Published in final edited form as:

Cogn Sci. 2004 ; 28(2): 241–258. doi:10.1016/j.cogsci.2003.12.002.

Mapping visual attention with change blindness: new directions for a new method

Peter U. Tse *

Department of Psychological and Brain Sciences, Dartmouth College, HB 6207, Moore Hall, PBS, Hanover, NH 03755, USA

Abstract

Change blindness provides a new technique for mapping visual attention with unprecedented spatial and temporal resolution. Change blindness can occur when a brief full-field blank interferes with the detection of changes in a scene that occur during the blank. This interference can be overcome by attending to the location of a change. Because changes are detected at attended locations, but not at unattended locations, detection accuracy provides an indirect measure of the distribution of visual attention. The likelihood of detecting a new element in a scene provides a measure of the occurrence of attention at that element's location. Potential new directions, advantages, and problems with this method are considered.

Keywords

Visual attention; Change blindness; Cognitive neuroscience

1. Introduction

Determining how attention operates at the level of neuronal circuitry remains one of the central unresolved problems of cognitive neuroscience. In order to determine the role of attention in neuronal information processing, the basic spatiotemporal characteristics of attention must first be mapped in detail. Perhaps the most basic task is to characterize how the spatial distribution of attention changes over time. Surprisingly, this has only been done in a rudimentary way by past researchers. Previous psychophysical mappings of the distribution of visual attention during a task have sampled only a few points in the visual field (e.g., Hughes & Zimba, 1985; the sole exception being Bennett & Pratt, 2001). This may in part account for the lack of consensus regarding how attention is spatially distributed and how that distribution changes over time. Because past maps of the spatiotemporal dynamics of attention have been so low in spatial and temporal resolution, numerous models of attention have arisen that are consistent with this poor data, but inconsistent with each other. Thus, researchers debate over whether attention should be modeled as a “spotlight,” “window,” or “zoom lens,” and they debate whether it moves in an analog or quantal fashion (e.g., Eriksen & St. James, 1986; Posner, Snyder, & Davidson, 1980). They debate whether attention can operate on features other than the locations of visual space, such as motion, shape, or color (e.g., Driver & Baylis, 1989). And they debate whether attention adheres primarily to moving objects or their locations (Duncan, 1984). My coworkers and I (Tse, Sheinberg, & Logothetis, 2003) have recently introduced a new method for mapping attention with high spatial and temporal resolution that may help settle some of these old debates. This method exploits the fact that attention has to

*E-mail address: Peter.Tse@dartmouth.edu (P.U. Tse).

be allocated to the location of a change in order for that change to be detected in a change blindness paradigm (Rensink, O'Regan, & Clark, 1997). The data described here have been described elsewhere (Tse et al., 2003), but will be summarized again here in order to clarify the method. In this paper I focus on the advantages, limitations, and possible future applications of this new method for mapping visual attention. This new method allows researchers to carry out the detailed mapping of visual attention necessary to distinguish among and generate new models of visual attention.

Change blindness can occur in many ways. Although variants of the phenomenon have been known for many years (e.g., Neisser & Becklen, 1975; Pashler, 1988; Sperling, 1960), there has been a recent revival of research on change blindness. Renewed interest followed the observation that changes in a scene that occur during a saccade are typically not noticed (McConkie & Currie, 1996). Others soon noted that changes that occur during a full-field blank or mask (Rensink et al., 1997) or partial field mask (O'Regan, Rensink, & Clark, 1999) are also difficult to detect. Others have shown that changes need not be masked by a blank or by a saccade at all. Scene changes that occur in plain view, but which are not attended (Becklen & Cervone, 1983; Haines, 1989; Mack & Rock, 1998; Rensink et al., 1997; Simons & Chabris, 1999) or which occur too slowly to be picked up by motion-energy detectors (Auvray & O'Regan, 2003; Simons, Franconeri, & Reimer, 2000) are also not typically noticed. Perhaps the simplest version of this phenomenon has been published as a game in magazines since the advent of photography. Two nearly identical photographs are placed side by side, and the viewer must localize the one thing that is different between them. This is not typically thought of as change blindness, but the difficulty of this task arises for the same reason that change localization is difficult in more modern paradigms. The location of the change is not detected by motion-energy detectors, which operate in parallel, so must be localized by serial comparisons made across successive scenes.

In the method described here there is a brief high-luminance full-field blank between two successive, overlapping images that differ (Rensink et al., 1997). An observer might not notice a prominent change in the image during the global blank even when searching for the change, because the global transient introduced by the blank undermines detection of the local transients used by the visual system to locate sudden changes. In particular, the first-order-motion (i.e., motion-energy) detectors that are usually used by the visual system to detect sudden motions, onsets, or changes in luminance are overwhelmed by the fact that a change has happened simultaneously at multiple locations. If many motion-energy detectors are activated by a global change, an isolated motion-energy detector's signal cannot "pop out" on a salience map and capture attention. A global flash interferes with change detection because the visual system's low-level change detectors have been incapacitated. When motion-energy detectors fail to provide a unique signal at a single location, differences between the pre- and postblank images are only detected at a given location when attention has been allocated to that location in both images, because only then can corresponding elements at that location be compared (Rensink, 2000b; Rensink et al., 1997; Simons, 1996; Simons & Levin, 1998). This comparison or monitoring of object states or locations over time is a defining characteristic of attention.

Attention permits serial comparisons because observers are able to select a subset of all the information in iconic memory for further operations, such as reporting (Sperling, 1960) or tracking (Pylyshyn & Storm, 1988). Without this selection of a location for an extended duration, information in the iconic store is simply lost as incoming sensory input continually replaces existing information in the iconic buffer. Although attention is generally thought to bind low-level features into spatially coherent objects (for reviews, see Treisman, 1998, 1999), attention can more accurately be thought to temporally bind successive spatially indexed representations into objects that can be tracked through space-time. Without this spatio-temporal binding, changes could not be noticed because without attention the contents of the

visual iconic store would be disjoint (Horowitz & Wolfe, 1998) and could therefore not be compared with later contents of that buffer. Once selected and bound by attention, however, new information can be processed as a change in a monitored location or tracked figure (Kahneman, Treisman, & Gibbs, 1992). While an array of first-order motion (i.e., motion-energy) detectors can rapidly detect the location of a single change because detection occurs in parallel, attention-based change localization is not parallel, and therefore not rapid. Attention cannot be allocated everywhere to the same degree, and comparisons must be made with past stimuli stored in the iconic buffer.

Within this theoretical framework, attention can be defined as enhanced processing over a limited subset of sensory information that has been selected for monitoring or tracking over time. Unattended information is either not monitored or it may be actively suppressed (e.g., Tipper, 2001). Because the global blank used in the standard change blindness paradigm (Rensink et al., 1997) undermines the visual system's (normally) automatic ability to detect changes on the basis of local transients, it can only detect changes by relying on attentional monitoring. The key assumption behind the present method is the view that the probability of change detection at a given location indicates the degree of attention at that location.

The basic experimental design, shown schematically in Fig. 1, exploits the hampered change detection that occurs in a flicker-induced change-blindness paradigm to create a map of how attention is redistributed in response to a peripherally flashed cue. Tse et al. (2003) used this method to answer the following question: How does the distribution of visual attention change in response to a task-irrelevant cue that flashes suddenly in the periphery?

2. Method

The method has been described in greater detail elsewhere (Tse et al., 2003), but is summarized here. Four participants (24–38 years of age) carried out the experiment. Three were paid observers who were naive regarding the purpose of the experiment, and one was the author. The naive observers completed approximately 100 h of training and testing over a 4-month period.

3. Apparatus

The visual stimulator was a dual-processor Pentium II workstation running Windows NT 4.0. The screen resolution was 1152×864 pixels, and the frame rate 85 Hz. Observers rested their chin in a chin rest. Fixation was ensured using an eyetracker (Sensomotoric Instruments, Berlin, Germany; Tse, Sheinberg, & Logothetis, 2002a, 2002b). Any time the observer's monitored eye was outside a fixation window with a 1.5° radius, the trial would be automatically aborted, and a new trial would be chosen at random from those remaining. If three trials were aborted in a row, the state system automatically reverted control to the eyetracker's calibration program. Once calibration was completed, the experiment resumed with a random trial.

4. Stimuli

The onset of a trial was indicated by the offset and immediate re-onset of the fixation point against the black background. The fixation point was a yellow circle 0.15° in diameter. The circular background, shown in Fig. 1b, was 30° in diameter and uniform black (<1 cd/m²). It spanned the height of the monitor at a viewing distance of 57 cm. The background outside this circular region was dark gray. A 23×23 array of square positions fit within a $30^\circ \times 30^\circ$ square that circumscribed this circular "window." On a given trial, half of these positions were occupied by red and green squares ($0.69^\circ \times 0.69^\circ$); the shades of red and green used were equiluminant as measured by a photometer (Minolta CRT color analyzer CA-100). The probability that a square was red or green was 50%. Squares never overlapped. Their centers

were at least 1.25° apart, and the orientation of each square was randomized on each trial. Squares at the edge of the window could be partly occluded.

In no-cue trials, the array was present for 565 ms, after which the screen turned entirely white for 47 ms. The array then reappeared with a new square in one of 149 positions (indicated by yellow circles in Fig. 1a); these positions represented a subset of the overall grid that fit within an imaginary 25° circle centered at fixation. The new square always appeared in a random position that had been unoccupied by a red or green square on previous frames within that trial. In cued trials, after 506 ms of the static array, 12.5 on the plus or minus x - or y -axis. The small outward apparent motion induced by this offset enhanced the salience of this cue. Cuing was followed by a return to the static array for 12, 82, 153, or 447 ms, and then a full-screen white blank that lasted 47 ms. After the blank, the array returned with a new square, as on no-cue trials.

The observers' task was to maintain fixation and report whether the new element was red or green. Four temporal intervals between the onsets of the cue and new square were tested. Each trial could have either no cue or a cue at one of four positions and one of four temporal intervals, selected at random. This design resulted in a total of 25,330 test trials per observer: $[1(\text{no cue}) + 4(\text{cue positions}) \times 4(\text{temporal intervals})] \times 149(\text{test positions}) \times 10(\text{trials per position})$. The training phase included half that number of trials. Each of 10 blocks of 2533 test trials was broken down into eight sessions of 316 or 321 trials. The intertrial interval was approximately 3 s in order to minimize possible effects of afterimages. A session typically lasted 35–45 min. Data were stored and later sorted and analyzed off-line.

5. Procedure

Observers were instructed to attend to the entire circular array of red and green squares and report the color of the new square in the final frame of each trial by pressing the appropriate button. There was no feedback on the correctness of responses. The instructions emphasized that the probability that this new square would appear at the cued location was the same as the probability that it would appear at any other tested location, and that therefore there was no advantage to attending to or ignoring the cues or their locations.

6. Practice phase

The blank (white) fourth frame (see Fig. 1) was 12 ms in duration during the 12,665 practice trials for the three naive observers. This was the shortest duration permitted by the refresh rate of the monitor, and was chosen to make learning the task as easy as possible. All observers were initially unable to see the new element if it appeared further than approximately 3° away from fixation. With practice, however, two of the observers learned to see the new square even when it was as far away as 12.5° from. The third observer did not learn to “ignore” the global blank, and continued to perform at chance beyond approximately 3° from fixation. He was disqualified because he could not perform the task over the whole area of the visual field tested, and because it became apparent that he was not paying full attention to the task, at times performing at chance even near fixation. Instead, the author served as the 3rd observer in the test phase of the experiment. He had received several thousand practice trials in the course of coding the experiment, but these practice data were not saved.

A blank of 47 ms was chosen for the test phase of the experiment because trained observers were neither perfect nor at chance at detecting the new square at this duration (observer 1: 72.8% of 25,330 trials correct; observer 2: 67.2% correct; observer 3: 68.4% correct).

7. Results

Fig. 2 shows the pooled data for the three observers from the test phase of the experiment. Change detection was significantly above chance for almost the entire 25° diameter circular area of the visual field tested, presumably because observers were instructed to attend to the whole array of squares. This corroborates the finding that outside the main focus of attention observers maintain a minimum degree of visual awareness with which they can detect and discriminate pop out classes of information at little or no cost to central attentional processing (Braun & Julesz, 1998). The attentional hot spot, corresponding to regions of near-perfect performance (yellow and white regions in Fig. 2), tended to change shape depending on the position of the cue, even though observers knew that the position of the cue was irrelevant to their task.

The no-cue case (Fig. 2, upper left) indicates the distribution of attention in the absence of a cue. This horizontal, elliptical hot spot can be thought of as the default shape of attention for this task. After the cue flashed, the hot spot elongated significantly (for statistics, see Tse et al., 2003) along the cue-fixation axis away from its default shape for most conditions. This is consistent with the finding that an abrupt onset causes an automatic allocation of attention to the location of that onset (Jonides & Yantis, 1988; Yantis & Egeth, 1999), particularly when attention has been set (Everling & Munoz, 2000; Folk, Remington, & Johnston, 1992) to detect changes, as here.

When the cue was to the left or right of fixation, the hot spot of attention tended to elongate along the horizontal axis, whereas when the cue was above or below fixation, the hot spot tended to elongate along the vertical axis. Also, for all observers, attention tended to pool not only at the cued location, but also at the corresponding location in the opposite hemifield. The hot spot spread out in the diametrically opposite direction from the flashed location along the cue-fixation axis, facilitating the detection of new elements up to 19° away from the cued location, relative to performance in the no-cue condition. The bilaterally elongated shape of the hot spot was most pronounced at 176 and 247 ms after cue onset, and diminished after that as the effects of cuing diminished.

8. Discussion of results

For most cue locations and stimulus onset asynchronies (SOAs), attention was enhanced (relative to the no-cue case) not only at the cued location, but also opposite the cued location. Opposite pooling of attention was stronger along the horizontal than the vertical axis, but this may be because the hot spot of attention was already biased to extend along the horizontal axis in the no-cue condition. Another asymmetry between the horizontal and vertical axes is that accuracy at the cued location was lower for cues on the vertical than on the horizontal axis. Why this should be is unclear, but it may be that the cue can interfere with the detection of changes at its location more strongly at upper vertical-axis locations, where attention has relatively low resolution (He, Cavanagh, & Intriligator, 1996; Intriligator & Cavanagh, 2001).

The data are consistent with previous results indicating that attention grows in an analog manner away from fixation toward a peripherally cued location (Shulman, Remington, & McLean, 1979; Tsai, 1983). However, the data do not rule out the possibility that there may also be non-analog shifts of attention under some conditions (Remington & Pierce, 1984; Sagi & Julesz, 1985; Sperling & Weichselgartner, 1995). Because the hot spot is aligned along the cue-fixation axis, it is possible that attention deforms in an analog manner along this axis, but need not traverse intermediate points between points that do not lie on this axis; this would be a compromise between temporally analog (Eriksen & St. James, 1986; Posner et al., 1980) and

quantal (Remington & Pierce, 1984; Sagi & Julesz, 1985; Sperling & Weichselgartner, 1995) theories of attentional shifts.

It may be that the “unnatural” maintenance of fixation played a role in the tendency for attention to pool opposite the cued location. Future work should reveal whether opposite pooling occurs even when observers are permitted to saccade to the cue. These and other experiments will have to determine why attention pools opposite the cued location, and what role cuing and fixation play in triggering this pooling. Opposite pooling of attention should also be observable using other measures of attentional allocation and performance, such as detection thresholds.

9. General discussion of method

This method exploits the fact that change detection is hampered in a flicker-induced change-blindness paradigm to map how spatial attention changes in response to an irrelevant peripheral cue. The basis of this method is the assumption that the likelihood of detecting a change at a given location corresponds to the occurrence of attention at that location. On theoretical and experimental grounds (Rensink, 2000b; Rensink et al., 1997; Simons, 1996), it can be argued that the distribution of judgment accuracy correlates directly with the distribution of attention. In essence, change-detection accuracy affords an indirect measure of attentional occurrence because change-detection accuracy must increase and decrease with increases and decreases in the perceptual signal-to-noise ratio. If spatial attention is operationally defined as the mechanism that boosts this ratio (e.g., Lee, Itti, & Braun, 1999; Treue & Maunsell, 1999), then the present method measures a correlate of attentional occurrence at each location. Change detection accuracy can thus be added to other indirect measures of attention, such as reaction time (RT) and sensitivity.

9.1. The meaning of the hotspot

Traditional metaphors that refer to attention as a moving spotlight or window fail to capture several properties of attention revealed by mapping the accuracy of change detection in a change-blindness paradigm. The high-resolution maps described here reveal that spatial attention can be graded to cover virtually the entire visual field. Attention has a hot spot that deforms in response to a peripherally flashed cue, even when observers know that this cue is uninformative for the task at hand. This result is consistent with the finding that abrupt onsets invoke an automatic and mandatory capture of attention (Jonides & Yantis, 1988; Yantis & Egeth, 1999). The combination of a global-but-low-resolution distribution of attention with a local-but-high-resolution distribution of attention is consistent with models of attention (Treisman & Gelade, 1980; Treisman & Sato, 1990; Wolfe & Gancarz, 1996) that allow for pop out of targets that are highly distinguishable from distractors, but require serial search for targets that are not easily distinguishable without the high-resolution hot spot.

Unlike a moving spotlight, the hot spot of attention tends to remain centered at fixation (Wolfe, O'Neill, & Bennett, 1998). This is not surprising because observers were required to maintain fixation and because acuity is best in the fovea. The hot spot also tends to be a single connected blob, as opposed to several disconnected blobs, at least when there is only one peripheral cue, as in the present study. This may account for results showing that attention is more easily allocated to compact regions than to multiple smaller regions of equivalent area (Podgorny & Shepard, 1983). Future experiments using two or more simultaneous cues can be carried out to determine how attention is split between two or more cues (cf. Bichot, Cave, & Pashler, 1999).

There is an averaging of attention inherent in the process of subsampling target locations. Assume for the moment that the analog distribution of attention differs on each trial. On average it will be stronger at some locations than at others. The attentional maps shown here do not

reveal the distribution of attention on any given trial because on each trial only a single position is probed. They reveal the average distribution of attention across all trials for a given condition. More precisely, they reveal a probability density function that attention was allocated to a given location for a given condition. Because data were averaged across trials, they cannot distinguish between two alternative accounts of how attention is spatially allocated: Either it is (a) allocated to each location on each trial, albeit to varying degrees, or it is (b) allocated to a few objects or even one object (Rensink, 2000a) at a time. If the former is the case, changes were detected because attention was present at all locations; different detection rates reflect different degrees of attending to different locations. The maps in Fig. 2 would then correspond to average attentional strength at each location on each trial. If the latter is the case, changes may be represented implicitly (Thornton & Fernandez-Duque, 2000; but see Mitroff, Simons, & Franconeri, 2002), and these implicit representations may be what then attracts attention (Smilek, Eastwood, & Merikle, 2000). The maps in Fig. 2 would then indicate the likelihood of the all-or-none presence of attention on each trial. The first account is more likely because the hot spot was typically 15–20° across. In this hot spot, detection rates were more or less uniformly 100%. This is evidence that attention was globally distributed on each trial at least for this large hot spot. However, even outside the hot spot, detection rates were typically better than 60% over a wide area. If attention were allocated only in a spatially piecemeal or quantal fashion on any given trial, it would still have to be distributed over an average of 60% of the squares on each trial to account for this result. Rensink's (2000a) notion that attention can only be allocated to one or at most very few objects at a time seems unlikely to be the case. There is, however, another possibility that is consistent with both Rensink's data and the present data. It could be that the seemingly incompatible results arise from measuring different aspects of attention. For example, there may be operators (e.g., FINSTs; see Pylyshyn & Storm, 1988) that bind features into object representations, and these operators may have a limited spatial range over which they can create object representation. It could be that Rensink's data reflect the inability of these operators to operate on more than a small number of objects at a time. In contrast, the present data could reflect the changing range of these operators.

9.2. Accuracy versus reaction time

Perhaps the most refined use of reaction times to measure the spatial distribution of attentional allocation to date has been the study of Bennett and Pratt (2001). They used reaction times to map out the spatial distribution of attention under cuing conditions that generate inhibition of return (IOR). In IOR, RTs are faster at a cued location up to approximately 150 ms after cue onset, followed by decrements in processing 300 ms or more after cue onset, although increasing task difficulty can move this crossover point to beyond 600 ms (Klein, 2000). This is generally interpreted to mean that cued locations are inhibited from again receiving full attention starting approximately 300 ms after having been attended (Klein, 2000; Posner & Cohen, 1984). In addition, recent evidence has shown that IOR involves a relative speeding of RTs opposite an attended location (Bennett & Pratt, 2001; Pratt, Spalek, & Bradshaw, 1999). IOR is thought to enhance the efficiency of visual search by inhibiting attentional shifts to already searched locations and enhancing the likelihood of shifts to new locations. According to the model proposed by Bennett and Pratt, the distribution of reaction times that they observe cannot be used to infer anything about the size of the attentional spotlight. Instead, reaction times, according to their model, reflect changes in the time to attend to a particular location. In contrast, the detection accuracy data here is assumed to reflect the strength of attention at a given location at a given instant. As such the data described here can be used to make inferences about the shape and size of the attentional spotlight, whereas the Bennett and Pratt data cannot. This suggests that detection accuracy offers a superior means of mapping attentional allocation than reaction time, at least for this change blindness paradigm.

Although the finding that performance is enhanced opposite a cue may seem similar to IOR, it is unlikely that IOR underlies the strengthening of attention opposite the cue, at least as measured within the present change blindness paradigm (compare Snyder, Schmidt, & Kingstone, 2001). The opposite enhancement in change detection accuracy was apparent already at the shortest SOA tested (106 ms), which is faster than the latency typically reported for IOR. Reaction times were collected (for data see Tse et al., 2003) but revealed no clear relationships between the distribution of reaction times and the location of the cue. In particular, there was no statistically significant evidence of attentional pooling (i.e., faster reaction times) at or opposite the cue location in the RT data. Thus, the present results about attention could not have been obtained by using the standard measure of attention: reaction time. It could be that the maximum tested duration (541 ms) was not long enough to demonstrate evidence of IOR in the reaction time (RT) data or evidence of attentional fading at the cued location. Another possible explanation of the lack of IOR in the RT data is that the intervening global flash effectively wiped out evidence of IOR in the reaction time data. Future experiments within an IOR paradigm (cue, delay, target) should be able to determine whether RTs opposite the cued location speed up by approximately 100 ms, as predicted by the data of this experiment. The visual system may orient to the cued and opposite locations along the cue-fixation axis in such a way that attention fades more rapidly at the cued location than at the opposite location. Note, however, that there was no evidence of a more rapid decline in accuracy at the cued than the opposite location.

Detection accuracy is quite likely a better measure of the degree to which attentional resources have been allocated to a location because attention may just be the mechanism that modulates gain, and therefore sensitivity in neurons. Such modulation would affect detection thresholds directly, but would only affect reaction times indirectly, as a consequence of having modified detection thresholds. There are several additional sources of intervening noise between detection and motor response that could potentially enter reaction time data, effectively “blurring” measurements of the spatiotemporal distribution of attention. If it is true that detection accuracy affords a less noisy means of probing attention than reaction times, then it is problematic that the majority of claims about the distribution and timecourse of attention have been made on the basis of reaction time measures.

9.3. Drawbacks of the method

Observers must receive extensive training in the present paradigm before they are able to see changes away from fixation despite the blinding blank (see Fig. 2). One concern is that highly trained observers may allocate attention differently than untrained observers. This is unlikely because the basic pattern of hot-spot elongation along the cue-fixation axis (with attention pooling at and opposite the cued location) was already apparent in the training phase for the two naive observers who continued on to the test phase. Training is known to improve performance on attentionally demanding tasks. Training may function to lessen the difficulty of a task, by increasing its familiarity (Braun, 1998), reducing the amount of attention required to carry out the task (Norman & Bobrow, 1975), or increasing the efficiency or strength of attentional allocation (Cave & Zimmerman, 1997). But there is no evidence that the spatial distribution of attention itself changes with practice in a cuing task. Future versions of this paradigm should involve less training, be less time-consuming, and more closely resemble real-world situations in which observers naturally allocate and shift attention.

The present experiments did not distinguish conscious change detection from mere unconscious but accurate guessing, because observers did not have to localize the change. One way to address this problem would be to have observers indicate the location of a change. For example, observers could be required to saccade to the location of a change. In this case, a correct response would not only have to be the correct color, but also the correct visual quadrant.

This would move baseline or chance performance down from 50% (chance performance on red versus green), as it is in the present experiment down to 12.5% (chance performance on red versus green in four possible quadrants). This lower baseline would allow more precise specification of the hotspot of attention because the likelihood of making a correct decision decreases in proportion to the baseline. The specification of change location could be accomplished using saccades, button presses, a joystick, pointing, or even verbal statements. Even though observers would be explicitly indicating the location of a change, the representation of change location could in principle still be non-conscious. However, this is inconsistent with what observers report. Observers report either seeing the change clearly or not. When they phenomenally experience the change they can specify its location, and when they do not experience it, they cannot report its location. The issue of conscious versus unconscious information processing of change location can be addressed head on by having observers make a confidence rating that they in fact saw the target on each trial. Observers will presumably be more confident about those trials where they consciously experienced the location of the change. Plots of attentional allocation would include only trials whose color and location was correctly specified; In addition, corresponding plots could be made for high confidence versus low confidence correct responses. If the spatiotemporal distribution of accurate responses differs as a function of confidence, we may be in a position to address the relationship of attention to conscious experience.

9.4. Future directions

This method has the capacity to resolve some long-standing disputes in the study of attention. Most promising is the potential to establish the size and shape of the hotspot of attention as a function of time. Although Posner (1978) found that the size of the attentional window did not vary as a function of eccentricity, others have argued that this size increases with eccentricity. Estimates of the size of the attentional focus vary widely even for authors who all used variants of Posner's original cuing paradigm (Nakayama & Mackeben, 1989, experiment 5, 8 arcmin; Posner et al., 1980, 1°; Sagi & Julesz, 1985, 1.5° at 2° eccentricity; Steinman, Steinman, & Lehmkuhle, 1995, 10–12°; Hughes & Zimba, 1985, whole visual hemifields). One problem with traditional cuing paradigms is that observers have nothing to attend to after a cue has disappeared. If attention is allocated to objects rather than space, then it is problematic to attempt to measure the shape (Sagi & Julesz, 1985) or size of the attentional window in the absence of any object to which one can attend. The present method gets around this problem by having a target that appears and remains present at its onset location. If attention requires an object to adhere to, then attention should be able to adhere to this target. These findings support the conclusion that the size of the attentional window is substantially larger than the size of the target. This is consistent with the finding that attention has a much coarser spatial resolution than visual resolution (He et al., 1996; Intriligator & Cavanagh, 2001). Moreover, these results are consistent with work on visual crowding (He et al., 1996) that implies that the shape of attention is like a spoke emanating from the point of fixation. These data would suggest that at least for the case of exogenous attention under conditions of fixation, the spoke emanates away from fixation both towards and away from the cued location.

It may be that observers learned to ignore or discount the blinding blank. Some observers were able to transfer this ability to other change-detection tasks that involved a full-field blank. For example, in transformational apparent motion a square becomes a rectangle (Tse & Logothetis, 2002). The percept is of a smooth animated transition between these two (in fact) discrete states. If an array of mixed squares and rectangles is alternated with an array identical in all regards except that one element has changed shape, the changed element will perceptually pop out at the moment of the change. But if a global flash is added in between the frames, pop out of this element is eliminated. Novices drop down to slopes (plotting reaction time against the number of distractors) comparable to those found in serial search. Remarkably, the percept of

transformational apparent motion returns once an observer has attended to the location of the change, whether they are fixating that location or not. However, preliminary data suggests that practiced observers appear to start at lower slopes. How observers learn to do this is not clear. Perhaps they learn to hold information in iconic memory (Sperling, 1960) for longer than untrained observers, or learn to transfer the contents of the iconic buffer to another protected buffer, or learn to stop the blank frame from fully flushing out that iconic buffer. Another possibility is that they learned to allocate attention in a way they could discount the flash by, for example, turning down the gain on motion-energy detectors. However this is accomplished, it is an interesting phenomenon in its own right. It is an ability that appears to be learned and one which appears to differ from observer to observer. This would explain the inability of one of the observers in the present experiment to learn to discount the global flash. As long as the learning that appears to take place does not involve changes in the spatial allocation of attention, results obtained with this method should also be informative about how untrained observers allocate attention in everyday circumstances.

Change blindness need not be induced by a global flash. A basic strategy for circumventing or incapacitating the motion-energy detectors that are the visual system's first line of defense in detecting changes is to swamp the motion-energy detector array with changes, so that no single detected change will pop out as an isolated peak of activation on the saliency map that drives the allocation of attention in a stimulus-driven manner. A global flash may simply be the best way to swamp the motion-energy detection array, because a global high-luminance flash should effectively activate all motion-energy detectors. However, multiple discrete changes are sufficient to incapacitate change localization (O'Regan et al., 1999), even though such "mudsplashes" are not nearly as effective at inducing change blindness as a full-field blank. The method described here should presumably work with even a few discrete changes rather than a global flash. This might be an improvement in the method, because then no time will be lost to the duration of a global flash. This could permit attentional mapping even when the lag between some attention-grabbing event and the change to be detected is very short. On the other hand, if change detection rates are much higher when using mudsplashes instead of a full-field blank, it might be difficult to obtain useful maps of attention, because the hotspot of attention can only be specified relative to other areas of the visual field that have lower accuracy rates. If the whole visual field were close to ceiling performance, the method would not be useful. For this reason mudsplashes may be less than ideal. One way to lower the baseline accuracy rate would be to increase the number of locations at which mudsplashes occur. In the extreme, every location could sustain a change at the moment that a new element appeared. In the context of the present experiment, this would mean that the new red or green square appeared at the moment that every "distractor" red or green square either rotated or changed position slightly. So many local changes dispersed across the test area might successfully mask the signal caused by the onset of a new element nearly as well as a global flash, without the temporal cost associated with having a global flash.

A second strategy for circumventing stimulus-driven change detection is to allow changes to occur so gradually that they do not trigger motion-energy detectors (Auvray & O'Regan, 2003; Simons et al., 2000) or a concomitant specification of location on a salience map. Because attention is spatially redistributed quickly in response to a salient event, it is unlikely that gradual changes can be exploited to map attention in a change blindness paradigm. However, the slow-change method of inducing change blindness could help answer another interesting question, namely: What is the duration over which attention can permit comparisons between two successive states of an object? Changes are probably hard to detect in the slow-change paradigm for two reasons. First, changes occur below the threshold of motion-energy detectors. Second, changes are so slow that they occur below the threshold specified by the maximum duration over which attention can permit a comparison of successive object states. If the dwell time of attention is too short, successive states may not differ sufficiently to count as different.

Even when a change is happening across the majority of an image, as when a scene gradually darkens in the late afternoon, no change will be noticed, because each attention-based comparison will yield the answer that there is no difference. These two thresholds (one for motion-energy detection and the other for attention-based change detection) that specify the amount of change that can be detected automatically need not be the same. If the rate of change could be sped up such that it still remained subthreshold for motion-energy detection, yet “popped out” upon first allocation of attention, then it could be shown that these thresholds differ. Such an attentional threshold for change detection would in turn specify the temporal span over which comparisons are made and can be made.

Observers sometimes noticed apparent motion between the cues and the target. If both the cue and target are attended, then they can be bound together as a single object and generate an apparent motion signal. In addition to first-order (luminance contrast-based) and second-order (pattern- or form-based) motion, it has been suggested that attentional tracking provides a third type of signal to high-level motion sensitive units (Cavanagh, 1992; Lu & Sperling, 1995). Because first- or second-order motion signal cannot underlie apparent motion in this change blindness paradigm, the apparent motion experienced by some observers must be attention-based motion. The method described here can perhaps be enhanced to exploit this tendency to see apparent motion. Attention-based apparent motion would provide another useful probe of attentional allocation and strength, because the perceived strength of apparent motion presumably correlates with the strength of attention allocated to the two locations of the objects that undergo apparent motion.

Unlike most natural scenes, the display shown in Fig. 1 was unstructured. Therefore, there may have been no recourse but to allocate attention in a purely spatial fashion. Several researchers have, however, provided evidence that attention tends to be allocated more to objects than to locations (e.g., Davis, Russell, Turatto, & Freeman, 2001; Duncan, 1984; Vecera & Farah, 1994). Using the change blindness method with objects may reveal a hot spot that adheres to a surface or a contour rather than being uniformly distributed across a large area. While attention was initially thought to be location-based (Broadbent, 1982; Downing & Pinker, 1985; Eriksen & St. James, 1986; Posner, 1980; Tsal, 1983; Tsal & Lavie, 1988), more recent studies support the view that attention operates in an object-centered reference frame rather than a location-centered one. Driver and Baylis (1989) showed that attention can be allocated to stimulus features other than location, including motion, color, and form. Duncan (1984) showed that attention can be bound to objects or their surfaces rather than their locations. Several behavioral (Baylis & Driver, 1992; Duncan, 1984; Kahneman et al., 1992; Kramer & Jacobson, 1991; Lavie & Driver, 1996) and neuropsychological (Behrmann & Moscovitch, 1994; Driver & Halligan, 1991; Humphreys & Riddoch, 1993) studies have since been published that support the view that attention is object-based rather than space-based. For example, attention appears to speed the processing of objects that are on the same surface relative to objects on different surfaces, even when the objects on different surfaces are closer to one another on the retina or in space (He & Nakayama, 1994; Nakayama, He, & Shimojo, 1995). Similar experiments can be designed using the present method. A change can occur on an attended surface or an unattended surface. Theories of object-centered attention would predict that the hotspot of attention should encompass the attended surface more so than the nonattended surface. Location-based and object-based accounts of attention can be pitted against one another here by having changes occur at identical locations, but on an attended or non-attended surface. Location-based theories would predict no difference in change detection accuracy rates, whereas object-based theories would predict a difference.

It should be possible to employ this method to map how attention is distributed when one is tracking, say, a moving disk with sustained or endogenous attention. This is an experiment currently being carried out in the author's laboratory. Observers must maintain fixation and

then track a white outline of a disk which moves along a circular trajectory around the fixation point. To guarantee that observers are in fact attentively tracking this disk, they are required to press a button each time there is a slight change in the shape or color of its outline. The disk moves over an array of red and green squares as shown in Fig. 1. Again observers must specify the color of the new element that appears during the global flash. In addition they must specify its location. The outcome of this experiment should clarify what exactly happens when observers attend to and track an object. It may turn out that the hotspot of attention is enhanced ahead of or behind a moving target, or along a spoke leading away from the point of fixation. No matter what the outcome of this experiment is, it will provide strong constraints on possible cognitive and neuronal models of attentional processing.

References

- Auvray M, O'Regan JK. L'influence des facteurs sémantiques sur la cécité aux changements progressifs dans les scènes visuelles. *Année Psychologique* 2003;103:9–32.
- Baylis G, Driver J. Visual parsing and response competition: The effect of grouping factors. *Perception & Psychophysics* 1992;51:145–162. [PubMed: 1549433]
- Becklen R, Cervone D. Selective looking and the noticing of unexpected events. *Memory & Cognition* 1983;11(6):601–608.
- Behrmann M, Moscovitch M. Object-centered neglect in patients with unilateral neglect: Effects of left-right coordinates of objects. *Journal of Cognitive Neuroscience* 1994;6:1–16.
- Bennett PJ, Pratt J. The spatial distribution of the inhibition of return. *Psychological Science* 2001;12:76–80. [PubMed: 11294232]
- Bichot NP, Cave KR, Pashler H. Visual selection mediated by location: Feature-based selection of noncontiguous locations. *Perception & Psychophysics* 1999;61:403–423. [PubMed: 10334090]
- Braun J. Vision and attention: The role of training. *Nature* 1998;393:424–425. [PubMed: 9623997]
- Braun J, Julesz B. Withdrawing attention at little or no cost: Detection and discrimination. *Perception & Psychophysics* 1998;60:1–23. [PubMed: 9503909]
- Broadbent D. Task combination and selective intake of information. *Acta Psychologica* 1982;50:253–290. [PubMed: 7124432]
- Cavanagh P. Attention-based motion perception. *Science* 1992;257(5076):1563–1565. [PubMed: 1523411]
- Cave KR, Zimmerman JM. Flexibility in spatial attention before and after practice. *Psychological Science* 1997;8:399–403.
- Davis G, Russell C, Turatto M, Freeman E. Segmentation, attention and phenomenal visual objects. *Cognition* 2001;80(12):61–95. [PubMed: 11245840]
- Downing, CJ.; Pinker, S. The spatial structure of visual attention. In: Posner, MI.; Marin, OSM., editors. *Attention and performance XI*. London: Erlbaum; 1985. p. 171-187.
- Driver J, Baylis GC. Movement and visual attention: The spotlight metaphor breaks down. *Journal of Experimental Psychology: Human Perception and Performance* 1989;15:448–456. [PubMed: 2527954]
- Driver J, Halligan P. Can visual neglect operate in object-centered coordinates? *Cognitive Neuropsychology* 1991;8:475–496.
- Duncan J. Selective attention and the organization of visual information. *Journal of Experimental Psychology: General* 1984;113:501–517. [PubMed: 6240521]
- Eriksen CW, St. James JD. Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics* 1986;40:225–240. [PubMed: 3786090]
- Everling S, Munoz DP. Neuronal correlates for preparatory set associated with pro-saccades and anti-saccades in the primate frontal eye field. *Journal of Neuroscience* 2000;20(1):387–400. [PubMed: 10627615]
- Folk CL, Remington RW, Johnston JC. Involuntary covert orienting is contingent upon attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance* 1992;18:1030–1044. [PubMed: 1431742]

- Haines, RF. A breakdown in simultaneous information processing. In: Obrecht, G.; Stark, LW., editors. *Presbyopia research: From molecular biology to visual adaptation*. New York: Plenum Press; 1989. p. 171-175.
- He S, Cavanagh P, Intriligator J. Attentional resolution and the locus of visual awareness. *Nature* 1996;383:334–337. [PubMed: 8848045]
- He ZJ, Nakayama K. Perceived surface shape not features determines correspondence strength in apparent motion. *Vision Research* 1994;34:2125–2135. [PubMed: 7941410]
- Horowitz TS, Wolfe JM. Visual search has no memory. *Nature* 1998;6(3946693):575–577. [PubMed: 9707117]
- Hughes HC, Zimba DL. Spatial maps of directed visual attention. *Journal of Experimental Psychology: Human Perception and Performance* 1985;11(4):409–430. [PubMed: 3161984]
- Humphreys, G.; Riddoch, M. Interactions between object and space vision revealed through neuropsychology. In: Meyer, D.; Kornblum, S., editors. *Attention and performance XIV*. Hillsdale, NJ: Erlbaum; 1993. p. 139-168.
- Intriligator J, Cavanagh P. The spatial resolution of visual attention. *Cognitive Psychology* 2001;43:171–216. [PubMed: 11689021]
- Jonides J, Yantis S. Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics* 1988;43:346–354. [PubMed: 3362663]
- Kahneman D, Treisman A, Gibbs BJ. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 1992;24:175–219. [PubMed: 1582172]
- Klein RM. Inhibition of return. *Trends in Cognitive Science* 2000;4(4):138–146.
- Kramer AF, Jacobson A. Perceptual organization and focused attention: The role of objects and proximity in visual processing. *Perception & Psychophysics* 1991;50:267–284. [PubMed: 1754368]
- Lavie N, Driver J. On the spatial extent of attention in object-based visual selection. *Perception & Psychophysics* 1996;58:1238–1251. [PubMed: 8961834]
- Lee DK, Itti L, Koch C, Braun J. Attention activates winner-take-all competition among visual filters. *Nature Neuroscience* 1999;2:375–381.
- Lu ZL, Sperling G. Attention-generated apparent motion. *Nature* 1995;377(6546):237–239. [PubMed: 7675109]
- Mack, A.; Rock, I. *Inattention blindness*. Cambridge, MA: MIT Press; 1998.
- McConkie GW, Currie CB. Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception & Performance* 1996;22(3):563–581. [PubMed: 8666953]
- Mitroff SR, Simons DJ, Franconeri SL. The siren song of implicit change detection. *Journal of Experimental Psychology: Human Perception and Performance* 2002;28(4):798–815. [PubMed: 12190251]
- Nakayama, K.; He, ZJ.; Shimojo, S. Visual surface representation: A critical link between lower-level and higher-level vision. In: Kosslyn, SM., editor. *An invitation to cognitive science*. Cambridge, MA: MIT Press; 1995. p. 1-70.
- Nakayama K, Mackeben M. Sustained and transient components of focal visual attention. *Vision Research* 1989;29:1631–1647. [PubMed: 2635486]
- Neisser U, Becklen R. Selective looking: Attending to visually specified events. *Cognitive Psychology* 1975;7(4):480–494.
- Norman DA, Bobrow DG. On data-limited and resource-limited processes. *Cognitive Psychology* 1975;7:44–64.
- O'Regan JK, Rensink RA, Clark JJ. Change blindness as a result of 'Mudsplashes'. *Nature* 1999;398:34. [PubMed: 10078528]
- Pashler HE. Familiarity and visual change detection. *Perception & Psychophysics* 1988;44:369–378. [PubMed: 3226885]
- Podgorny P, Shepard RN. Distribution of visual attention over space. *Journal of Experimental Psychology: Human Perception and Performance* 1983;9:380–393. [PubMed: 6223978]
- Posner, MI. *Chronometric explorations of the mind*. Hillsdale, NJ: Erlbaum; 1978.

- Posner MI. Orienting of attention. *Quarterly Journal of Experimental Psychology* 1980;32:3–25. [PubMed: 7367577]
- Posner, M.; Cohen, Y. Components of visual orienting. In: Bouma, H.; Bouwhuis, D., editors. *Attention and performance X*. London: Erlbaum; 1984. p. 531–556.
- Posner MI, Snyder CRR, Davidson BJ. Attention and the detection of signals. *Journal of Experimental Psychology: General* 1980;109:160–174.
- Pratt J, Spalek T, Bradshaw E. The time to detect targets at inhibited and non-inhibited locations: Preliminary evidence for attentional momentum. *Journal of Experimental Psychology: Human Perception and Performance* 1999;25:730–746.
- Pylyshyn ZW, Storm RW. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision* 1988;3:179–197. [PubMed: 3153671]
- Remington R, Pierce L. Moving attention: Evidence for time-invariant shifts of visual selective attention. *Perception & Psychophysics* 1984;35:393–399. [PubMed: 6739275]
- Rensink RA. Seeing, sensing, and scrutinizing. *Vision Research* 2000a;40:1469–1487. [PubMed: 10788653]
- Rensink RA. Visual search for change: A probe into the nature of attentional processing. *Visual Cognition* 2000b;7:345–376.
- Rensink RA, O'Regan JK, Clark JJ. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science* 1997;8:368–373.
- Sagi D, Julesz B. Fast noninertial shifts of attention. *Spatial Vision* 1985;1:141–149. [PubMed: 3940055]
- Shulman GL, Remington RW, McLean JP. Moving attention through visual space. *Journal of Experimental Psychology: Human Perception and Performance* 1979;5:522–526. [PubMed: 528957]
- Simons DJ. In sight, out of mind: When object representations fail. *Psychological Science* 1996;7:30–35.
- Simons DJ, Chabris CF. Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception* 1999;28(9):1059–1074. [PubMed: 10694957]
- Simons DL, Franconeri SL, Reimer RL. Change blindness in the absence of a visual disruption. *Perception* 2000;29(10):1143–1154. [PubMed: 11220207]
- Simons DJ, Levin DT. Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review* 1998;5:644–649.
- Smilek D, Eastwood JD, Merikle PM. Does unattended information facilitate change detection? *Journal of Experimental Psychology: Human Perception and Performance* 2000;26:480–487. [PubMed: 10811158]
- Snyder J, Schmidt WC, Kingstone A. Attentional momentum does not underlie inhibition of return. *Journal of Experimental Psychology: Human Perception and Performance* 2001;27:1420–1432. [PubMed: 11766934]
- Sperling G. The information available in brief visual presentations. *Psychological Monographs* 1960;74:1–29.
- Sperling G, Weichselgartner E. Episodic theory of the dynamics of spatial attention. *Psychological Review* 1995;102:503–552.
- Steinman BA, Steinman SB, Lehmkuhle S. Visual attention mechanisms show a center-surround organization. *Vision Research* 1995;35:1859–1869. [PubMed: 7660592]
- Thornton IM, Fernandez-Duque D. An implicit measure of undetected change. *Spatial Vision* 2000;14:21–44. [PubMed: 11334179]
- Tipper SP. Does negative priming reflect inhibitory mechanisms? A review and integration of conflicting views. *Quarterly Journal of Experimental Psychology A* 2001;54(2):321–343.
- Treisman A. Feature binding, attention and object perception. *Philosophical transactions of the royal society of London B biological sciences* 1998;29(3531373):1295–1306.
- Treisman A. Solutions to the binding problem: progress through controversy and convergence. *Neuron* 1999;24(1):105–110. 111–125. [PubMed: 10677031]
- Treisman A, Gelade G. A feature-integration theory of attention. *Cognitive Psychology* 1980;12:97–136. [PubMed: 7351125]

- Treisman A, Sato S. Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance* 1990;16:459–478. [PubMed: 2144564]
- Treue S, Maunsell JH. Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *Journal of Neuroscience* 1999;19:7591–7602. [PubMed: 10460265]
- Tsal Y. Movements of attention across the visual field. *Journal of Experimental Psychology: Human Perception and Performance* 1983;9:523–530. [PubMed: 6224890]
- Tsal Y, Lavie N. Attending to color and shape: The specific role of location in selective visual processing. *Perception & Psychophysics* 1988;44:15–21. [PubMed: 3405725]
- Tse PU, Sheinberg DL, Logothetis NK. The duration of 3-d form analysis in transformational apparent motion. *Perception*. *Perception & Psychophysics* 2002a;64(2):244–265. [PubMed: 12013379]
- Tse PU, Sheinberg DL, Logothetis NK. Fixational eye movements are not affected by abrupt onsets that capture attention. *Vision Research* 2002b;42:1663–1669. [PubMed: 12079794]
- Tse PU, Sheinberg DL, Logothetis NK. Attentional enhancement opposite a peripheral flash revealed by change blindness. *Psychological Science* 2003;14(2):1–8. [PubMed: 12564746]
- Vecera SP, Farah MJ. Does visual attention select objects or locations? *Journal of Experimental Psychology: General* 1994;123:146–160. [PubMed: 8014610]
- Wolfe, JM.; Gancarz, G. Guided search 3.0: A model of visual search catches up with Jay Enoch 40 years later. In: Lakshminarayanan, V., editor. *Basic and clinical applications of vision science*. Dordrecht, The Netherlands: Kluwer Academic Publishers; 1996. p. 189-192.
- Wolfe JM, O'Neill P, Bennett SC. Why are there eccentricity effects in visual search? Visual and attentional hypotheses. *Perception & Psychophysics* 1998;60:140–156. [PubMed: 9503918]
- Yantis S, Egeth HE. On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance* 1999;25:661–676. [PubMed: 10385983]

Fig. 1.

Experimental design. The trial sequence is illustrated in (b). On cued trials, a spot appeared in Frame 2, as shown, and Frame 3 was presented for one of three durations. On noncued trials, there was no spot in Frame 2, and Frame 3 always lasted 12 ms. In the final frame of the sequence (labeled 5 here) on both cued and noncued trials, a new square was shown in one of 149 positions of the array (a); these positions represented a subset of the overall grid that fit within an imaginary 25° circle centered at fixation. The new square always appeared in a random position that had been unoccupied on previous frames within a given trial. The final frame always contained one and only one new element. The observers' task was to maintain fixation and report whether the new element was red or green. SOA = stimulus onset asynchrony.

Fig. 2.

Average attentional maps for three observers. Pooled percentage correct for no-cue trials and cued trials with each SOA and each cued position (white stars). Each of the 149 tested locations shown in Fig. 1a could have a maximum of 30 correct responses. Bilinear interpolation was used to obtain values for nontested positions, and results were smoothed using a rotationally symmetric 3×3 Gaussian low-pass filter with a standard deviation of 0.5 grid units. Values were then mapped to the $25^\circ \times 25^\circ$ color maps shown. The probability that observed values were obtained by chance is indicated by the p values to the left of the bar. These values were obtained using the binomial test, where $n = 30$ and $p = .5$. The test region was approximately an imaginary circle that touched the borders of these square maps. Black regions outside these circles were not tested. The fixation point is indicated at the center of each map.