replication and non-extension of Law, Pratt and Abrams (1995) *Percept. Psychophys.* 60, 1452–1456
59 Watson, D.G. and Humphreys, G.W. Visual marking: evidence for inhibition using a probe-dot detection paradigm. *Percept. Psychophys.* (in press)
60 Pratt, J. *et al.* (1999) The time to detect targets at inhibited and noninhibited locations: Preliminary evidence for attentional momentum. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 730–746
61 Rafal, R.D. *et al.* (1994) Effects of inhibition of return on voluntary and visually guided saccades. *Can. J. Exp. Psychol.* 48, 284–300

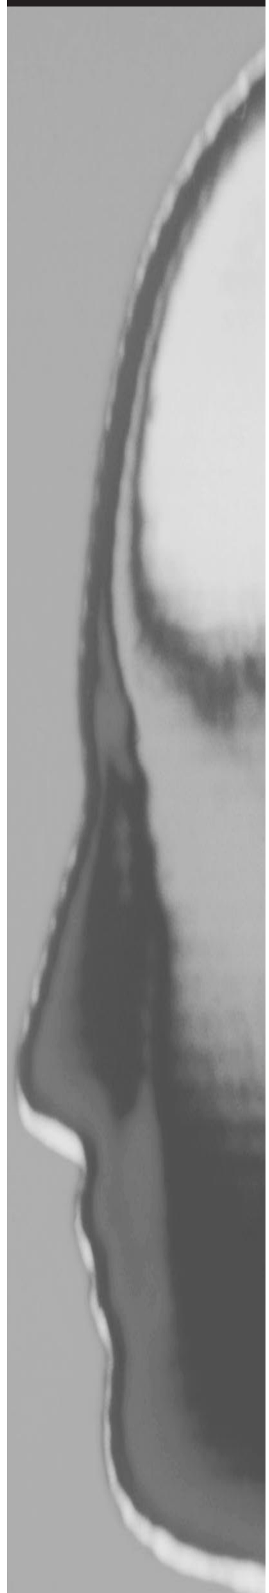

Attentional capture and inattentional blindness

Daniel J. Simons

Although we intuitively believe that salient or distinctive objects will capture our attention, surprisingly often they do not. For example, drivers may fail to notice another car when trying to turn or a person may fail to see a friend in a cinema when looking for an empty seat, even if the friend is waving. The study of attentional capture has focused primarily on measuring the effect of an irrelevant stimulus on task performance. In essence, these studies explore how well observers can ignore something they expect but know to be irrelevant. By contrast, the real-world examples above raise a different question: how likely are subjects to notice something salient and potentially relevant that they do not expect? Recently, several new paradigms exploring this question have found that, quite often, unexpected objects fail to capture attention, a phenomenon known as 'inattentional blindness'. This review considers evidence for the effects of irrelevant features both on performance ('implicit attentional capture') and on awareness ('explicit attentional capture'). Taken together, traditional studies of implicit attentional capture and recent studies of inattentional blindness provide a more complete understanding of the varieties of attentional capture, both in the laboratory and in the real world.

In order to study whether unattended stimuli can capture attention, researchers must first operationally define a way to measure the capture of attention. Two classes of definitions have been used in the study of attentional capture. 'Explicit attentional capture' occurs when a salient and unattended stimulus draws attention, leading to awareness of its presence. 'Implicit attentional capture' is revealed when a salient and irrelevant stimulus affects performance on another task, regardless of whether or not subjects are aware of the stimulus. The first, explicit attentional capture, is perhaps the more intuitive conception: when someone across a room says our name or waves at us vigorously, the stimulus signal segregates itself from the background and we become aware of its source. Typically, studies adopting the explicit approach determine whether capture has occurred by asking subjects whether they noticed the critical stimulus. Although reporting the presence of an unexpected object may not conclusively demonstrate attentional capture, the failure to notice it suggests that the object failed to capture attention explicitly. Several recent studies of explicit attentional capture have found that when observers are focused on some other object or event, they often fail to notice salient and distinctive objects, a phenomenon that is termed 'inattentional blindness'. Although the use of explicit reports was one of the first approaches used to study attentional capture (in the study of divided and selective attention), and despite a recent resurgence of interest in inattentional blindness, most recent studies have focused on implicit attentional capture. That is, such studies make the critical stimulus irrelevant to the primary task and infer capture based on different patterns of response times or eye movements. This review considers evidence for attentional capture in both implicit and explicit paradigms. Together, these findings raise the intriguing possibility that salient stimuli, including the appearance of new objects, might not always capture attention in the real world.

Implicit measures of attentional capture

Most recent studies of attentional capture have adapted methodologies used extensively in the study of visual search. Four distinct paradigms have been used to explore implicit

D.J. Simons is at the Department of Psychology, Harvard University, 33 Kirkland Street, #820, Cambridge, MA 02138, USA.

tel: +1 617 495 5387
fax: +1 617 495 3728
e-mail: dsimons@wjh.harvard.edu

1364-6613/00/$ – see front matter PII: S1364-6613(00)01455-8

Table 1. Paradigms used in studies of attentional capture

Paradigm	Method	Evidence for capture
Additional Singleton	Subjects perform a search task and an additional distinctive item appears in the display.	Slowed overall search performance in the presence of the additional item, often referred to as the cost of filtering out the additional irrelevant item.
Oculomotor Capture	Subjects perform a search task while eye movements are monitored. An additional distinctive item appears in the display.	An eye movement toward the additional item.
Irrelevant Feature Search	An irrelevant distinctive item appears at either a target or distractor location in a search display.	Speeded performance when the irrelevant item happens to be the target.
Pre-cueing	An irrelevant spatial cue precedes a search display.	Slowed search performance when the cue appears in a location subsequently occupied by a distractor.

attentional capture by measuring the effects of an irrelevant stimulus on performance of a primary task (see Table 1).

The logic underlying the 'Additional Singleton' task is the most intuitive of these paradigms: subjects perform a visual search task, and one item in the search display has a unique, distinctive feature (the 'singleton') that is unrelated to the search task and this item is never the target item (see Fig. 1). For example, subjects might search a display for a uniquely colored item, and one item in the display would have a delayed onset relative to the rest of the display. Attention is thought to be captured if performance is slowed by this irrelevant singleton relative to performance on trials with no irrelevant singleton. Observers have no strategic reason to search for the irrelevant singleton because it is never the target, so attending to it will always slow performance[1].

The 'Oculomotor Capture' paradigm adopts essentially the same approach, but capture is reflected by an inappropriate eye movement toward the irrelevant item rather than by a slowed response to the target. Any saccades toward the late-onset item provide evidence for attentional capture, even if subjects are unaware that they moved their eyes away from the target of their search[2]. In the Additional Singleton task, when attention is distributed across the entire display, both color and onset features lead to slowed performance[1]. Similarly, late onsets seem to draw the eyes away from the target of the search, possibly contributing to the response-time effect[2]. The effects on response time are present whenever observers are searching for a unique item (e.g. the only circle among squares) and the additional singleton is unique on a different dimension (e.g. the only blue item when all others are red).

The 'Irrelevant Feature Search' paradigm uses displays similar to those of the Additional Singleton paradigm (see Fig. 2) with one important difference: the irrelevant feature can also be the target of the search. Consequently, the dependent measure used to infer attentional capture is somewhat more complicated. Rather than simply measuring an overall reduction in search speed, capture is inferred from a change in the slope of the function relating response rate to the number of distractors in the display. Subjects perform a relatively difficult visual search task for which search latency increases as the number of distractors increases (e.g. searching for a rotated L among rotated Ts)[3]. On each trial, one of the items in the search display has a feature that differs from all the other items in the display (e.g. a different color or luminance). Yet, this distinctive feature in no way predicts the target location, and subjects know this. Any item in the display, including the target, is equally likely to have the distinctive feature; the feature appears in the same location as the target no more often than in any given distractor location. When there is no irrelevant feature or when the irrelevant feature is a distractor, search latency will increase as the number of distractors increases (the difference between these two conditions should be relatively small as subjects are likely to examine a distractor first in either case). By contrast, when the distinctive item also happens to be the target of the search, if it captures attention, response latency should be relatively unaffected by the number of distractors – the distinctive item should be the first one examined regardless of the number of distractor items.

In this paradigm, only the sudden onset of a new object consistently captures attention[4–6] and other features (e.g. color, orientation, etc.) do not affect search slopes[4,7] (but see discussion below). Yantis and colleagues[3,5,8] argue that attentional capture occurs because an onset signals the presence of a new object and requires the formation of a new object file[9,10]. Other sorts of distinctive features, including luminance changes[3], do not signal the existence of a new object, hence they do not capture attention[4,7]. As in the Additional Singleton and Oculomotor Capture paradigms, observers need not be aware of the new onset for it to influence search slopes. In fact, Yantis reports that during debriefing, observers rarely report noticing the existence of onsets in these displays[11], even though their response latencies are affected.

Results from all three of these paradigms suggest that onsets exogenously capture attention. However, findings differ for color and other salient features. In the Additional Singleton paradigm, the presence of a uniquely colored item increases search times, suggesting that it captured attention. In contrast, in the Irrelevant Feature Search paradigm, search rates are no faster when the target happens to have a unique color. Whether or not we should infer that color captures attention depends on the interpretation of the demands of these tasks. In typical variants of the Additional Singleton paradigm, subjects actively look for a unique item and another distinctive item captures attention. Perhaps if subjects adopt an attentional set for 'unique items', attention will be drawn by other unique items[12,13], suggesting that capture by the additional singleton results from the demands of searching for a singleton target. In fact, when observers perform an Additional Singleton task and the target is *not* a singleton but instead is defined by a particular value on a feature dimension (e.g. the only circle among squares and diamonds), then an additional irrelevant feature singleton (e.g. a uniquely colored item) can fail to capture attention[12,14,15]. This finding suggests that attention is not automatically drawn to the additional singleton. Rather, the singleton affects performance only when observers are searching for a singleton, which reflects the importance of top-down influences on performance. A strong form of this view, the 'contingent capture' hypothesis, suggests that all attentional capture depends on the attentional set of the observer[16]: stimuli will only capture attention if observers adopt the relevant attentional set. Most studies testing the contingent capture hypothesis have adopted the fourth paradigm that I discuss here: the 'Pre-cueing' paradigm (Fig 3).

In the Pre-cueing paradigm, subjects view an uninformative spatial cue prior to performing a search task. The cue is no more likely than chance to predict the location of the target in the search task (i.e. it can be a valid cue or an invalid cue), and subjects are aware that the cue is uninformative[14,17]. Attentional capture is inferred when performance is speeded if the cue happens to appear at the target location (valid cue) and slowed if it appears at a distractor location (invalid cue). Observers have no strategic reason to shift attention voluntarily to the pre-cue, so any attention shift must have been determined by the properties of the cue.

In this paradigm, the attentional set of the observer plays a critical mediating role in attentional capture, even by abrupt onsets[16,18–20]. For example, color pre-cues only capture attention if subjects are searching for a color target[16]. If the attentional set is for dynamic stimuli (e.g. a motion or late-onset target), only dynamic pre-cues capture attention and if the attentional set is for static stimuli, only static pre-cues capture attention[17]. Furthermore, when observers are searching for a specific feature value (e.g. green), only pre-cues of the same value (e.g. green but not red) capture attention[14]. In general, attention can be captured by any singleton when observers are in a singleton search mode, but if they are searching for a particular feature value, only cues with the same value will capture attention[14]. If attentional capture is contingent on the attentional set of the perceiver and on the demands of the task, then the evidence for bottom-up attentional capture by a visual feature is called into question – stimuli do not automatically capture attention owing to their salience alone[16–19,21].

Summary of the implicit attentional capture debate

The primary debate in the literature on implicit attentional capture focuses on which features, if any, automatically draw attention regardless of the expectations and attentional set of the observer. Evidence from the Pre-cueing paradigm suggests that attentional capture does not occur in the absence of the appropriate attentional set. Findings from each of the other paradigms suggest that stimulus-driven capture can occur, particularly by the abrupt onset of a new object. All of these studies explore the degree to which subjects can ignore something they know to be irrelevant.

During performance of these tasks, observers often do not even notice the irrelevant feature despite its effects on their search performance[11]. In fact, even distinctive features that are presented below a subjective threshold for awareness can

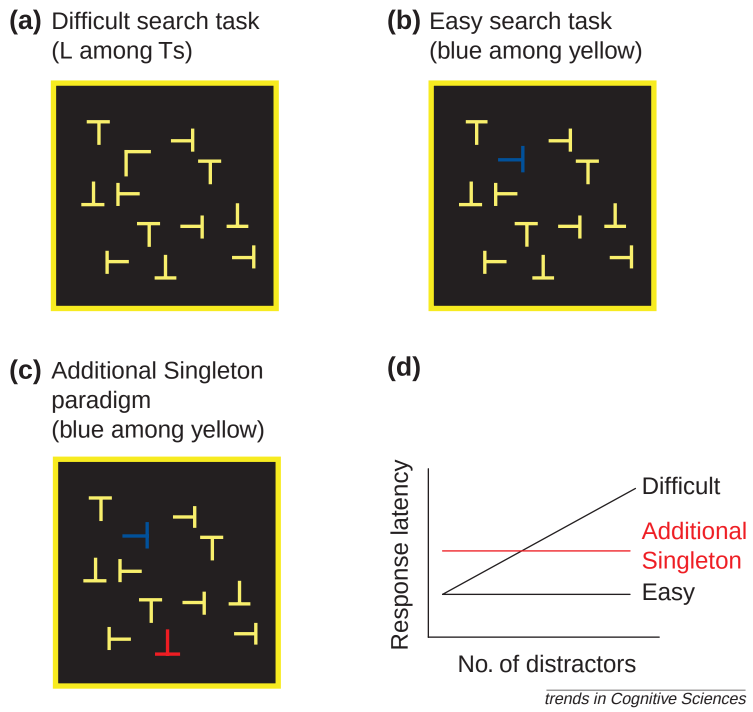

Fig. 1. Visual search and the 'Additional Singleton' paradigm. Three search conditions are shown, with idealized results. The primary measurement in visual search involves the increase in response time that results from adding additional distractor items to the display. The slope of the function relating target detection latency to increases in the number of distractor items provides an estimate of the efficiency of visual processing. **(a)** a difficult visual search in which subjects must find a rotated L among rotated Ts. A difficult search produces an increase in response latency as the number of distractor items increases **(d)**, presumably because each additional item demands some quantity of attentional focus. Easy search tasks, such as **(b)**, in which the subject must find the blue target, demand fewer attentional resources, thus response latency is relatively unaffected by the number of distractor items in the display. Studies of attentional capture rely on these consistent patterns of performance in order to examine the effects of a critical feature on either the overall speed of search or on the relationship between search latency and the number of distractor items. **(c)** The 'Additional Singleton' paradigm used to study attentional capture. In this task, subjects search for a blue target among yellow distractors, with one distractor item being uniquely colored (in this case, red). If the irrelevant item captures attention, performance should be slowed relative to that in the easy search. **(d)** Ideal results are shown for trials when the target is present in the display for all three types of search task.

implicitly capture attention and affect performance[22]. Evidence for implicit attentional capture is critical to understanding the mechanisms underlying visual search and for determining whether a perceptual event can automatically influence performance. Furthermore, such implicit effects can have a dramatic influence on our accomplishment of real-world tasks and goals. For example, much of our driving performance probably reflects implicit detection of salient events (e.g. cars turning or slowing down) leading to corresponding adjustments to our behavior. A good proportion of perception occurs without awareness, and given the capacity limitations on attention and memory, we need to be able to adjust our behavior without necessarily becoming aware of the cause or even the need for adjustment.

Yet, these implicit effects on behavior might not embody all aspects of attentional capture, particularly the intuitive notion that attentional capture should lead to awareness. Will we not automatically become aware of salient events in our visual environment, particularly events that have behavioral consequences? Would we automatically become aware of a salient new object if it unexpectedly appeared in front of us?

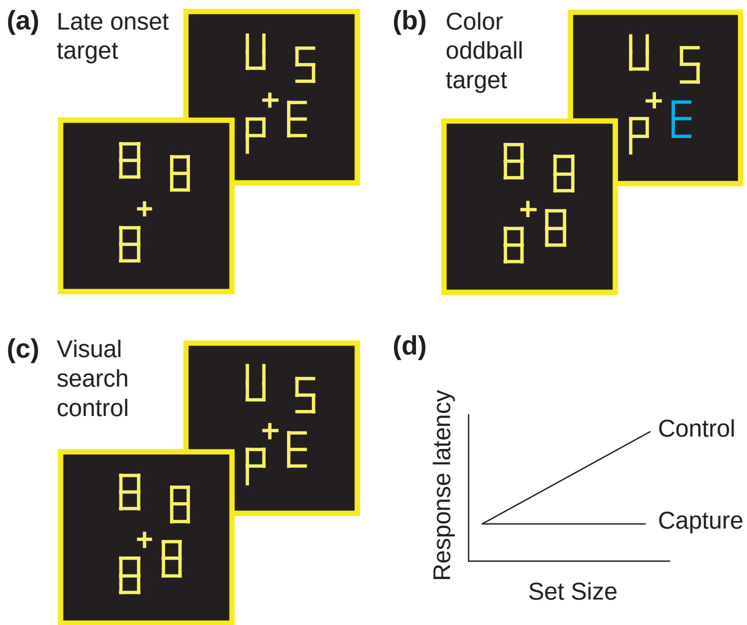

Fig. 2. The 'Irrelevant Feature Search' paradigm. (a–c) Three conditions of this attentional capture paradigm are shown, with idealized results. The control condition **(c)** is a typical visual search task. The initial frame consists of an array of 'figure 8' symbols. These symbols are then replaced with letters by deleting line segments. Observers must determine whether an E or an H is in the display (other tasks can also be used). In the late onset version of this search task **(a)**, when the figure 8s are replaced with letters, an additional letter appears in a location that had not previously contained a figure 8. This late-onset item is sometimes a target and sometimes a distractor. It does not predict the location of the target letter and subjects know it is not predictive. Similarly , a distinctive irrelevant feature **(b)** also does not predict the target location. Attentional capture is indicated by a reduced search slope when the distinctive item happens to be the target letter **(d)**.

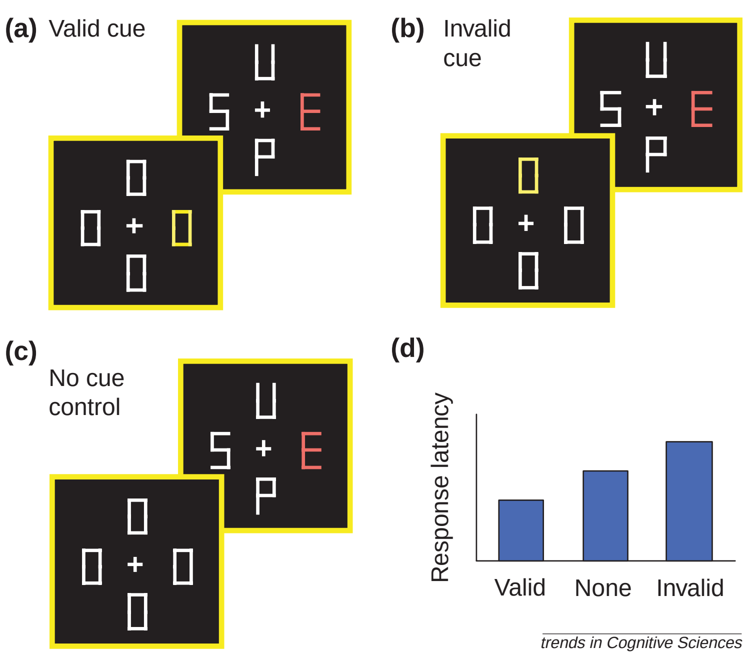

Fig. 3. The 'Pre-cueing' paradigm. Three critical conditions of the Pre-cueing paradigm are shown, with idealized results. In this paradigm, the task is to determine whether the unique item in the second frame is an E or an H. There is always a unique item and it is always either an E or H. In the valid cue condition **(a)**, the initial display contains a color cue in the location that will be the target on the next display. In the invalid cue condition **(b)**, the color cue is in a location that will be occupied by a distractor on the next display. In the no-cue control condition **(c)**, all of the cues are the same color. Attentional capture is indicated by a faster decision on valid cue trials than on control trials as well as by slower decisions on invalid cue trials **(d)**. As in the Irrelevant Feature Search paradigm (Fig. 1), across all trials the cue does not predict the target location, and observers know it is unpredictive.

And, if so, would attentional capture allow us not only to modify our behavior implicitly in order to accomplish an existing goal but also to select a new behavioral goal?

Explicit attentional capture (inattentional blindness)

Although we might intuitively believe that unusual, unexpected and salient objects will capture attention, leading to awareness, they often do not. Perhaps you have had an automobile accident and the other driver claimed he did not see you even though you were right in front of him. Of course, the driver's performance might have been affected if your car implicitly captured attention, but that would do little to resolve the question of why he did not *see* you and it probably could not have prevented the collision. In most real-world settings, the critical question of interest is not whether an object will implicitly affect performance, but whether it will explicitly capture attention and reach awareness, thereby allowing us to modify our behavior and select new goals. Although much, if not most, of perception and performance occurs without awareness, we feel that when salient events occur, we should become aware of them so that we can intentionally change our behavior.

The implicit paradigms explore how well observers can ignore something they *expect* but know to be *irrelevant*, whereas in explicit attentional capture, the critical question is how likely subjects are to notice something that is potentially *relevant*, but that they do *not* expect (see Box 1 for a discussion of the role of expectations). Recent studies of explicit attentional capture reveal a surprising degree of blindness to salient or unusual events that we might expect to capture attention. For example, observers often fail to notice surprisingly large, but unexpected changes to their visual world, such as a change to the identity of the central actor in a brief motion picture[23–25]. Most subjects intuitively believe that such changes should capture attention and be detected, both because of their magnitude and their potential behavioral relevance[26]. More relevant for the present review, observers sometimes fail to notice an unexpected object or event altogether – a phenomenon now known as 'inattentional blindness' (IB)[27–32]. Studies of inattentional blindness are among the few direct explorations of explicit attentional capture by complex visual events.

Inattentional blindness for distinctive objects

Recently, Arien Mack, Irvin Rock and their colleagues developed a paradigm for studying IB and explicit attentional capture[27–30] (here referred to as the 'static IB paradigm'). In their task, subjects decided which arm of a briefly presented cross was longer. After several such trials, subjects viewed a critical trial during which another object unexpectedly appeared along with the cross. Afterwards, subjects were asked whether they had noticed anything that had not been present on the previous trials. When the cross appeared at fixation and the unexpected object appeared away from fixation, approximately 25% of subjects were inattentionally blind: the unexpected object did not explicitly capture attention and they did not notice it. Interestingly though, when the cross appeared away from fixation and the unexpected object appeared at fixation, nearly 75% were inattentionally blind[28]. Even when the object was a different color or moved stroboscopically, observers

Box 1. What is an unattended stimulus?

In order to study how objects capture our attention, researchers must first find a way to present an unattended stimulus that could potentially draw attention. But how can we be certain that an object is not attended to? How can we operationally define the absence of attention?

One approach has been to define attention in terms of its consequences for perception. On this definition, the absence of such consequences implies the absence of attention. For example, we might argue that attention reduces the threshold for perceiving motion. For a near-threshold stimulus, attention might be necessary for the perception of motion; if subjects then fail to detect motion, the absence of detection implies the absence of attention. However, given that researchers are often interested in the consequences of inattention for perception, such definitions risk circularity.

A related approach is to link attentional capture to more easily identified or operationalized behaviors. In studies of dichotic listening, for example, subjects actively attend to a speech stream presented to one ear (typically by repeating each word immediately after hearing it) while ignoring supposedly irrelevant words presented to the other ear (Refs a–c). Researchers assume that listeners do not attend to the ignored stream, and for the most part, its semantic content is usually not retained. However, findings from dichotic listening studies also suggest that this 'unattended' information does receive substantial processing even though most of the content is not remembered (Refs d,e). For example, subjects do perceive some physical properties of the ignored stream (e.g. pitch) and they often detect highly salient words (e.g. their own name; Ref. c). The contents of the ignored stream can also influence interpretations of words in the attended stream (Ref. f), and if the attended message unexpectedly switches from one ear to the other, observers temporarily attend to the message from the wrong ear (Ref. b).

Yet, even if we accept the premise that deliberately ignored information is typically unattended, the inference that unattended information is processed semantically might be unfounded. In order to draw this conclusion with certainty, researchers must demonstrate that the ignored stream was completely ignored throughout the task (Ref. d). Even if a stream is supposed to be ignored, that does not guarantee that it will be ignored in practice. Listeners might shift attention from one stream to the other rapidly, thereby periodically attending the 'intentionally ignored' stream. Without a way to verify that a stimulus was consistently ignored, we cannot be certain it was unattended, and inferences about the perception of unattended stimuli will be invalid.

The implicit attentional capture paradigms (see main text) do not rely on the subject's ability to ignore part of a display continuously, but they do presuppose that a stimulus that is known to be irrelevant will be unattended; any stimulus that is known by the perceiver to be irrelevant to the primary task will be unattended because the subject knows that attention to that stimulus would not improve task performance. Thus, any effect of this unattended stimulus on performance can be attributed to attentional capture. However, as in dichotic listening, one problem with this assumption is that observers might still periodically shift attention to the stimulus voluntarily even though they know it is irrelevant to the task (Ref. d). In other words, whenever subjects expect a stimulus to appear, we cannot eliminate the possibility that they periodically attend to that stimulus. Of course, if they are *unaware* that a distinctive stimulus has appeared – as they seem to be in some implicit attentional capture tasks – such voluntary shifts are less likely.

Consequently, the ideal way to study attentional capture by an unattended item is to use a stimulus that is both unrelated to the primary task and also *unexpected*. If a feature is unexpected, it cannot influence task strategies. Assuming that salient or unusual stimuli will be remembered and reported once they have been attended – and given that the postulated function of explicit attentional capture is to bring unusual and potentially important objects or events to awareness – if an unexpected object does explicitly capture attention, it is likely to be reported later (for a view that rejects this assumption, see Box 3). Similarly, if a salient unexpected object fails to capture attention explicitly, it should be unavailable for conscious report. This logic underlies work on inattentional blindness and explicit attentional capture (see main text).

References

a Cherry, E.C. (1953) Some experiments upon the recognition of speech, with one and with two ears. *J. Acoust. Soc. Am.* 25, 975–979
b Treisman, A. (1964) Monitoring and storage of irrelevant messages in selective attention. *J. Verbal Learn. Verbal Behav.* 3, 449–459
c Moray, N. (1959) Attention in dichotic listening: affective cues and the influence of instructions. *Q. J. Exp. Psychol.* 11, 56–60
d Holender, D. (1986) Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: a survey and appraisal. *Behav. Brain Sci.* 9, 1–66
e Kahneman, D. (1973) *Attention and Effort*, Prentice-Hall
f MacKay, D.G. (1973) Aspects of the theory of comprehension, memory, and attention. *Q. J. Exp. Psychol.* 25, 22–40

were often inattentionally blind[28]. These findings show that a salient new object does not always explicitly capture attention. Of course, even in the absence of explicit attentional capture, the object may still implicitly affect performance (see Box 2).

Selective looking and the perception of unexpected events

Although the findings of IB suggest that novel, distinctive objects do not necessarily explicitly capture attention, perhaps attentional capture failed in these experiments because the objects were static and presented too briefly. During the 1970s and 1980s, the 'selective looking' paradigm was developed as a visual analog of dichotic listening[33–35] to explore the detection of sustained, dynamic, unexpected visual events[32,36–40]. In a selective-looking task, two simultaneous events are presented in a single display, and observers monitor one of them. In the first study to use this paradigm[36], observers viewed a hand-slapping game superimposed on a display of three people passing a basketball (both displays were partially transparent). After several trials of monitoring one event or the other, subjects viewed several critical trials in which an unexpected event occurred in the ignored display. On one trial, the players in the hand game stopped and shook hands. In another, the players in the basketball game 'lost' the basketball and continued to play, pretending to make passes. In yet another trial, the hand game players stopped and briefly passed a ball back and forth before resuming their hand game. When subjects were attending to one event, they typically did not notice the unexpected event in the ignored stream. They were 'inattentionally blind', and the events did not explicitly capture attention.

Box 2. Implicit attentional capture in the absence of explicit attentional capture

Although most studies of explicit attentional capture focus on whether or not observers notice an unexpected object, even in the absence of awareness the object might still influence performance. In other words, the object might implicitly capture attention even when it fails to do so explicitly. Several studies using the static inattentional blindness paradigm (see main text) have explored this question by examining whether observers show priming for the unexpected object that they did not notice. For example, observers are more likely to complete word fragments with a word that had appeared in the display even if they had not reported seeing it (Ref. a). Furthermore, even unattended background information can influence performance. Moore and Egeth used a variant of the static IB paradigm in which the cross was replaced by two horizontal lines and subjects were asked to judge which was longer (Ref. b). Random dots appeared in the background of the display on each of the initial trials. On the critical trial, the dots were arranged to produce the inducing elements of either the Ponzo or the Müller–Lyer illusion. Although subjects rarely noticed the pattern in the dots, their judgments of line length were clearly influenced by the illusions. These effects suggest that attention was implicitly, but not explicitly, captured by the unexpected object. Although subjects could not report the configuration of the dots and in fact never noticed that they were grouped to form the illusion, their judgments were still influenced by the dot configuration. Further studies are needed to explore implicit attentional capture in the absence of explicit attentional capture, especially in the context of selective-looking paradigms (see main text).

References

a Mack, A. and Rock, I. (1998) *Inattentional Blindness*, MIT Press

b Moore, C.M. and Egeth, H. (1997) Perception without attention: evidence of grouping under conditions of inattention. *J. Exp. Psychol. Hum. Percept. Perform.* 23, 339–352

An even more dramatic example of inattentional blindness (IB) comes from a selective-looking study that used a display with two superimposed teams, each playing a ball game[40]. When observers monitor one of the two overlapping teams and not the other (e.g. the three players wearing white shirts and not the three players wearing black shirts), they often failed to see a woman with an open umbrella appear from one side of the screen and walk across the display[32,40]. The appearance of this new, salient object did not capture attention.

Recently, Christopher Chabris and I set out to replicate and extend these studies and to revive the selective-looking paradigm as a tool for the study of attentional capture[31,41]. As in the basketball-game studies[40], our subjects counted the passes made by either the white team or the black team. The two teams and the unexpected event were filmed separately and then superimposed into a single video display to replicate the original displays. After about 45 seconds of the display, while the subjects were performing the counting task, a woman carrying an open umbrella walked across the display and exited the other side five seconds later (Fig. 4a). As in the earlier study, many subjects did not notice the umbrella woman (57% in our study and approximately 70% in the original). We also completed another set of conditions with a different unexpected object: a person wearing a gorilla suit. Again, we found a great deal of IB: on average, 73% of subjects failed to notice the partially transparent gorilla walk across the display.

Fig. 4. Still frames from the Simons and Chabris selective-looking paradigm. (a) Three frames from the transparent 'umbrella woman' event. The umbrella woman enters from the left and proceeds across the display, exiting on the right. **(b)** Three frames from the opaque 'gorilla' event. In both cases, the unexpected object was visible for 5 s as it traversed the display. (Adapted from movies used in Ref. 31.)

Although these studies suggest that salient new objects in complex displays do not explicitly capture attention, the degree of inattentional blindness could have been due to some oddity of the displays. Partially transparent displays are not typical of our real-world visual experience, so they may have impaired subjects' ability to detect the unexpected object. Therefore, we also tested subjects with a set of displays in which all of the players and the unexpected object were opaque and could occlude each other (Fig. 4b). The events were filmed in a single 'take' from the same camera position used to make the transparent films. If IB in the earlier studies and in our replication were due to some oddity caused by the transparent displays, then subjects should easily detect the umbrella woman and gorilla in these opaque conditions. However, they did not. On average, approximately 35% of subjects did not see the fully visible umbrella woman and gorilla. In one extra condition, the opaque gorilla stopped halfway across the display, turned to face the camera, thumped its chest, and then exited on the other side of the screen. Even in this condition, half of the observers did not see it! In fact, when we showed the video again after explaining what had occurred, observers were often shocked, sometimes even exclaiming, 'I missed *that*?!' Most observers intuitively believe that unusual events will explicitly capture attention[26], hence their surprise when learning of the unexpected event in our studies (shortened, compressed versions of these displays can be downloaded from http://www.wjh.harvard.edu/~viscog/lab/demos.html). Across all conditions

Box 3. Inattentional blindness or inattentional amnesia?

Recently, Wolfe and Horowitz have raised a potential alternative explanation for findings of change blindness and inattentional blindness (Refs a,b). They suggest that failures to report changes or unexpected objects reflect not a failure of perception or of attentional capture, but a failure of memory. Subjects might attend to an object, consciously perceive it, and then forget it by the time they are asked about it. Although we can safely argue that a stimulus was attended to if subjects can recall it, we cannot necessarily infer that it was unattended to if it was not recalled. In practice, these two explanations, blindness and amnesia, might be empirically inseparable. No matter how quickly subjects can be asked about a critical event, the questioning will still occur after the event. If observers fail to report it, proponents of the inattentional amnesia hypothesis can claim a failure of memory rather than a failure of perception.

In the case of the dramatic findings of inattentional blindness using the selective-looking paradigm (Refs c,d), the amnesia account seems somewhat less plausible – it is somewhat hard to imagine seeing a gorilla that was visible for up to 9 seconds, and then forgetting it immediately after the task (see Fig. 4 in main article). Of course, a possible rejoinder to this argument is that observers might see the gorilla, but not encode it as a gorilla. Under these conditions, it would seem slightly more plausible that they might consciously perceive something present in the display and still not recall it immediately after the task. Essentially, this amounts to inattentional agnosia (thanks to Jeremy Wolfe for suggesting this alternative).

Although the amnesia account may not be the most parsimonious explanation for the selective-looking results, it does raise a more important issue. Any measure of performance that requires attentive encoding (e.g. a verbal report of noticing) is likely to underestimate the amount of information that has actually been encoded. As findings from the implicit attentional capture paradigms suggest (see main text), even in the absence of explicit attentional capture by, and conscious awareness of, the critical object or feature, performance might still be affected. Performance measures based on verbal reports are even more likely to underestimate the amount of information *attended* because some subset of attended information may not be fully encoded. And, even if we could establish what information had been attended, other information might be implicitly represented without attention. In other words, the failure to report an object does not mean that the object did not implicitly capture attention. It need not be forgotten – it could just be consciously inaccessible.

This position is entirely consistent with other findings in the attention literature. For example, patients with visual neglect show evidence of processing of objects they claim not to have seen (Ref. e), observers show priming for stimuli that were presented during an attentional blink (Ref. f) and for a changed object or location even when they did not report the change (Refs g,h). Our research is currently exploring the possibility of such implicit representations in the face of inattentional blindness in the selective-looking paradigm.

References

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

b Wolfe, J.M. (1999) Inattentional amnesia. In *Fleeting Memories: Cognition of Brief Visual Stimuli* (Coltheart, V., ed.), pp. 71–94, MIT Press

c Simons, D.J. and Chabris, C.F. (1999) Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception* 28, 1059–1074

d Simons, D.J. and Mitroff, S. The incidental stance. In *Vision and Attention* (Harris, L.R. and Jenkin, M., eds), Springer-Verlag (in press)

e Driver, J. and Mattingley, J.B. (1998) Parietal neglect and visual awareness. *Nat. Neurosci.* 1, 17–22

f Shapiro, K.L., Arnell, K.A. and Raymond, J.E. (1997) The attentional blink: a view on attention and a glimpse on consciousness. *Trends Cognit. Sci.* 1, 291–296

g Fernandez-Duque, D. and Thornton, I.M. (2000) Change detection without awareness: do explicit reports underestimate the representation of change in the visual system? *Visual Cognit.* 7, 323–344

h Williams, P. and Simons, D.J. (2000) Detecting changes in novel, complex three-dimensional objects. *Visual Cognit.* 7, 297–322

that we tested, approximately 46% missed the unexpected event. Subjects were more likely to notice the opaque events (65%) than the transparent ones (42%), but even fully visible, dramatic, and sustained events did not consistently capture attention. (Of course, it should be remembered that they might have implicitly captured attention; see Box 2.)

Conclusions

In the static IB paradigm, observers often fail to notice the onset of a new, unexpected object in the display. In some respects, this finding is consistent with findings from the Irrelevant Feature Search paradigm showing that when attention is focused on some other part of a display, an abrupt onset might not implicitly capture attention[42]. Implicit attentional capture in the Irrelevant Feature Search paradigm requires that attention must not be focused elsewhere. The static IB results are consistent with this notion and suggest that when attention is engaged elsewhere, new objects can fail to explicitly capture attention as well. However, the selective-looking results raise some problems for this explanation for the failure of attentional capture. In the selective-looking paradigm, observers are focusing on multiple objects and the unexpected object literally passes through the attended locations. Attention is distributed across the display, but focused on other objects and events[32,43]. Thus, failed attentional capture cannot be attributed to spatially focused attention[42].

However, the more general notion of attentional engagement may help to explain both types of failed attentional capture. In both implicit and explicit paradigms, when attention is engaged, the likelihood of capture is reduced. In the static IB case and in the implicit search tasks, attention is often focused on a clearly defined spatial region and in selective-looking tasks, attention is engaged by objects and events. Do these two types of attentional engagement, location-based and object/event-based, have equivalent effects on capture?

In most real-world settings, observers are actively engaged in some task or goal, and the degree of attentional engagement can vary substantially. For example, driving a car in traffic in a rainstorm will probably limit the focus of attention to a relatively small region, perhaps increasing the degree of engagement relative to driving under normal conditions. The degree of engagement may well influence the probability of both implicit and explicit attentional capture. Yet, no studies have assessed the effects of varying the nature or degree of

Outstanding questions

- Is the lack of explicit attentional capture specific to vision? Visual attention can focus exclusively on a single object or location, but perhaps auditory attention is less focused. If so, we might not need an attentional set to detect auditorily relevant stimuli. Real-world attentional capture often occurs in modalities other than vision, hence, studies of 'inattentional deafness' and attentional capture might be important.
- Are there any cases in which explicit attentional capture consistently occurs in the real world? In other words, would more ecologically relevant events be more likely to capture attention, either implicitly or explicitly? Most studies of attentional capture have focused on simple events, such as onsets, or simple feature differences, such as color. Perhaps more complex but important visual events would automatically draw attention. For example, a looming stimulus might attract attention because it implies an imminent collision.
- Can differences between the implicit and explicit attentional capture paradigms be attributed to the degree or type of attentional engagement? Abrupt onsets fail to capture attention when observers are focused on another location. Would the same be true if they were focused on another object or event? Would salient objects always capture attention if attention were distributed across the entire display and not focused on any particular object, event or location?
- To what extent do selective attention and inattentional blindness require active ignoring? The static IB paradigm requires only that subjects attend to the target cross, and no elements in the display are ignored. However, all studies using the selective-looking paradigm have included both attended and non-attended items. Would levels of IB increase in the static task if subjects were required to ignore other stimuli? Would levels of IB decrease in the selective-looking task if subjects had nothing to ignore?

attentional engagement on capture. Future studies are clearly needed to explore implicit and explicit attentional capture while systematically varying the degree and type of attentional engagement.

Taken together, the similarity of the results from the static IB paradigm and the selective-looking paradigm suggests that IB may be a pervasive aspect of visual perception (for an alternative explanation of IB, see Box 3). More importantly, the results suggest that the appearance of a new object does not automatically capture attention – at least not explicitly. These findings raise the intriguing possibility that explicit attentional capture by a new visual object simply does not occur in the real world. Unless subjects adopt an attentional set for the appearance of a new object or they are not focused on any other objects, events or locations, it is unlikely to capture attention exogenously. This somewhat radical hypothesis would suggest that our intuitions about attentional capture reflect a metacognitive error: we do not realize the degree to which we are blind to unattended and unexpected stimuli and we mistakenly believe that important events will automatically draw our attention away from our current task or goals. Although such events might implicitly capture attention, thereby affecting the performance of our current task, they might not explicitly capture attention. If true, this hypothesis would require a rethinking of the rationale for studies of attentional capture.

The study of attentional capture is often motivated by a desire to explain how, for example, an animal detects the appearance of a predator, or how we notice when a pedestrian steps in front of our car. Explorations of the causes of automobile accidents are consistent with the claim that such events do not explicitly capture attention: nearly 50% of fatal automobile accidents are attributed to some driver-related factor, including inattention and distraction[44]. In other words, drivers often do not see salient and important objects. This fact can be rephrased in terms of attentional capture: if observers are attending to their driving (e.g. the car in front of them, road signs, etc.), and if they do not expect pedestrians to step in front of the car, they are unlikely to see them.

In order to understand more fully the conditions that lead to attentional capture, further studies are needed that explore not just the effects of implicit attentional capture on performance, but also the interaction between the observer's expectations, the degree of attentional engagement, and the likelihood of explicit attentional capture. The literature on implicit attentional capture has focused on the influence of attentional set on capture. Findings of IB suggest that expectations and set may play an important role in explicit attentional capture as well. Ideally, these two distinct approaches can be combined to reveal conditions necessary for involuntary encoding of distinctive features, as well as the factors necessary to bring those features to awareness.

Acknowledgements

Many thanks to Christopher Chabris, Steve Most, Rebecca Reimer, Brian Scholl, and Steve Yantis for comments on earlier versions of this manuscript. The writing of this manuscript was supported in part by NSF grant no. BCS-9905578 and by a Research Fellowship from the Alfred P. Sloan Foundation.

References

1 Theeuwes, J. (1994) Stimulus-driven capture and attentional set: selective search for color and visual abrupt onsets. *J. Exp. Psychol. Hum. Percept. Perform.* 20, 799–806
2 Theeuwes, J. *et al.* (1998) Our eyes do not always go where we want them to go: capture of the eyes by new objects. *Psychol. Sci.* 9, 379–385
3 Yantis, S. and Hillstrom, A.P. (1994) Stimulus-driven attentional capture: evidence from equiluminant visual objects. *J. Exp. Psychol. Hum. Percept. Perform.* 20, 95–107
4 Jonides, J. and Yantis, S. (1988) Uniqueness of abrupt visual onset in capturing attention. *Percept. Psychophys.* 43, 346–354
5 Yantis, S. and Jonides, J. (1984) Abrupt visual onsets and selective attention: evidence from visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 10, 601–621
6 Remington, R.W., Johnston, J.C. and Yantis, S. (1992) Involuntary attentional capture by abrupt onsets. *Percept. Psychophys.* 51, 279–290
7 Hillstrom, A.P. and Yantis, S. (1994) Visual motion and attentional capture. *Percept. Psychophys.* 55, 399–411
8 Yantis, S. (1998) Control of visual attention. In *Attention* (Pashler,H., ed.), pp. 223–256, Psychology Press
9 Treisman, A. (1993) The perception of features and objects. In *Attention: Selection, Awareness, and Control. A Tribute to Donald Broadbent* (Baddeley, A. and Weiskrantz, L., eds), pp. 5–35, Clarendon Press
10 Kahneman, D. *et al*. (1992) The reviewing of object files: object-specific integration of information. *Cognit. Psychol.* 24, 175–219
11 Yantis, S. (1993) Stimulus-driven attentional capture and attentional control settings. *J. Exp. Psychol. Hum. Percept. Perform.* 19, 676–681
12 Bacon, W.F. and Egeth, H.E. (1994) Overriding stimulus-driven attentional capture. *Percept. Psychophys.* 55, 485–496
13 Theeuwes, J. and Burger, R. (1998) Attentional control during visual search: the effect of irrelevant singletons. *J. Exp. Psychol. Hum. Percept. Perform.* 24, 1342–1353
14 Folk, C.L. and Remington, R. (1998) Selectivity in distraction by irrelevant featural singletons: evidence for two forms of attentional capture. *J. Exp. Psychol. Hum. Percept. Perform.* 24, 847–858
15 Yantis, S. and Egeth, H.E. (1999) On the distinction between visual salience and stimulus-driven attentional capture. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 661–676

16 Folk, C.L. *et al*. (1992) Involuntary covert orienting is contingent on attentional control settings. *J. Exp. Psychol. Hum. Percept. Perform.* 18, 1030–1044

17 Folk, C.L. *et al*. (1994) The structure of attentional control: contingent attentional capture by apparent motion, abrupt onset, and color. *J. Exp. Psychol. Hum. Percept. Perform.* 20, 317–329

18 Folk, C.L. *et al*. (1993) Contingent attentional capture: a reply to Yantis (1993). *J. Exp. Psychol. Hum. Percept. Perform.* 19, 682–685

19 Gibson, B.S. and Kelsey, E.M. (1998) Stimulus-driven attentional capture is contingent on attentional set for displaywide visual features. *J. Exp. Psychol. Hum. Percept. Perform.* 24, 699–706

20 Warner, C.B. *et al*. (1990) Voluntary allocation versus automatic capture of visual attention. *Percept. Psychophys.* 48, 243–251

21 Gibson, B.S. (1996) Visual quality and attentional capture: a challenge to the special role of abrupt onsets. *J. Exp. Psychol. Hum. Percept. Perform.* 22, 1496–1504

22 McCormick, P.A. (1997) Orienting attention without awareness. *J. Exp. Psychol. Hum. Percept. Perform.* 23, 168–180

23 Levin, D.T. and Simons, D.J. (1997) Failure to detect changes to attended objects in motion pictures. *Psychonomic Bull. Rev.* 4, 501–506

24 Simons, D.J. and Levin, D.T. (1997) Change blindness. *Trends Cognit. Sci.* 1, 261–267

25 Simons, D.J. (2000) Current approaches to change blindness. *Visual Cognit.* 7, 1–15

26 Levin, D.T. *et al.* (2000) Change blindness blindness: the metacognitive error of overestimating change-detection ability. *Visual Cognit.* 7, 397–412

27 Newby, E.A. and Rock, I. (1998) Inattentional blindness as a function of proximity to the focus of attention. *Perception* 27, 1025–1040

28 Mack, A. and Rock, I. (1998) *Inattentional Blindness*, MIT Press

29 Mack, A. *et al.* (1992) Perceptual organization and attention. *Cognit. Psychol.* 24, 475–501

30 Rock, I. *et al.* (1992) Perception without attention: results of a new method. *Cognit. Psychol.* 24, 502–534

31 Simons, D.J. and Chabris, C.F. (1999) Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception* 28, 1059–1074

32 Neisser, U. (1979) The control of information pickup in selective looking. In *Perception and its Development: A Tribute to Eleanor J. Gibson* (Pick, A.D., ed.), pp. 201–219, Erlbaum

33 Cherry, E.C. (1953) Some experiments upon the recognition of speech, with one and with two ears. *J. Acoust. Soc. Am.* 25, 975–979

34 Moray, N. (1959) Attention in dichotic listening: affective cues and the influence of instructions. *Q. J. Exp. Psychol.* 11, 56–60

35 Treisman, A. (1964) Monitoring and storage of irrelevant messages in selective attention. *J. Verbal Learn. Verbal Behav.* 3, 449–459

36 Neisser, U. and Becklen, R. (1975) Selective looking: attending to visually specified events. *Cognit. Psychol.* 7, 480–494

37 Littman, D. and Becklen, R. (1976) Selective looking with minimal eye movements. *Percept. Psychophys.* 20, 77–79

38 Stoffregen, T.A. and Becklen, R.C. (1989) Dual attention to dynamically structured naturalistic events. *Percept. Mot. Skills* 69, 1187–1201

39 Neisser, U. (1976) *Cognition and Reality: Principles and Implications of Cognitive Psychology*, W.H. Freeman

40 Becklen, R. and Cervone, D. (1983) Selective looking and the noticing of unexpected events. *Mem. Cognit.* 11, 601–608

41 Simons, D.J. and Mitroff, S. The incidental stance. In *Vision and Attention* (Harris, L.R. and Jenkin, M., eds), Springer-Verlag (in press)

42 Yantis, S. and Jonides, J. (1990) Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *J. Exp. Psychol. Hum. Percept. Perform.* 16, 121–134

43 Haines, R.F. (1991) A breakdown in simultaneous information processing. In *Presbyopia Research: From Molecular Biology to Visual Adaptation* (Obrecht, G. and Stark, L.W., eds), pp. 171–175, Plenum Press

44 NHTSA (1999) *Traffic safety facts 1998: A compilation of motor vehicle crash data from the Fatality Analysis Reporting System and the General Estimates System.* DOT HS 808 983, National Highway Traffic Safety Administration.