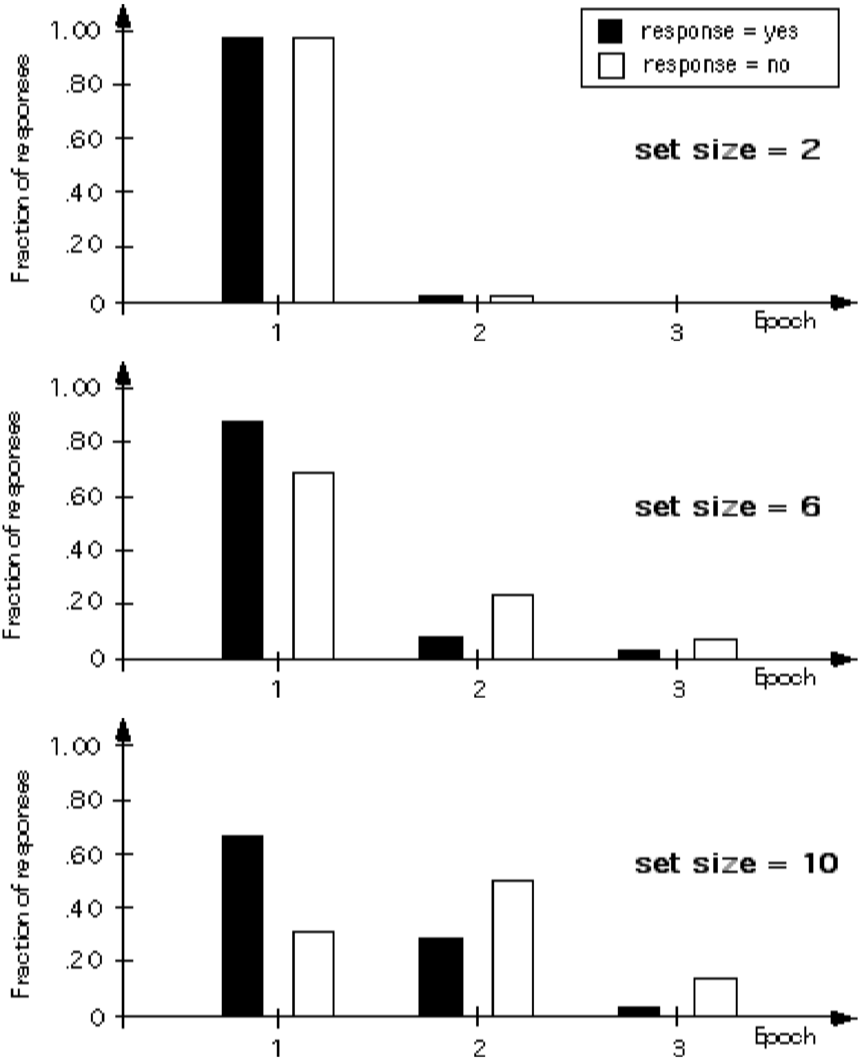

FIG. 9. Distribution of responses for high-capacity observers. For set sizes 2 and 6, most responses occur during epoch 1. For set size 10, target-present responses begin to appear in epoch 2, whereas more target-absent responses are found in epoch 2 than in epoch 1. This pattern is consistent with a median capacity estimate of 8.0 items, but not 16.0 items.

falls to 14.6%. As such, this pattern indicates that the median capacity for high-capacity observers is much closer to 8 than it is to 16.

Distributions for the low-capacity observers are shown in Figure 10. As before, most responses for set size two fall in epoch 1. Importantly, effects due to capacity limits already begin to appear at set size six: Most target-present

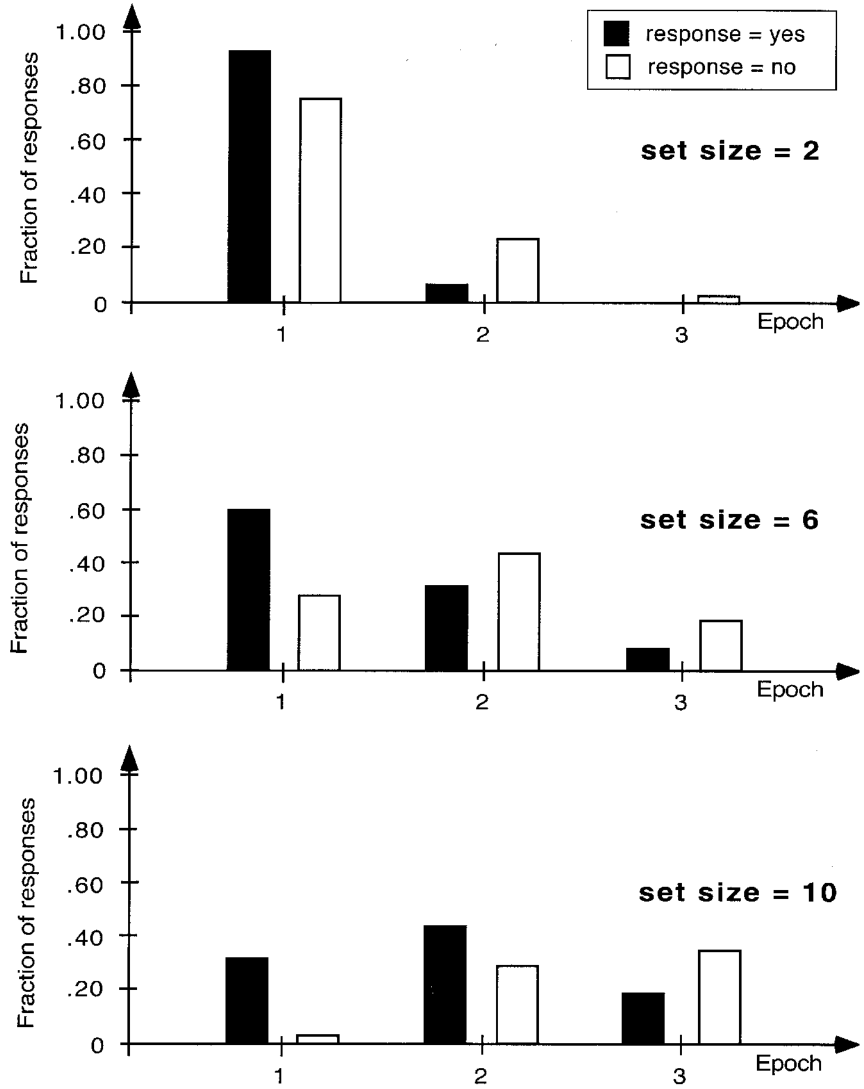

FIG. 10. Distribution of responses for low-capacity observers. For set size 2, most responses occur during epoch 1. For set size 6, target-present responses begin to appear in epoch 2, whereas more target-absent responses are found in epoch 2 than in epoch 1. This pattern is amplified for set size 10, where most target-present responses are spread over the first three epochs, and target-absent responses occur most often in epoch 3. This pattern is consistent with a median capacity estimate of 4.2 items, but not 8.4 items.

responses fall in epoch 1 (59.0%), but a considerable number now fall in epoch 2 (32.2%); meanwhile, target-absent responses are less common in epoch 1 (27.5%) than in epoch 2 (43.2%). Both trends are heightened for set size 10. Fewer target-present responses fall in epoch 1 (31.4%) than in epoch 2 (43.7%), and there is now an appreciable presence in epoch 3 (18.0%). Target-absent responses are almost completely absent from epoch 1 (2.8%). Interestingly, fewer responses occur in epoch 2 (28.4%) than in epoch 3 (33.5%), with a considerable number also falling in epochs 4 and 5 (34.2%).[9] This pattern is consistent with observers having a median capacity of 4.2, but not 8.4.

Discussion

The distributions of responses across epochs show that the estimates of capacity derived from measured search slopes are at least approximately true. At the very least, capacities are certainly not twice the estimated values. Thus, the identification of change-detection capacity with attentional capacity can be maintained. And so the position developed here remains consistent.

But if the capacity estimates are valid, this implies that measured search speed is the same as true search speed. And this in turn implies that the attentional process examined here does have a memory of where it has been.

This conclusion appears to be at odds with work showing that visual search involves no memory trace (Horowitz & Wolfe, 1998). However, the task here is different in several ways. To begin with, set sizes are smaller: Whereas Horowitz and Wolfe used sizes as high as 16, sizes here extended only to 10. It may be that a limited amount of memory exists, enough to store several locations; if so, such a memory would be useful for the relatively small displays used here, but not for larger ones. Alternatively, the critical factor could be the nature of the task itself. Horowitz and Wolfe used a static target, which required determining shape but not location. In contrast, the task here requires observers to compare successive items at the same location. For this, information about the exact location of an item is vital, and it may be that obtaining this information somehow produces a memory trace. However this situation is resolved, the outcome should cast new light on our understanding of attentional processing.

SUMMARY

The experiments presented here provide considerable evidence for the proposal that focused attention is required to see change. Experiment 1 showed that visual search for change is a self-terminating process that requires time linear in

---

[9] The distributions are normalized by a count that includes errors (such as responses made before epoch 1) and long RTs (responses made after epoch 5). Consequently, the values for the epochs described can add up to less than 100%.

the number of items in the display. This establishes that search slopes and baselines are valid measures of performance for this task, so that it is formally similar to the attentional search for complex static patterns.

The results of subsequent experiments showed that this similarity is more than just formal. First of all, processing times for orientation and contrast polarity were found to be approximately constant for on-times between 80 and 640 msec, with a value of about 100–120 msec/item. Although these values are higher than those for most attentional searches of static displays, this is to be expected: The detection of change requires operations not needed for detecting static patterns (e.g. loading information into vSTM), and so processing each item simply requires more time.

Another result indicating that search for change is mediated by focused attention is the finding that only a limited amount of information can be held across a temporal gap. The estimate found here for orientation—5.5 items—is consistent with other estimates of attentional capacity (e.g. Pylyshyn & Storm, 1988). Interestingly, estimates based on polarity change lead to an apparent "supercapacity" of at least nine items. Presumably, polarity enables some form of grouping, so that even though only five items are held, each may be a "chunk" that collects information from several structures in the display. But whatever the explanation, the finding that properties similar in regards to processing speed are not similar in regards to capacity indicates that simple features are not treated alike by all visual mechanisms. Evidently, a divide of some kind exists.

The final piece of evidence stems from the ability of observers to ignore variations in irrelevant properties. It was found that both static and dynamic variations in irrelevant properties could be effectively filtered out, leading to search slopes similar to those for non-varying displays. Given that these four aspects of the search process are compatible with known aspects of focused attention, it would appear that the mechanisms involved in visual search for change are largely—if not entirely—the same as the attentional mechanisms involved in visual search for complex static patterns.

As a demonstration that the granular nature of these experiments can be used to explore other, less-understood aspects of focused attention, the distribution of responses across different alternations was analysed. It was found that the estimates of attentional capacity made here were consistent with those derived via other types of study. This was then used to argue that the attentional process involved in change detection does create a memory trace of where it has been.

Taken together, the implication of these results is clear: Visual search for change involves a limited-capacity process that is not just formally similar to the attentional process used for static patterns, but involves many of the same underlying mechanisms. Visual search for spatiotemporal patterns can therefore be considered to be a direct extension of visual search for static patterns. As such, we not only have a sound theoretical framework that links change blindness to focused attention—we also have a sound methodological framework that can let us use this phenomenon to explore the various attentional processes involved in our perception of the world.

REFERENCES

Chun, M.M., & Wolfe, J.M. (1996). Just say no: How are visual searches terminated when there is no target present? *Cognitive Psychology, 30,* 39–78.

Cowan, N. (1998). Evolving conceptions of memory storage, selective attention, and their mutual constraints with the human information-processing system. *Psychological Review, 104,* 163–191.

Garner, W.R. (1974). *The processing of information and structure*. Potomac, MD: Lawrence Erlbaum Associates Inc.

Horowitz, T.S., & Wolfe, J.M. (1998). Visual search has no memory. *Nature, 394,* 575–577.

Julesz, B. (1984). A brief outline of the texton theory of human vision. *Trends in Neuroscience, 7,* 41–45.

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

Luck, S.J., & Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–280.

Pashler, H. (1988). Familiarity and visual change detection. *Perception and Psychophysics, 44,* 369–378.

Phillips, W.A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics, 16,* 283–290.

Pylyshyn, Z.W., & Storm, R.W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision, 3,* 179–197.

Rensink, R.A. (1996). The attentional capacity of visual search under flicker conditions. *Perception, 25*(Suppl.), 2.

Rensink, R.A. (1997). How much of a scene is seen? The role of attention in scene perception. *Investigative Ophthalmology and Visual Science, 38,* 707.

Rensink, R.A. (this issue). The dynamic representation of scenes. *Visual Cognition, 7,* 17–42.

Rensink, R.A., & Enns, J.T. (1995). Preemption effects in visual search: Evidence for low-level grouping. *Psychological Review, 102,* 101–130.

Rensink, R.A., & Enns, J.T. (1998). Early completion of occluded objects. *Vision Research, 38,* 2489–2505.

Rensink, R.A., O'Regan, J.K., & Clark, J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368–373.

Sagi, D., & Julesz, B. (1985). “Where” and “what” in vision. *Science, 228,* 1217–1219.

Schneider, W., Dumais, S.T., & Shiffrin, R.M. (1984). Automatic and control processing and attention. In R. Parasuraman & D.R. Davies (Eds.), *Varieties of attention* (pp. 1–27). Orlando, FL: Academic.

Simons, D.J. (1996). In sight, out of mind: When object representations fail. *Psychological Science, 7,* 301–305.

Simons, D., & Levin, D. (1997) Change blindness. *Trends in Cognitive Sciences, 1,* 261–267.

Townsend, J.T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science, 1,* 46–54.

Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95,* 15–48.

Ulrich, R., & Giray, M. (1989). Time resolution of clocks: Effects on reaction time measurements—good news for bad clocks. *British Journal of Mathematical and Statistical Psychology, 42,* 1–12.

Wolfe, J.M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review, 1,* 202–238.

Wolfe, J.M. (1996). Post-attentive vision. *Investigative Ophthalmology and Visual Science, 37,* S214.

Zelinsky, G.J. (1997). Eye movements during a change detection search task. *Investigative Ophthalmology and Visual Science, 38,* 373.

APPENDIX A: Sensitivity of Performance to Off-time

In the experiments here, off-times (ISIs) are always 120 msec. Earlier studies with grid-like patterns showed that if this value is used, differences in the motion signals created by the changes are not large enough to draw attention to the target (Phillips, 1974). Indeed, studies with letter arrays show that off-times as short as 67 msec may be sufficient for this (Pashler, 1988). But the displays used here are somewhat different than these, and so may still be affected by low-level motion signals, either as motion energy picked up by low-level sensors or as apparent motion created by the low-level correspondence of successive items. Note that the simple existence of such signals is not the issue here: rather, it is whether the signals generated by the flickering targets are sufficiently *distinct* from those generated by the flickering distractors. If so, the increased salience of the targets might cause search to be speeded up, thereby skewing estimates of attentional properties.

To determine if this might be the case, a set of control experiments had observers search for change in an array of rectangles, much as in Experiment 1. On-times were always 160 msec; off-times were 80, 120, or 160 msec. Ten observers participated in each experiment. Between-observer comparisons were carried out on both target-present and target-absent slopes.

The results are shown in Figure A. For orientation, off-times of 120 msec led to slopes of 53.4 msec/item (target present) and 101.7 msec/item (target absent). Reducing off-times to 80 msec led to slopes of 47.9 and 82.0 msec/item, an insignificant difference ($p > .6$; $p > .4$). Increasing off-times to 160 msec led to slopes of 56.0 and 115.0 msec/item, which again were not significantly different ($p > .7$; $p > .6$). Even the differences between the 80 and 160 msec conditions were not found to be significant ($p > .4$; $p > .15$). A one-way ANOVA showed a similar lack of effect both for target-present slopes, $F(2,27) = 0.06$; $p > .9$, and target-absent slopes, $F(2,27) = 0.88$; $p > .4$.

A similar pattern was found for polarity. Off-times of 120 msec led to slopes of 48.3 and 66.8 msec/item. Reducing off-times to 80 msec yielded slopes of 45.4 and 64.2 msec/item, an insignificant difference ($p > .8$; $p > .8$). Increasing off-times to 160 msec led to slopes of 49.5 and 90.4 msec/item. Although the target-absent slopes here were somewhat larger, neither slope differed significantly from its corresponding 120 msec value ($p > .9$; $p > .2$). The 80 msec and 160 msec conditions were not found to be significantly different ($p > .7$; $p > .1$). A one-way ANOVA also indicated no significant effect of off-time either for target-present or target absent slopes, $F(2,27) = 0.32$; $p > .7$; $F(2,27) = 1.42$; $p > .2$, respectively.

Taken together, these results show that differences in off-times do not cause large differences in speed for the particular cadences tested here. This is unlikely to occur if motion signals affect search rate—in particular, it would not be expected that speeds for the 80 msec and 160 msec conditions would differ so little. Thus, for off-times of 120 msec there would seem to be little influence of the motion signals generated by the flickering display items.

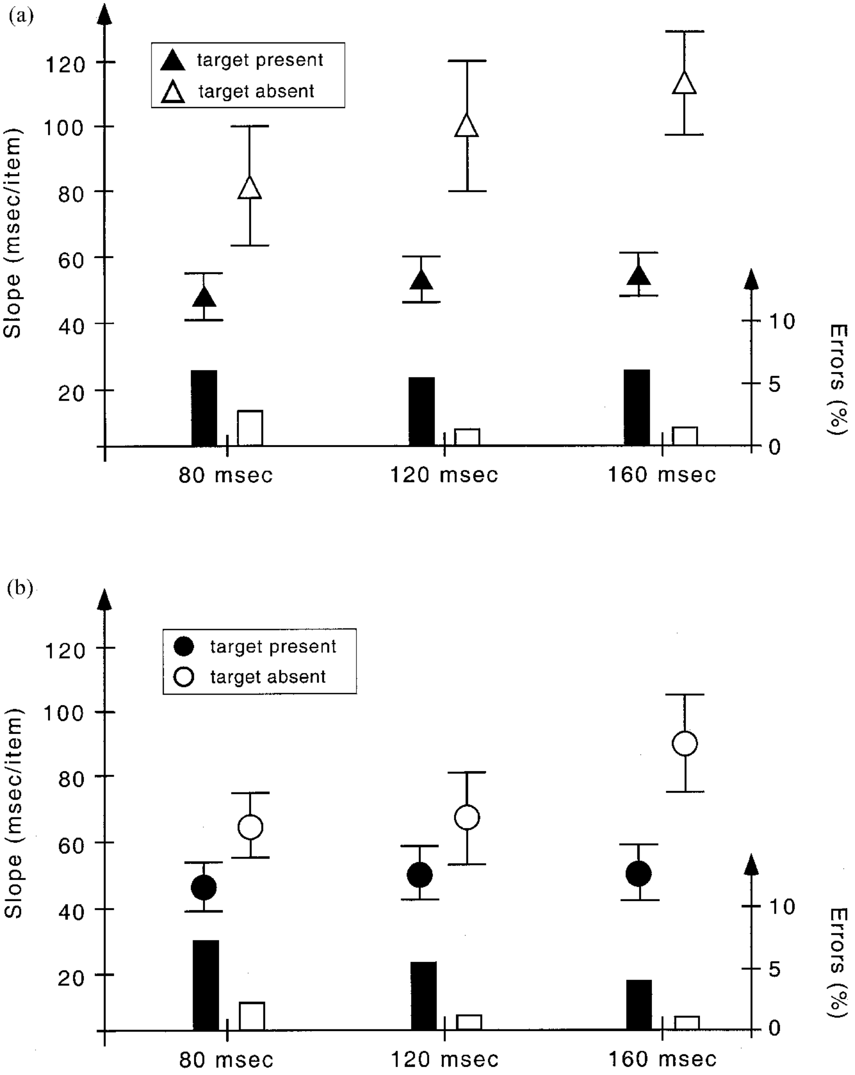

FIG. A. Change-detection performance as a function of off-time. (a) Orientation change. For all three off-times tested, target-present slope was essentially the same. Target-absent slope tended to increase somewhat with off-time, but this increase was not significant. (b) Polarity change. Again, target-present slopes for all three off-times were much the same, with target-absent slopes showing a slight (though insignificant) increase with off-time.

APPENDIX B: Speed–Accuracy Trade-offs

In the experiments carried out here, some conditions are particularly difficult. Two general strategies are usually available to an observer to handle such a situation: decrease speed or decrease accuracy. Since either of these will do, speed and accuracy can be traded off in various proportions. This presents a potential problem for the analyses of the results here, since the determination of some aspects of attention (e.g. capacity) depends upon the absolute values of the search slopes, and not just a relative speedup or slowdown. As such, it is important to determine if the search rates in the experiments here are affected by a speed-accuracy trade-off.

The most obvious possibility in this regard is that the observer only checks a sub-set of the items and then guesses a response, without having actually seen the target. The errors found in the main set of experiments are mostly false negatives (target misses), something that might be expected if the observer checks a subset of the items and then responds "no" if the target is not seen. If this were the strategy used, the speeds measured in the experiments would be artificially high: Although an increase in RTs would be caused by the greater number of items scanned, this number would be less than that actually in the display.

To determine if such a trade-off is being made, 12 observers were tested on two variants of a difficult search experiment. Here, items were rectangles similar to those used in the main experiments. Half the items were vertical and half horizontal, and half were black and half white, with all combinations being represented equally. To make this task as demanding as possible, *distractors* continually changed orientation, whereas *targets* maintained a constant orientation. In the *detection* variant, a constant-orientation target was present half the time and absent half the time. Observers simply reported whether the display did or did not contain a target. In the *identification* variant, a constant-orientation target was always present; half the time it was white, and half the time it was black. In this variant, observers reported the colour of the target. For both variants, on-times were 640 msec and off-times 120 msec.

The detection variant is essentially a search task similar to that in the main experiments reported here, except with a target chosen to make search as difficult as possible. The results of this experiment are shown in Figure B. As is evident, this was an extremely difficult task, with target-present slopes averaging 254 msec/item, and target-absent slopes 468 msec/item. As for case of the main experiments, errors were mostly misses; errors were as high as 16% when 10 items were displayed.

To see whether the high error rate was due to observers skipping items, the same subjects were run (in counterbalanced order) on the identification task. In this variant, a non-changing item was always present, so that a definite termination condition always existed. As Figure B shows, search here was again difficult: Target-present slopes for both types of target averaged 268 msec/item. Importantly, error rates here were low: Average error was only 3.6%; even in the worst condition, errors did not exceed 5%. Thus, speed in this variant was not obtained by skipping items.

Comparing performance for identification against that for detection shows that identification is generally about 300 msec faster. However, within-observer comparison indicated that search slopes did not differ significantly ($p > .25$). Thus, the speed for the detection (a task requiring most of the same operations as identification) was unlikely to be due to the high error rate. Rather, errors would appear to be due to observers attending to all the items, but not always processing them sufficiently to detect the presence (or absence) of change. As such, the search slopes measured in the main experiments here need not be corrected to compensate for a speed–accuracy trade-off.

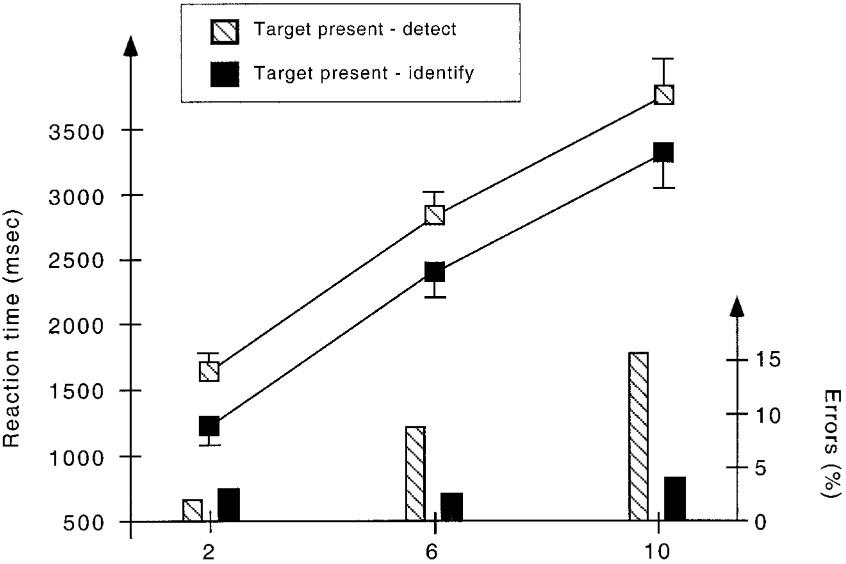

FIG. B. Detection vs. identification. For both tasks, target-present slopes are largely the same, whereas baselines for identification are lower by about 300 msec.