Available online at www.sciencedirect.com
ScienceDirect
Cognition 107 (2008) 904–931
COGNITION
www.elsevier.com/locate/COGNIT

Attentional resources in visual tracking through occlusion: The high-beams effect ☆

Jonathan I. Flombaum [a,*], Brian J. Scholl [a,*], Zenon W. Pylyshyn [b]

[a] *Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA*
[b] *Rutgers University, Center for Cognitive Science, Psychology Building Annex, 152 Frelinghuysen Road, Piscataway, New Jersey 08854-8020, USA*

Received 19 August 2006; revised 8 December 2007; accepted 24 December 2007

Abstract

A considerable amount of research has uncovered heuristics that the visual system employs to keep track of objects through periods of occlusion. Relatively little work, by comparison, has investigated the online resources that support this processing. We explored how attention is distributed when featurally identical objects become occluded during multiple object tracking. During tracking, observers had to detect small probes that appeared sporadically on targets, distracters, occluders, or empty space. Probe detection rates for these categories were taken as indexes of the distribution of attention throughout the display and revealed two novel effects. First, probe detection on an occluder's surface was better when either a target or distractor was currently occluded in that location, compared to when no object was behind that occluder. Thus even occluded (and therefore invisible) objects recruit object-based attention. Second, and more surprising, probe detection for both targets and distractors was always bet-

☆ For helpful conversation and/or comments on earlier drafts, we thank Todd Horowitz, Allan Kugel, Josh New, and two anonymous reviewers. J.I.F. was supported by an individual NSF Graduate Research Fellowship and an NIH National Research Service Award predoctoral fellowship. B.J.S. was supported by NSF #BCS-0132444. Z.W.P. was supported by NIH 1-R01-MH60924. A subset of these experiments was presented at the 2006 annual meeting of the Vision Sciences Society.

* Corresponding authors. Tel.: +1 203 432 2518; fax: +1 203 432 7172.
E-mail addresses: Jonathan.Flombaum@yale.edu (J.I. Flombaum), Brian.Scholl@yale.edu (B.J. Scholl).
URL: http://www.yale.edu/perception/.

doi:10.1016/j.cognition.2007.12.015

ter when they were occluded, compared to when they were visible. This new *attentional high-beams* effect indicates that the ability to track through occlusion, though seemingly effortless, in fact requires the active allocation of special attentional resources.

Keywords: Multiple object tracking; Object-based attention; Object persistence; Inhibition; Occlusion

1. Introduction

During early stages of processing, the visual system dedicates a considerable amount of resources to recovering the three-dimensional structure of the environment that is lost on the retina during optical transmission (Marr, 1982). The resulting representations of surfaces and their occlusion relationships go on to constrain further visual processing and cognition (Nakayama, He, & Shimojo, 1995) – for instance, in visual search (He & Nakayama, 1992) and in the perception of apparent motion (He & Nakayama, 1994). This prioritization of surface representations, and the accompanying body of research exploring how these representations are constructed, can make it appear as though occlusion is a problem that is opened and closed, so to speak, quickly and early in vision. The occlusion relationships that are identified at these early stages seem to form the basis for later cognition, rather than pose a problem for later cognition.

Nearly all of this research has employed static displays, and if occlusion relationships remained stable over time, this perspective might be accurate. In the real world, however, where objects move, and where occlusion relationships are constantly in flux, this view may underestimate the extent to which occlusion presents a challenge to visual cognition. In particular, occlusion complicates the process of maintaining representations of objects as the same persisting individuals through time and motion. Such complications may require the frequent online application of extra resources to resolve, although this processing may not always be phenomenologically apparent. In the present study, we introduce a new phenomenon – *attentional high-beams* – that provides a window into the special resources that are in fact required to maintain persisting object representations through occlusion.

1.1. The problem of object persistence

Determining depth relationships in early vision involves solving a type of inverse mapping problem: three-dimensional relationships exist in the world itself, but are collapsed during optical transmission, and must then be recovered. A similar type of inverse problem characterizes the determination of object *persistence*: objects in the world maintain their identities as the same persisting individuals over time and motion, but these identity relationships are not explicit in the fragmented visual input that our eyes receive, and must be inferred for visual experience to be coherent.

Research with a variety of paradigms and phenomena has uncovered some of the principles that our visual system uses to meet this challenge. Perhaps the most central generalization of this work is that computations of object persistence prioritize *spatiotemporal* information over information about objects' surface features (Scholl, 2001). This is in some ways surprising since it is not how we tend to *reason* about objects: if we see a large blue object and then later see a small green object, we typically judge that they must be distinct individuals. Not so in visual processing: despite such judgments, we may very readily *see* them as successive states of the same persisting individual. In apparent motion, for instance, observers perceive two featurally dissimilar flashes as subsequent stages of a single object so long as they occur near enough in space and time (e.g., Burt & Sperling, 1981), and any differences in the spatial proximity of two objects will always trump any effects of their featural similarity (e.g., Dawson, 1991). Such effects are in no way limited to the discrete nature of apparent motion. In the Tunnel Effect, for example, observers see an object that disappears behind an occluder, and then automatically – and seemingly effortlessly – they perceive the motion of a single enduring individual when an object emerges from the other side of the occluder along a consistent trajectory (Burke, 1952; Michotte, Thinès, & Crabbé, 1964/1991). This percept of persistence obtains even if the object that reemerges looks strikingly different from the object that initially disappeared – for instance, if it dramatically changed its color (e.g., from green to red; Flombaum & Scholl, 2006) or its kind (e.g., from a kiwi to a lemon; Flombaum, Kundey, Santos, & Scholl, 2004).

1.2. Out of sight, out of mind?

Spatiotemporal priority is a robust computational principle that the visual system employs in many different contexts when drawing a correspondence between two views of an object separated in time and space. This kind of correspondence is inherently retrospective in that the visual system must take a current view of an object and link it with previous views (Kahneman, Treisman, & Gibbs, 1992), for example views of an object just before and after a period of occlusion. But what happens while an object is actually occluded, before it reappears – i.e., before such a correspondence can take place? What kind of active processing occurs during the moments that an object is not visible? One possibility is that the visual system does nothing active during these moments, simply matching the memory trace of an occluded object with the present percept of a disoccluding object, if and when an object disoccludes. Related work on static amodal completion is consistent with this hypothesis. In these experiments, it has sometimes been observed that while object-based attention can spread through occluded regions of an object (e.g., Behrmann, Zemel, & Mozer, 1998; Moore, Yantis, & Vaughan, 1998), the attentional benefits that apply to the object regions on both sides of the occluder do not always apply to the occluded regions themselves (e.g., Davis & Driver, 1997; Moore & Fulton, 2005; but see Haimson & Behrmann, 2001). Similarly, perhaps attention 'forgets' about objects when they become briefly occluded.

This seems unlikely, however, since many disparate measures – including single-unit recordings in macaque cortex (Assad & Maunsell, 1995), BOLD signal responses in adults (Olson, Gatenby, Leung, Skudlarski, & Gore, 2003), scalp electrophysiology from human infants (Kaufman, Csibra, & Johnson, 2005), and even

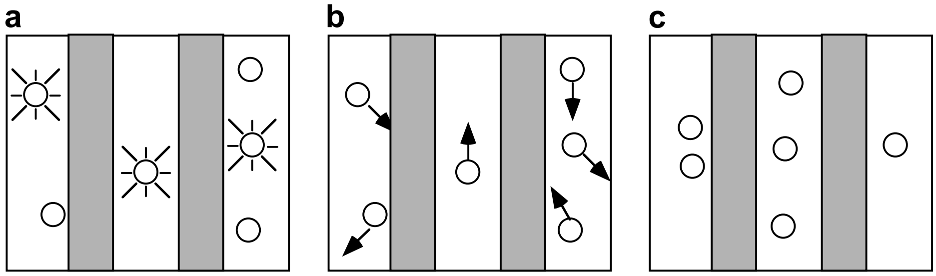

Fig. 1. A schematic depiction of a multiple object tracking trial with occluders. (a) Several identical items are presented, and a subset flash several times identifying them as targets. Next (b), all the objects move randomly and unpredictably in the display, becoming occluded whenever they happen to pass behind one of the two surfaces in the foreground. Finally (c), all the objects stop moving, and participants use the mouse to identify the targets that flashed in the beginning. (The occluders in the actual experiment were unshaded, defined only by their outlined contours on the screen.)

saccade patterns in human infants (Johnson, Amso, & Slemmer, 2003) – provide signatures of "keeping an object in mind" over even brief disappearances. In one study, for instance, cells in monkey parietal cortex continued firing in response to a moving target when the target disappeared briefly from the display (Assad & Maunsell, 1995). Related work with human adults (Olson et al., 2003) and nonhuman infants (Kaufman et al., 2005) has extended these findings, revealing patterns of neural activity associated with the maintenance of moving objects, specifically through occlusion (as opposed to other kinds of disappearance). Accordingly, the present study was designed to ascertain the degree to which special attentional resources might be dedicated to the processing of occluded objects.

1.3. Multiple object tracking

We explored object persistence with a *multiple object tracking* (MOT) task (Pylyshyn & Storm, 1988). In a typical MOT experiment (see Fig. 1), observers first see several (usually six to ten) featurally identical items (e.g., discs). At the start of a trial a subset of these items blink, identifying them as targets (and the items that do not blink as distractors). All of the items then move around the display following haphazard trajectories, and when they stop moving at the end of a trial (typically 5–20 s later), participants must identify the target items. Participants perform this task successfully for as many as five targets among a total of ten featurally identical items (i.e., among five distractors), indicating that they can simultaneously maintain an index of selection on each individual target (Pylyshyn & Storm, 1988; Scholl, Pylyshyn, & Feldman, 2001).[1]

[1] MOT is clearly attentionally demanding and effortful, leading most researchers to talk of MOT as an attentional process (e.g. Cavanagh & Alvarez, 2005). However, Pylyshyn (1989, 1994, 2001, 2003, 2007) has proposed that this task may involve several stages, and that the mechanism responsible for tracking the continuing identity of individual objects could itself be preattentive. Pylyshyn has hypothesized such a mechanism, called a visual index or 'FINST', that individuates objects and keeps track of their identity in a data-driven manner, despite changes in their properties or locations. We do not discuss this framework here, though for extensive discussion see Pylyshyn (2007) and Scholl (in press).

Several features of this paradigm make it especially useful for our purposes in the present study. First, MOT requires the maintenance of persisting representations for the target items without relying in any way on their appearance. Second, unlike paradigms that employ spatial cueing, MOT requires sustained and active attention over time, rather than brief shifts of attention and passive vigilance. Finally, previous work has demonstrated that tracking performance remains unimpaired when targets (and distractors) move behind surfaces in the display and become briefly occluded (Scholl & Feigenson, 2004; Scholl & Pylyshyn, 1999). In these ways, MOT furnishes a unique opportunity to ask about the resources that are required in order to maintain representations of several persisting objects even when they are not always visible. Based on the data reported below, we suggest that the apparent effortlessness with which we succeed in MOT through occlusion (relative to MOT without occlusion) belies the allocation of additional attentional resources to occluded objects. Much in the way that the brighter-than-usual 'high-beam' lights on a car provide extra illumination during special circumstances when the normal headlights are insufficient, the visual system appears to provide extra facilitation for tracking occluded objects – a task which seems to require these special resources.

1.4. The present study

Participants in these experiments tracked either two or three targets among either two or three distractors, in a typical MOT design. In addition, however, all of the objects were frequently occluded whenever they passed behind two 'walls' (vertical strips that divided the screen into three equally sized portions). Previous experiments with very similar displays have indicated that MOT performance is largely unaffected by the addition of such occluders: the presence of the gradual accretion and deletion cues serve as a cue to the visual system to maintain those object representations despite their disappearances (Scholl & Feigenson, 2004; Scholl & Pylyshyn, 1999).[2]

During each 20-s MOT trial, participants also had to detect sporadic probes that appeared briefly and haphazardly in various parts of the display – e.g., on an unoccluded target, on an occluder with nothing behind it, or on an occluder hiding a target or distractor (i.e., at the location of a currently occluded item). Probe detection performance thus provided a measure of the distribution of attention over the course of an MOT trial. Since its introduction by Posner and others (e.g., Downing & Pinker, 1985; Posner, 1980), the detection of transient probes has been perhaps the most common method for measuring the relative distribution of attention in a display, including attentional selection realized via inhibition of nontargets (e.g., Watson & Humphreys, 2000). This method has also previously been used with MOT to investigate the distribution of object-based attention along spatially extended objects (Alvarez & Scholl, 2005) and to investigate the relative allocation of attention to targets and distractors

[2] Such cues may be particularly useful – or even required – in online visual tracking when items disappear frequently and asynchronously. In contrast, occlusion cues may be less important when items must be encoded by spatial memory through other types of interruptions (Horowitz, Birnkrant, Fencsik, Tran, & Wolfe, 2006).

(Pylyshyn, 2006; Pylyshyn, Haladjian, King, & Reilly, submitted for publication; Sears & Pylyshyn, 2000). In the present experiments, we were particularly interested in what probe detection would reveal about the degree to which attention is allocated to 'invisible' objects that were hidden behind an occluding surface.

2. Experiment 1: Tracking through occlusion

In this first experiment, probes appeared sporadically throughout a trial, and each probe was equally likely to appear on:

(a) the center of a randomly chosen unoccluded target (or distractor)
(b) in empty space, but near a randomly chosen unoccluded target (or distractor)
(c) on the surface of an occluder, at the location of a randomly chosen and currently occluded target (or distractor)
(d) on the surface of an occluder near (but not directly at) the location of a randomly chosen and currently occluded target (or distractor)
(e) on the surface of an occluder when no items were currently behind it

These various possible probe locations are each depicted in Fig. 2. Because these conditions and displays are inherently dynamic, animations depicting each probe type can be viewed online at http://www.yale.edu/perception/high-beams/.

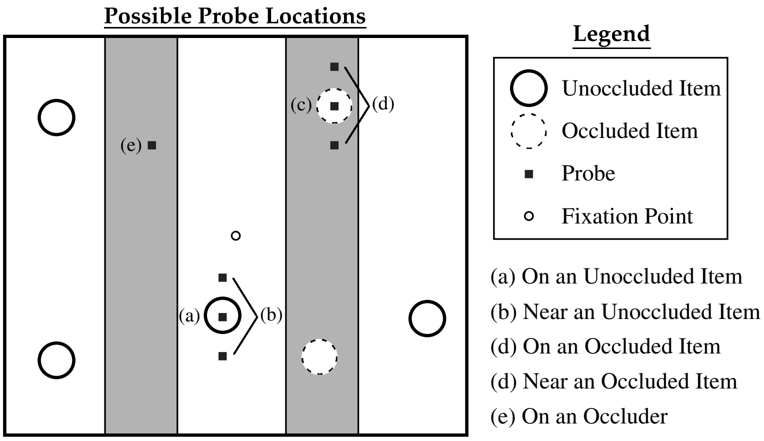

Fig. 2. A schematic depiction of the display and probe conditions (not to scale) in Experiment 1. Participants tracked three targets in a set of six identical items (targets and distractors), and they pressed a key whenever they saw a dim gray probe appear on the screen. Each probe was equally likely to appear on (a) unoccluded targets and distractors; (b) near unoccluded targets and distractors; (c) on the surface of an occluder at the location of a presently occluded target or distractor; (d) on the surface of an occluder *near* (but not directly at) the location of a presently occluded target or distractor; and (e) on the surface of an occluder when nothing was presently occluded by it. A small fixation circle was always present in the center of the display. (The occluders in the actual experiment were unshaded, defined only by their outlined contours on the screen.)

Two sets of comparisons in probe detection rates were of particular interest. First, we were interested in whether and how attention might be directed to occluded objects. Since an occluded object does not stimulate one's retinas and provides no direct cues to its present location, it is possible that attention simply cannot be directed to such an object in the first place. However, based on related neurophysiological work (Assad & Maunsell, 1995; Kaufman et al., 2005; Olson et al., 2003), we predicted that attention would be directed to occluded objects even when they are momentarily invisible. As a result, we expect better probe detection on occluders that are currently obscuring targets compared to occluders that are currently obscuring nothing.[3] We might intuitively predict no such effect for distractors, however, since distractors are presumably intentionally ignored during MOT. Second, we were interested in the relative detection of probes on occluded vs. unoccluded targets. Do occluded targets receive *less* attention than unoccluded ones, perhaps because it is difficult to direct attention to something that cannot be seen? Or do occluded targets actually receive *more* attention than unoccluded ones, perhaps in order to facilitate their representation through a visual interruption?

2.1. Method

2.1.1. Participants

Sixteen participants were tested in 50-min sessions and received course credit or a small monetary payment. The data from one participant were dropped from the final analysis due to an overall tracking accuracy (71%) that was more than two standard deviations ($SD = 6\%$) below the group mean (86%). All participants were naïve to the purposes of the experiment and reported normal or corrected-to-normal acuity.

2.1.2. Apparatus

The experiment was run on an iMac computer with a CRT display using custom software written in C with the VisionShell graphics libraries (Comtois, 2007). Participants sat approximately 50 cm from the screen without head restraint so that the entire display subtended $32.12° \times 27.84°$ of visual angle. The monitor refreshed at 117 Hz and all motion was perceptually smooth.

2.1.3. MOT displays

Each trial included six identical discs (subtending 1.74°) on a black background. Each disc had a black center surrounded by a .058° (one pixel) white outline. In addition to the discs, each display included a .58° fixation circle outlined in red, and two occluding surfaces that extended from the top of the display to the bottom, with widths of 3.48°, separated by 8.39°. Each occluder was defined by its .058° white bor-

[3] Note that the punctate probes used here do not appear to be presented at any particular depth, and so it is possible to see them as being presented on the invisible objects. In contrast, our predictions might not hold if the occluder and the objects were stereoscopically presented in different depth planes: there, we might expect that probes presented stereoscopically at the depth plane of the occluded objects would enjoy improved detection, but that probes appearing at the depth plane of the occluder itself might not.

der, with a black interior, and the two occluders were arranged so as to divide the display into three equally sized regions (as in Figs. 1 and 2). Eye movements were not monitored, and no special instructions were given concerning fixation since different fixation conditions have been found not to affect performance on this task.[4]

At the start of each trial, three of the six discs blinked on (57 ms) and off (57 ms) five times to identify them as targets. All the discs then moved around the display through haphazard paths for a total duration of 20 s. The discs moved through trajectories that were generated in advance and stored off-line; thus each subject observed the same set of trajectories, but in a different randomized order and with probes appearing at different times and places. Each trajectory file included the locations of each of the six items on each cycle of a 20-s trial. The trajectories were created as follows: at the start of a trial each item was assigned a random velocity vector that was updated on each frame. As each item moved through the display, it was 'repulsed' (via an inverse-distance-squared "force field") by each of the other five items and by the edges of the display (ensuring that items would not get too close to each other or intersect). This repulsion affected each item's velocity and direction causing approaching items to veer away from each other and sometimes to reverse direction. (See the online movies for examples.) The new position of an item on each frame was a function of the item's current position updated by its velocity vector, the sum of the repulsion vectors at that item, and a random adjustment to the item's velocity vector. In addition, starting positions for the items were constrained such that items could not begin occluded, within 3.48° of another item, or overlapping with fixation. The occluders in the display remained stationary and always occluded the moving items when they intersected.

2.1.4. Probe presentations and conditions

Attentional probes took the form of dim .12° gray squares presented for 100 ms each. Only one probe was presented at a time, and a random interval of 2500–6000 ms separated subsequent probes (as well as the latency to the first probe on each trial). Thus each trial included between three and seven probes that occurred an average of 4.3 s apart, resulting in an average of 347 probes for an entire experimental session composed of 80 trials.

Each probe was equally likely to occur in one of nine probe conditions: (1) *Target* probes appeared in the centers of randomly selected unoccluded target items. (2) *Distractor* probes appeared in the centers of randomly selected unoccluded distractors. (3) *Near Target* and (4) *Near Distractor* probes appeared one item radius (.87°) either above or below unoccluded targets and distractors, respectively. (5) *Occluded Target* and (6) *Occluded Distractor* probes appeared at the centers of targets and distractors

[4] Though the dynamics of eye movements during MOT are of interest (Doran, Hoffman, & Scholl, 2006; Fazl & Mingolla, 2007; Fehd & Seiffert, in press; Zelinsky, Neider, & Todor, 2007), constraints on fixation seem not to affect performance in any noticeable way. Pylyshyn and Storm (1988) eliminated trials on which subjects made eye-movements, and they obtained qualitatively identical results to other studies that employed no special constraints or instructions concerning fixation (e.g. Intriligator & Cavanagh, 2001; Scholl et al., 2001; vanMarle & Scholl, 2003; Yantis, 1992) or instructed subjects to maintain fixation but did not monitor eye-movements (e.g. Scholl & Pylyshyn, 1999).

when they were completely occluded – and thus always appeared on an occluder's surface. (7) *Near Occluded Target* probes and (8) *Near Occluded Distractor* probes appeared on an occluder's surface, one item radius above or below occluded targets and distractors. Probes in these eight conditions could not occur on or near an item whose nearest contour was closer than .87° above or below another item. (9) Finally, each *Occluder* probe appeared on the surface of an occluder that currently had no items behind it, with the constraint that probes had to be located at least .87° from each of the occluder's borders.

2.1.5. Procedure and design

Participants were given two tasks: (1) to keep track of each of the three target items for the duration of each 20-s trial, and (2) to press a key on the keyboard as soon as possible after detecting the appearance of a probe. Keypresses that occurred within one second of a probe's onset were recorded as hits, while keypresses that occurred at any other time were recorded as false alarms. The instructions stressed that participants should prioritize the tracking task, monitoring for probes as well as they could without sacrificing tracking accuracy. At the end of each 20-s trial, all the items stopped moving and participants used the mouse to click on the three items that they judged to be the targets. (Items that were occluded when a trial ended appeared on top of the occluders so that they could be selected as well.) As the mouse cursor moved around the display during the target selection period, the unselected item that was nearest to the mouse was highlighted in red, and previously selected items that were already selected were highlighted in green. No feedback was given about either tracking accuracy or probe detection.

Participants completed a total of 80 trials, after first completing 5 practice trials (the results of which were not recorded). In our analysis of these data, however, we treat each probe independently of the tracking trial that it came from. Because probe conditions depend on whether a probed item was a target or a distractor, our analysis of probe detection data includes only probes that occurred in perfect-tracking trials (i.e., trials for which all three targets were correctly identified).

2.2. Results

Overall tracking accuracy was high (86%) and commensurate with previous experiments (e.g., Scholl & Pylyshyn, 1999), indicating that occlusion does little to impair tracking performance. Similarly, participants had low false alarm rates, averaging only 17 false alarms over the course of 80 trials. A repeated measures ANOVA including the nine probe conditions revealed a main effect of probe kind on detection, $F(8,14) = 27.17$, $p < .001$, $\eta^2 = 1.94$. Table 1 lists the mean detection rate and standard error for each probe condition. We conducted additional planned comparisons motivated by pilot data (Flombaum, Scholl, & Pylyshyn, 2006) to examine three effects: distractor inhibition, attention to occluders, and attentional high-beams (i.e., attention to locations on occluders that hide objects).

Table 1
Probe detection by condition for each experiment, with mean detection % and standard errors in parentheses

Condition	Experiment 1	Experiment 2	Experiment 3	Experiment 4
Near Distractor	58.4 (6.2)	55.5 (5.9)	34.8 (3.7)	49.3 (4.2)
Distractor	31.3 (5.1)	26.8 (3.6)	27.1 (3.5)	38.5 (4.6)
Occluded Distractor	64.7 (6.0)	30.0 (4.0)	43.1 (2.0)	28.4 (5.4)
Occluder	55.0 (5.7)	43.8 (5.9)	31.3 (3.9)	54.3 (5.5)
Near Occluded Distractor	67.4 (5.8)	41.4 (5.2)	40.1 (3.4)	41.1 (5.0)
Near Target	68.5 (6.0)	60.2 (4.4)	56.6 (2.6)	64.4 (3.5)
Target	52.1 (5.1)	33.0 (3.9)	47.6 (4.4)	58.7 (3.6)
Occluded Target	74.1 (6.0)	38.0 (5.1)	60.0 (2.8)	45.1 (5.4)
Near Occluded Target	78.2 (4.9)	43.9 (5.2)	67.4 (2.8)	59.4 (3.7)

2.2.1. Distractor inhibition

Previous research has indicated that, contrary to intuition (and phenomenology), distractors are not just ignored in MOT, but are actively inhibited (Pylyshyn, 2006; Pylyshyn et al., submitted for publication). As depicted in Fig. 3a, our present results replicate these findings: probe detection was significantly worse for Distractor probes compared to Target probes ($t(14) = 5.49$, $p < .001$, $d = 1.05$), and also significantly worse for Distractor probes compared to Near Distractor probes ($t(14) = 7.26$, $p < .001$, $d = 1.23$).

2.2.2. Attention to occluders

This experiment also yielded evidence for object-based attention to currently invisible objects: as depicted in Fig. 3b, probe detection was significantly worse for Occluder probes, relative to both Occluded Target probes ($t(14) = 4.61$, $p < .001$, $d = .84$) and Occluded Distractor probes ($t(14) = 2.98$ s, $p = .01$, $d = .43$).

2.2.3. Attentional high-beams

Finally, the most novel and unexpected finding in our results, as depicted in Fig. 3c, was that probe detection was actually *better* for occluded items relative to unoccluded items, for both targets (Target vs. Occluded Target: $t(14) = 4.55$, $p < .001$, $d = 1.01$) and distractors (Distractor vs. Occluded Distractor: $t(14) = 10.10$, $p < .001$, $d = 1.55$).

2.3. Discussion

This experiment was designed to assess how (and whether) attention is allocated to moving objects that sometimes become momentarily hidden by a foreground surface. Our results replicated one previously discovered effect, and introduced two others:

2.3.1. Distractor inhibition

The phenomenology of MOT seems to involve selection of targets, but as is often the case with such selection, both facilitation and inhibition are involved. Such effects

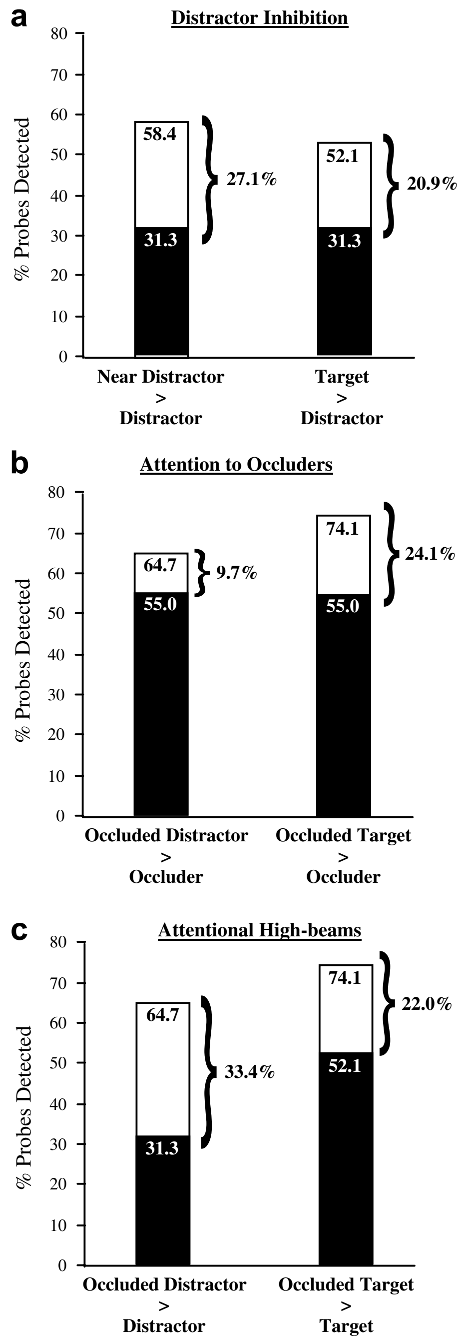

Fig. 3. The three primary effects from Experiment 1: (a) *Distractor inhibition*, wherein distractors are inhibited relative to both targets and empty space. (b) *Attention to occluders*, wherein probes on occluders are detected better when they are currently occluding either targets or distractors. (c) *Attentional high-beams*, wherein probes on both targets and distractors are detected better when occluded.

have previously been observed during MOT (Pylyshyn, 2006; Pylyshyn et al., submitted for publication), and were also seen here: detection of probes on distractors was not only impaired relative to targets (as expected), but also relative to empty space, suggesting that the distractors were actively inhibited.[5] As explored in Section 6, this effect is surprising, but consistent with the other results in this study – attention to occluders and attentional high-beams – suggesting that distractors are processed to a greater degree in MOT than intuition alone would suggest.

2.3.2. Attention to occluders

The primary purpose of this study was to explore how object-based attention during MOT interacts with momentary occlusion, and perhaps the most straightforward question in this regard concerns the degree of attention to the locations of occluded objects. The basic result of this experiment was simply that such attention exists: probes that occurred on occluders were detected better when the occluders were currently occluding objects. What is again surprising about this effect, however, was that it obtained for distractors in addition to targets. This implies that the visual system must 'know' each time that a distractor becomes occluded, which in turn means that distractors must be tracked to some degree – at a minimum, when they become occluded.

2.3.3. Attentional high-beams

Beyond comparisons between occluded objects and the occluders themselves, we also directly compared probe detection on objects when they were and were not occluded. We initially expected that maintained attention to objects during occlusion would manifest itself in equal detection in these conditions – in other words, that probe detection (at least for targets) would simply discount the occluders. In fact, however, this comparison yielded an unexpected result: probe detection for objects (again, *both* targets and distractors) was actually *better* when occluded. We dub this the *attentional high-beams* effect. Rather than discounting occlusion, attention seems to compensate for it. The maintenance of persisting object representations through occlusion appears to pose a special challenge, for which an added burst of attention is allocated, much like the way that one turns on the bright 'high-beam' lights on a car to illuminate the road when it is especially hard to see. This effect is especially striking given (1) that tracking through occlusion seems largely effortless (relative

[5] We note, in passing, that this experiment and the following ones include an additional surprising finding: better probe detection near targets than on targets. We have conducted follow-up experiments that explore the nature of this phenomenon (Flombaum & Scholl, 2007, submitted for publication), but we do not discuss them here since these effects do not interact with our occlusion results, and since they interact in subtle ways with aspects of motion processing that are beyond the scope of this paper. These probe conditions in empty space were included in the present study primarily to prevent participants from predicting the locations of forthcoming probes, and as such these probe conditions are not involved in the central comparisons, which concern only occluders, occluded objects, and unoccluded objects. As explored in Section 6, however, these results do lead to interesting complications for the nature of what has been called (and what we will continue here to call) 'distractor inhibition'.

to tracking without occlusion), and (2) that it does not feel like distractors are tracked at all, much less with special intensity during occlusion.

3. Experiment 2: Moving occluders

The results of Experiment 1 were interpreted with respect to the categorical status of the probe conditions. We argued that probe detection was better or worse because the probe had appeared *on an object*, *in empty space*, or *on an occluder*. There is an alternative to this interpretation, though: perhaps probe detection actually depends only on the relative visibility of the probes. Consider, for example, the comparison between probe detection on occluded vs. unoccluded objects. Which white contours are nearby in these comparisons to potentially interfere with one's seeing the probe? On unoccluded objects, the contours are very close by indeed since the probe always occurs within the confines of the visible object itself. For occluded objects, however, the object's contours are invisible, and the nearest contours (that form the boundary of the occluder itself) are a bit farther away. Thus the attentional high-beams effect might simply reflect a more pedestrian type of visual interference.

This possibility seems unlikely to us since none of the conditions seemed to produce a phenomenologically salient interference effect: it just does not seem like probes are hard to detect because other things are nearby. In addition, of course, this alternative hypothesis cannot account for enhanced probe detection on occluders when objects are currently occluded (since the relevant visible contours are identical in both cases). Controlling for this alternate explanation is tricky. While it would be possible in principle to adjust the sizes of the objects, and/or the widths of the occluders to better equate probe visibility, this would require an entirely different design, given the radically different shapes of the objects and the occluders (cf. Cepeda, Cave, Bichot, & Kim, 1998; Humphreys, Stalmann, & Olivers, 2004). As such, we took two different approaches to definitively rule out this alternate explanation in Experiments 2 and 3.

First, in the current experiment, we simply removed the need for attentional tracking by employing objects that remained stationary. Instead, observers simply had to remember which three disks flashed at the start of a trial. So that momentary occlusion of targets and distractors would still take place, the occluders translated back and forth throughout a trial, and we could thus contrast probe detection across the same probe conditions as in Experiment 1. If the attentional high-beams effect were due entirely to differences in probe visibility on occluded compared to unoccluded items, then the effect should obtain in the current displays as well, since they involved the same visual comparisons. In contrast, if the attentional high-beams effect reflects the unique demands of tracking moving objects through occlusion, then it should not be present in the current displays since those tracking demands did not exist.

3.1. Methods

Except as noted here, this experiment was identical to Experiment 1. We tested a new group of 15 naïve participants. After the three target discs flashed at the start of a trial, all six discs remained stationary, but the two occluders translated horizontally at a rate of 6.79° per second. The initial motion direction of each of occluder was set randomly at the start of each trial. The left occluder reversed direction whenever it moved within .29° of the leftmost edge of the display or within .87° of the display's vertical midline. The right occluder similarly reversed direction when it approached within .87° of the midline or within .29° of the display's rightmost edge. Instructions emphasized that participants should monitor for the appearance of transient probes without sacrificing accuracy for remembering the three targets. Probes appeared every 4.6 s, on average, and there were a total of about 310 probes across an experimental session.

3.2. Results and discussion

Participants' overall 'tracking' performance was high (95%), and they made an average of 18 false alarms across an experimental session. As depicted in Fig. 4 (with supporting data in Table 1), this experiment revealed no prioritization of attention to stationary occluded objects. A repeated measures ANOVA revealed a main effect of probe condition on detection, $F(8,14) = 5.21$, $p < .001$, $\eta^2 = .75$. Planned comparisons revealed evidence for distractor inhibition, but no evidence of attention to occluders, or of the attentional high-beams effect.

3.2.1. Distractor inhibition

Detection of Distractor probes was marginally worse than Target probes ($t(14) = 2.14$, $p = .05$, $d = .43$) and significantly worse than Near Distractor probes ($t(14) = 8.34$, $p < .001$, $d = 2.13$). It seems likely that this effect was actually due to differences in probe visibility for probes on the centers of distractors compared to probes in the empty space near distractors. In Experiment 3 we also find evidence for distractor inhibition while controlling for issues of probe visibility in a completely different way, and the original report of distractor inhibition in MOT (Pylyshyn, 2006) employed other methods to demonstrate that distractor inhibition effects cannot be explained entirely in terms of differences in low-level probe visibility.

3.2.2. Attention to occluders

This experiment yielded no evidence of attentional resources prioritizing foreground surfaces when those surfaces occlude targets or distractors: as depicted in Fig. 4b, probe detection was no different for Occluder probes relative to Occluded Target probes ($t(14) = 1.12$, $p = .28$, $d = .27$) and was actually worse for Occluded Distractor probes ($t(14) = 3.50$, $p = .004$, $d = .71$).

3.2.3. Attentional high-beams

As depicted in Fig. 4c, the targets and distractors in this experiment were not subject to attentional high-beams: probe detection was no different for occluded items relative to unoccluded items, for both targets (Target vs. Occluded Target: $t(14) = .88$, $p = .39$, $d = .29$) and distractors (Distractor vs. Occluded Distractor: $t(14) = .73$, $p = .48$, $d = .22$).

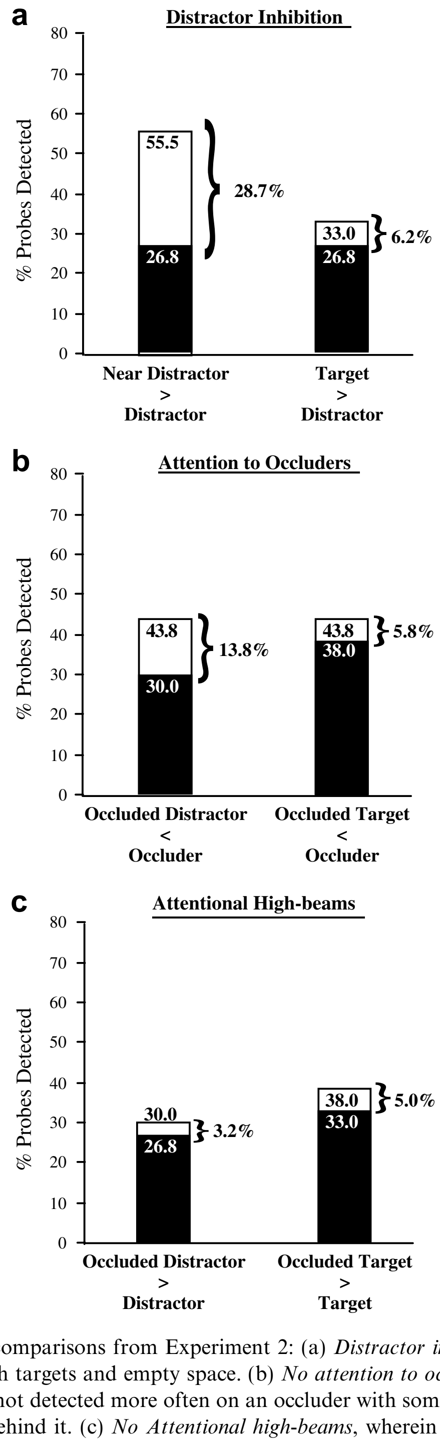

Fig. 4. The three primary comparisons from Experiment 2: (a) *Distractor inhibition*, wherein distractors are inhibited relative to both targets and empty space. (b) *No attention to occluders*, where in contrast to Experiment 1, probes were not detected more often on an occluder with something behind it compared to an occluder with nothing behind it. (c) *No Attentional high-beams*, wherein probes in the location of an occluded object were detected no more readily than probes on unoccluded objects.

In sum, simply removing the need to track the targets seems to have obviated the need for special attentional consideration of occluded objects relative to unoccluded ones (the high-beams effect), as well as occluding occluders relative to empty ones (attention to occluders). Thus these effects are not simply a product of low-level differences in probe visibility across conditions – since those same differences were present in this experiment, but the relevant effects were not. Instead, the special resources allotted to occluded objects and occluding surfaces appear to emanate directly from the tracking demands of MOT, reflecting special resources that mediate successful tracking through occlusion.

4. Experiment 3: Motion-defined objects with occlusion

In order to demonstrate further that the discoveries of Experiment 1 reflected the attentional demands of tracking through occlusion rather than lower-level visual interference, Experiment 2 employed the same conditions but without the demands of tracking, and showed that the relevant effects disappear. In this experiment we sought converging evidence from a complementary manipulation, by retaining the tracking task, but removing all potential visual differences across probe conditions. To do this, we simply removed the contours of the objects altogether: the background, the objects, and the interiors of the occluders were all defined by visual noise, with each pixel set randomly to either black or white (see Fig. 5). The regions of noise that constituted the objects then remained locally stable throughout their motion, so that these regions of noise temporarily replaced the background noise as the objects moved. Objects in such displays were thus defined solely by their motion and were invisible in any given static frame. Thus all local visual properties were equated directly for probes in all conditions.

4.1. Methods

Except as noted here, this experiment was identical to Experiment 1. We tested a new group of 13 naïve participants. The data from one participant were dropped from the final analysis because their number of false alarms (39) was more than two standard deviations ($SD = 11$) above the group mean (15).

The background of the display consisted of pure visual noise, in which each pixel was randomly set to either black or white at the start of a trial. The occluders were filled with visual noise, though they also possessed a .058° black perimeter throughout the trial. Probes were small red .23° squares that appeared for 110 ms. Over the course of 80 trials (with a new set of object trajectories), participants observed an average of 329 probes, and each probe was equally likely to occur in one of the same nine probe conditions used in Experiment 1.

They key difference in this experiment was that the objects themselves were defined solely by their motion during most of each trial: each of four identical 1.74° squares was defined by visual noise by randomly setting each pixel to either black or white. At the beginning of the trial, a black line (.17°) surrounded each

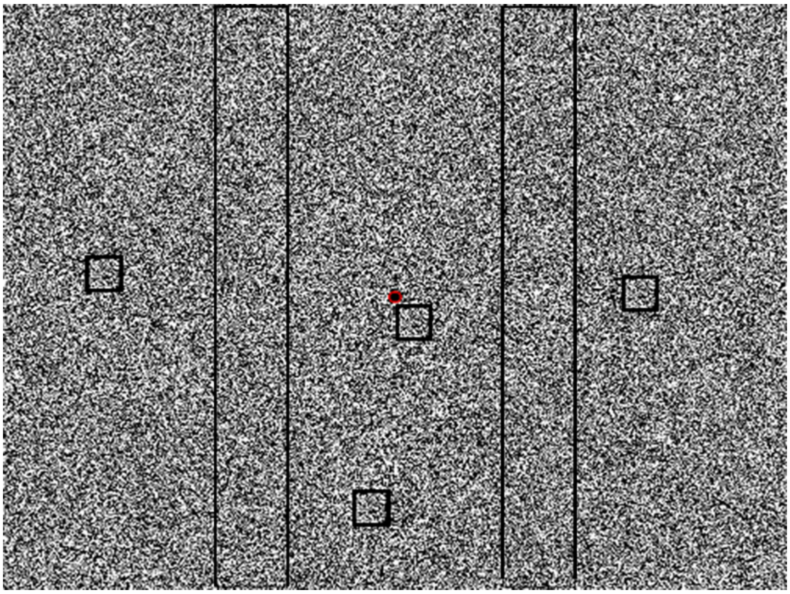

Fig. 5. A schematic depiction (not to scale) of the display at the start of a trial in Experiments 3 and 4. The background, occluders, and items were filled with a random pattern of black and white noise. One-pixel black lines identified the occluders' boundaries throughout the trial. At the very beginning of a trial, targets and distractors possessed a thin black boundary so that they could be located, but these borders disappeared during the tracking.

square so that the objects were visibly defined while they remained static. The two targets were then identified by alternating five times between completely black and completely white (108 ms for each alternation). They were then replaced with their random noise patterns – with the initial black lines now removed – and all objects began to move through the display as in Experiment 1. These regions of noise remained locally stable throughout their motion, so that they temporarily overwrote the background noise as they moved. This gave rise to salient motion-defined 'noisy' objects moving on a noisy background. (This display cannot be depicted in a static figure since the objects would be invisible, but the displays can be viewed online.) At the end of each trial, all objects immediately turned solid white while the participants selected the two targets.

4.2. Results and discussion

Pilot testing revealed that tracking motion-defined objects was possible, but more difficult than standard MOT. Since good tracking performance is critical in these experiments for the meaningful evaluation of target and distractor probe detection, we reduced the tracking load here from 3 among 6 to 2 among 4. As a result, tracking performance in this experiment was high, averaging 93%. Participants also made few false alarms, averaging only 15 over the course of 80 trials.

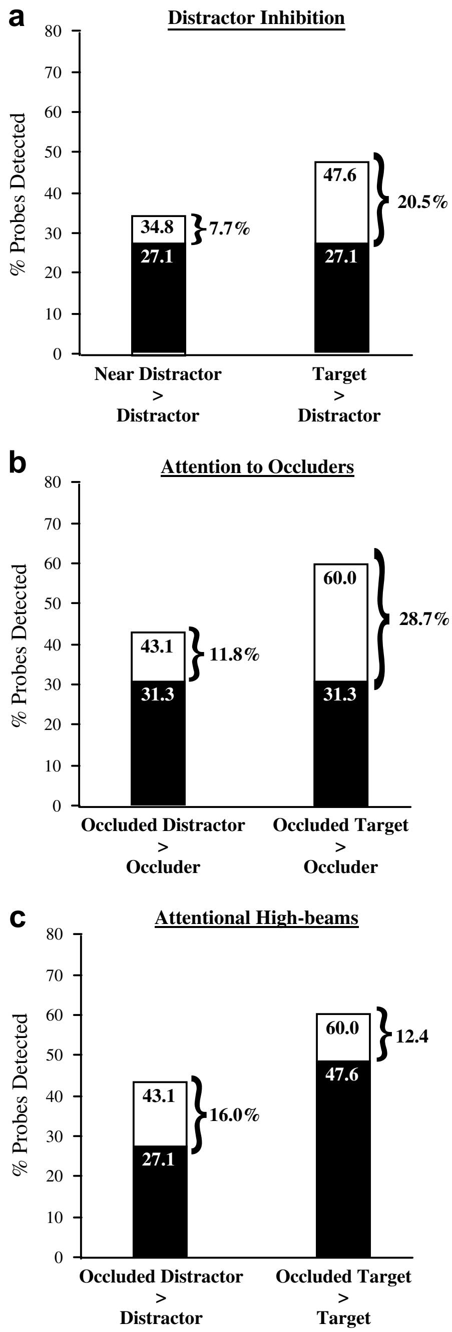

Fig. 6. The three primary effects from Experiment 3: (a) *Distractor inhibition*, wherein distractors are inhibited relative to both targets and empty space. (b) *Attention to occluders*, wherein probes on occluders are detected better when they are currently occluding either targets or distractors. (c) *Attentional high-beams*, wherein probes on both targets and distractors are detected better when occluded.

As depicted in Fig. 6 (with supporting data in Table 1), this experiment replicated each of the three main effects from Experiment 1. A repeated measures ANOVA revealed a main effect of probe condition on detection, $F(8, 11) = 23.73$, $p < .001$, $\eta^2 = 2.159$. Planned comparisons revealed evidence for distractor inhibition, attention to occluders, and attentional high-beams.

4.2.1. Distractor inhibition

Detection of Distractor probes was worse than both Target probes ($t(11) = 4.74$, $p < .001$, $d = 1.49$) and Near Distractor probes ($t(11) = 2.25$, $p = .041$, $d = .61$). As emphasized in Section 6, this evidence (depicted in Fig. 6a) is the strongest to date for distractor inhibition since confounds involving nearby object contour were removed.

4.2.2. Attention to occluders

This experiment again yielded evidence for object-based attention to currently invisible objects: as depicted in Fig. 6b, probe detection was significantly worse for Occluder probes relative to both Occluded Target probes ($t(11) = 5.32$, $p < .001$, $d = 2.45$) and Occluded Distractor probes ($t(11) = 2.75$, $p = .016$, $d = 1.12$).

4.2.3. Attentional high-beams

As depicted in Fig. 6c, the motion-defined objects in this experiment were subject to attentional high-beams: probe detection was again better for occluded items relative to unoccluded items, for both targets (Target vs. Occluded Target: $t(11) = 2.92$, $p = .011$, $d = .97$) and distractors (Distractor vs. Occluded Distractor: $t(11) = 4.58$, $p < .001$, $d = 1.63$).

The fact that these three effects all persisted with the motion-defined objects used here implies (1) that they are not artifacts of subtle differences in visibility between different conditions, and (2) that these effects generalize across both large superficial display differences, and also different degrees of attentional load.

5. Experiment 4: Motion-defined objects with background contours

Though the motion-defined objects used in Experiment 3 eliminate the possibility that the effects observed in this study were driven by differences in the visibility of the probes in various conditions, it remains possible that the results could be fueled by differences in the absolute *locations* of the probes. In Experiments 1 and 3, the occluded object probes (which were detected better than unoccluded object probes) always occupied certain narrowly confined regions of the display (viz. behind the narrow occluders). Might these regions be attentional ‘hot-spots’ for other reasons? If one could split attention to only two loci, for example, these might be good candidates since they are evenly distributed along the main axis of the display. This possibility seems unlikely to us for several reasons. First, we know of no previous evidence for such a distribution of attention. Second, many investigators interpret MOT in terms of attention being split among each of the objects themselves in the first place (e.g., Alvarez & Scholl, 2005; Cavanagh & Alvarez, 2005; Scholl, in press; Scholl et al., 2001). And third, this hypothesis could not account for the effect of attention to occluders in the first place, since these probe locations are largely equated. Nevertheless, in this experiment we ruled out the possibility of a purely spatial explanation for our results by replicating Experiment 3 with only one change: now the objects passed *over* the previously-occluding contours.

5.1. Methods

Except as noted here, this experiment was identical to Experiment 3. We tested a new group of 14 naïve participants. Data from two participants were not included in the final analysis because their overall levels of tracking performance (76% and 78%) were more than two standard deviations ($SD = 8\%$) below the group mean (96%). The occluders in this experiment – now defined only in terms of their thin black outlines – were functionally drawn in the background so that the objects simply passed over them. Thus the only occlusion present in this experiment was that the objects (moving squares of visual noise) occluded the thin black lines that outlined what in Experiment 3 had been 'occluders'.

5.2. Results and discussion

For simplicity, we continue to refer to the nine probe conditions by the same names as used in the previous two experiments. Thus Occluded Target, Occluded Distractor, Near Occluded Target, and Near Occluded Distractor probes all refer now to probes that appeared on or near fully visible items during the times that they moved completely over one of the two large vertical rectangles (that had previously served as occluders).

Tracking performance in this experiment was high (96%) and did not differ from that in Experiment 3 ($t(22) = 1.02$, $p = .32$, $d = .44$). This confirms again, in rather different circumstances, that occlusion has little impact on tracking performance in MOT (as initially reported in Scholl & Pylyshyn, 1999). As in the previous two experiments, participants made few false alarms, averaging only 9 over the course of 80 trials. A repeated measures ANOVA revealed a main effect of probe condition on detection ($F(8, 11) = 16.94$, $p < .0001$, $\eta^2 = 1.54$), and Table 1 lists mean detection rates and standard errors for each probe condition. Planned comparisons revealed evidence for distractor inhibition, but no evidence of attention to occluders, or of attentional high-beams:

5.2.1. Distractor inhibition

Because distractor inhibition refers to the inhibition of unoccluded distractors relative to unoccluded targets, we expected that drawing the "occluders" in the background in this experiment should make no difference. And, in fact, the distractor inhibition in this experiment was again robust, as depicted in Fig. 7a: detection of Distractor probes was worse than both Target probes ($t(11) = 7.07$, $p < .001$, $d = 1.41$) and Near Distractor probes ($t(11) = 3.26$, $p = .006$, $d = .71$).

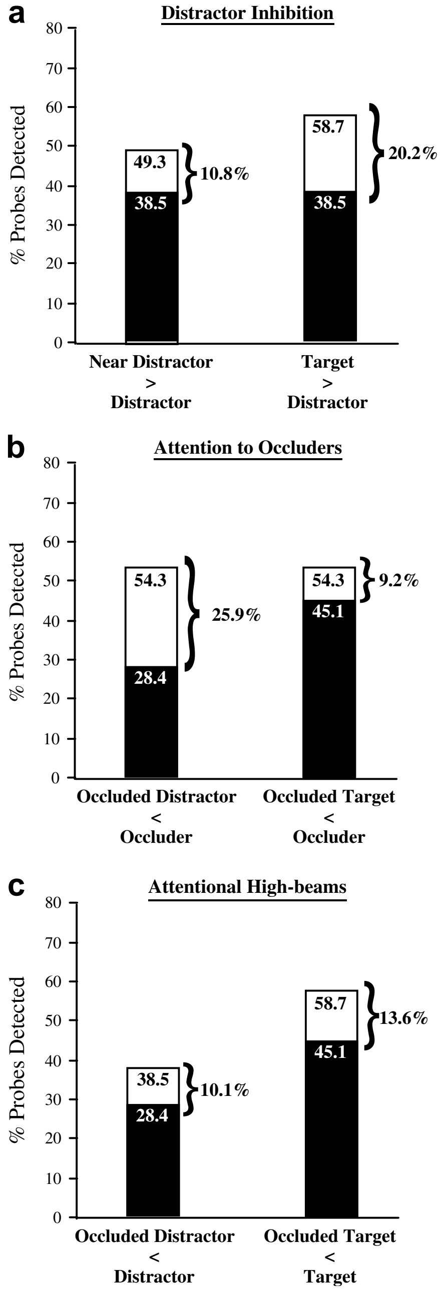

Fig. 7. The three primary comparisons in Experiment 4: (a) *Distractor inhibition*, wherein distractors are inhibited relative to both targets and empty space. (b) *No attention to occluders*: in contrast to Experiments 1 and 3, probes on occluders were detected better when the "occluders" did not intersect any targets or distractors. (c) *No attentional high-beams*: in contrast to Experiments 1 and 3, probes that appeared on targets and distractors were more easily detected on objects that did not currently intersect with an "occluder".

5.2.2. Attention to occluders

In contrast to Experiments 1 and 3, there was no evidence of preferential attention to "occluders" which overlapped objects. As depicted in Fig. 7b, there was no reliable difference between the detection of Occluded Target probes and Occluder probes ($t(11) = 1.56$, $p = .14$, $d = .49$). Moreover, we found significantly *worse* detection for Occluded Distractor probes compared to Occluder probes ($t(11) = 3.93$, $p = .002$, $d = 1.37$), indicating, as expected, that visible distractors continue to be inhibited as they pass over the "occluder" locations.

5.2.3. Attentional high-beams

Finally, in contrast to Experiments 1 and 3, we found no evidence in this experiment for attentional high-beams. In fact, detection of Occluded Target probes was significantly worse than detection of Target probes ($t(11) = 3.49$, $p = .004$, $d = .86$), and detection of Occluded Distractor probes was significantly worse than detection of Distractor probes ($t(11) = 5.08$, $p < .001$, $d = 58$).[6]

The disappearance of the attention-to-occluder and attentional high-beams effects in this experiment demonstrates that these effects depend on the presence of bona fide occlusion in the displays.

6. General discussion

The results presented here are in many ways a case study of one of the central lessons of cognitive science: that how effortless and natural a process feels is a poor guide to how much work is supporting that process beneath the surface. This lesson is most salient with processes such as recognizing a face or understanding speech – ubiquitous and seemingly effortless activities that nevertheless require sophisticated and specialized processing. The present study focused on an equally ubiquitous phenomenon of visual processing: the perception of objects as persisting through momentary occlusion. Like face or speech recognition, it is deceptively easy to assume that this situation does not pose any special difficulty at all. The perceived persistence of objects through occlusion seems to be something that just happens, without intent or noticeable effort. The central lesson of this study, however, is that additional processing resources – realized via the allocation of attention – are recruited in just these instances, presumably because occlusion presents a notable

[6] These two effects remain somewhat mysterious. One possible explanation for this finding involves the fact that our analysis of probes on objects not intersecting with an occluder includes those objects that occupied the most central portion of the display, in between the two occluders. If probes in this region are detected more easily than probes elsewhere in the display, this could account for the relative advantage for probes on 'unoccluded' objects compared to probes on 'occluded' ones. In Experiments 1 and 3, in contrast, the high-beams effect may have overwhelmed this effect, producing the opposite pattern of results; in this case, of course, the results of Experiments 1 and 3 would actually be underestimating the strength of attentional high-beams. In any case, however, it is clear that this experiment succeeded in ruling out the possibility that absolute spatial effects could explain attentional high-beams, since these effects are in the opposite direction from that predicted by the spatial hypothesis.

challenge to coherent visual perception. This phenomenon was most directly illustrated here via the attentional high-beams effect.

6.1. Attentional high-beams and object persistence

The initial studies of multiple object tracking through occlusion (Scholl & Pylyshyn, 1999) were designed to determine if such a thing was possible at all. The outcome seemed uncertain given the especially high load involved in MOT combined with the frequent disappearance of the items to be tracked. Nevertheless, the results of these experiments were clear: not only was tracking through occlusion possible, but it did not seem any more difficult than tracking without occlusion. Even when studies have found worse tracking performance with occlusion, the magnitude of this effect has been miniscule (Scholl & Feigenson, 2004). In retrospect, however, perhaps this result was not so surprising: after all, any tracking system that was unable to readily cope with occlusion would be of little worth in the real world. The same is true of the ability in MOT to operate in 3D spaces in addition to 2D displays (Viswanathan & Mingolla, 2002), and to discount observer motion (Liu et al., 2005).

Because of these results, however, it can seem dangerously easy to simply discount occlusion, assuming that any challenge it poses is dispatched much earlier in visual processing so that occluders essentially do not exist in higher-level object-based processing. The present experiments, however, indicate that this is not so. Even momentary instances of occlusion demand extra resources. This conclusion was fueled by two specific results:

First, we found that attention to occluders was amplified when those occluders obscured an object compared to when they did not. This effect was manifested here in terms of better probe detection on occluders when they occluded targets or distractors compared with occluders that currently occluded nothing (Experiments 1 and 3). Experiment 4 further demonstrated that this effect is due to occlusion per se, rather than to prioritized absolute spatial locations, since the effect disappeared (and actually reversed for distractors) when the “occluders” were drawn in the background. This result demonstrates maintained object-based attention to objects through occlusion, but it is consistent with the possibility, discussed earlier, that occluders are simply discounted in such processing.

The second effect we observed in this study, however, fueled the opposite conclusion: that occluders, far from being simply ignored, actually demand extra attention to those objects that they are currently occluding. This attentional high-beams effect was manifested in terms of better probe detection on occluded objects relative to unoccluded objects. This effect was replicated under different display conditions and tracking loads in Experiments 1 and 3, and was not due to psychophysical constraints on either probe visibility (since the effect occurred in Experiment 3 without any object contours at all, and disappeared in Experiment 2 when tracking was not required) or spatial locations (since the effect disappeared without functional occluders in Experiment 4). It could be, of course, that this additional attentional engagement during occlusion plays no role in visual processing. That seems unlikely, however, given (1) that MOT is already highly attentionally demanding, and (2) that the effect occurs so systematically, and despite the fact that the periods of occlusion in this study were fairly brief – on average only 1100 ms. A more exciting possibility, then, is that the additional momentary allocation of attention to occluded objects is somehow involved in computing what ends up feeling like relatively effortless tracking.

That this high-beams effect obtained for distractors (as well as for targets) attests to its importance. While distractors may be otherwise inhibited during tracking (as discussed in the next section), this inhibition is suddenly released during momentary occlusion. Thus the deployment of attentional high-beams to aid in tracking through occlusion even trumps the need to inhibit distractors. Similarly, the high-beams effect is not isolated to conditions where tracking is difficult: in Experiment 3, for example, it occurred even when tracking performance was well over 90%. Thus attentional high-beams seem to be recruited automatically, for all objects. In other words, the high-beams effect reflects the ever-present challenges of representing persisting objects through occlusion, rather than the unique challenges of the MOT paradigm.

6.2. Tracking distractors?

The fact that the 'occlusion advantage' for probe detection applied to occluded distractors as well as to occluded targets is important for establishing the nature and importance of attentional high-beams, but it is also important for our understanding of how distractors in MOT are processed. In the first place, of course, attentional high-beams with distractors demonstrates that they *are* processed – something not predicted in most theories of tracking. After all, attention cannot be amplified on an occluded distractor unless the visual system is sensitive to where that object is, and when it becomes occluded. In addition, however, our results also provide the most rigorous evidence to date for the existence of another type of distractor processing.

Distractor inhibition during MOT was first reported by Pylyshyn (2006) in terms of impaired probe detection on distractors relative to empty space. The difficulty with this conclusion, however, is that it potentially suffers from the same types of concerns that motivated our Experiment 2 and 3 with regard to the high-beams effect: it could be that with moving objects, probes are more readily detected without nearby contours. The initial studies controlled for this by having 'empty space' probes actually occur 1 object radius away from a distractor, but even this is problematic. Distractor probes are still entirely surrounded by proximal contours (1 radius away) in every direction, whereas empty-space probes have such a contour in only one direction. This simple contour-based explanation was then further disproved by Pylyshyn (2006), who demonstrated that distractor inhibition disappears when tracking is not required; that is, when participants only need to detect probes, without tracking, differences between moving items and the spaces near them disappear. We further controlled for such concerns in Experiments 3 and 4 in the present project since the objects did not have any visible contours. In other words, if you were to see static frames of Distractor probes and Near Distractor probes, you would be unable to distinguish them. Thus these experiments both provide especially careful tests of distractor inhibition during MOT, and also demonstrate that it transfers to different display types and tracking loads.

At the same time, our results complicate the notion of distractor inhibition in a different way. A careful examination of Table 1 reveals that similar effects seem to have occurred for the targets in some conditions as well (see footnote 3): Near-Target probes were sometimes detected more often than Target probes. This suggests that the explanation for these phenomena might not be a form of *inhibition* after all, since it is unclear how or why targets would be inhibited. We have nevertheless continued to use the term 'distractor inhibition' throughout the paper because we assume that they reflect the same real effects that have been described using this term by Pylyshyn (2006) (see also Pylyshyn et al., submitted for publication). Pylyshyn's studies, however, make the inhibitory role clearer: in particular, these studies did not find impairments for targets relative to empty space. For additional discussion of such effects, see Flombaum and Scholl (submitted for publication).

In some ways, the 'distractor inhibition' in this study may complement previous studies using other paradigms of maintained inhibition through occlusion. Inhibition of return, for example, can continue to apply to a stimulus that experiences a brief period of occlusion (Jeffries, Wright, & Di Lollo, 2005; Yi, Kim, & Chun, 2003), and visually 'marked' distractors remain marked even after a moment of occlusion (Kunar, Humphreys, Watson, & Smith, 2003). These previous studies demonstrated the maintenance of inhibition through occlusion, but the present study demonstrates that inhibition does not always simply discount occlusion. Rather, in this study, distractor inhibition was momentarily put 'on hold' during occlusion itself, via the operation of object-based attention and attentional high-beams.

6.3. Conclusions: A general principle of design

It is perhaps a general principle of design that the average amount of resources required to tackle most situations is considerably less than the amount of resources required to tackle the most difficult or exigent situations. As a result, simply building in the average amount of resources will not suffice and the system will fail in difficult circumstances. Building in the resources needed to cope with the most difficult situations, in contrast, is wasteful; those resources are simply not necessary all the time. A solution is thus to have two separate mechanisms of resource allocation: a 'standing' pool of resources and a reserve that can be brought online in especially demanding situations. A concrete example of this tradeoff in action is the headlights on a car. The amount of light they produce is sufficient for most driving conditions, but in especially difficult terrain or weather, you must turn your high-beams on to achieve the same degree of safety and efficiency. The result is an overall boost in the visibility of everything in view – both 'targets' (e.g., road signs) and 'distractors' (e.g., roadside trees). This example is perhaps especially apt given that attention has long been analogized to a spotlight (e.g., Cave & Bichot, 1999). Here we have suggested that the 'spotlight' of attention may not always provide uniform facilitation, but that attentional high-beams may also be brought online in especially challenging situations such as tracking through occlusion.

References

Alvarez, G. A., & Scholl, B. J. (2005). How does attention select and track spatially extended objects? New effects of attentional concentration and amplification. *Journal of Experimental Psychology: General, 134*, 461–476.

Assad, J. A., & Maunsell, J. (1995). Neural correlates of inferred motion in primate posterior parietal cortex. *Nature, 373*, 518–521.

Behrmann, M., Zemel, R. S., & Mozer, M. C. (1998). Object-based attention and occlusion: Evidence from normal participants and computational model. *Journal of Experimental Psychology: Human Perception & Performance, 24*, 1011–1036.

Burke, L. (1952). On the tunnel effect. *Quarterly Journal of Experimental Psychology, 4*, 121–138.

Burt, P., & Sperling, G. (1981). Time, distance, and feature trade-offs in visual apparent motion. *Psychological Review, 88*, 171–195.

Cavanagh, P., & Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends in Cognitive Sciences, 9*, 349–354.

Cave, K. R., & Bichot, N. P. (1999). Visuospatial attention: Beyond a spotlight model. *Psychonomic Bulletin and Review, 6*, 204–223.

Cepeda, N. J., Cave, K. R., Bichot, N. P., & Kim, M.-S. (1998). Spatial selection via feature-driven inhibition of distractor locations. *Perception & Psychophysics, 60*, 727–746.

Comtois, R. (2007). *VisionShell PPC* [Software libraries]. Cambridge, MA: Author.

Davis, G., & Driver, J. (1997). Spreading of visual attention across modally versus amodally completed surfaces. *Psychological Science, 8*, 275–281.

Dawson, M. (1991). The how and why of what went where in apparent motion: Modeling solutions to the motion correspondence problem. *Psychological Review, 98*, 569–603.

Doran, M. M., Hoffman, J. E., & Scholl, B. J. (2006). The role of eye fixations in amplification and concentration effects during MOT [Abstract]. *Abstracts of the Psychonomic Society, 11*, 134.

Downing, C., & Pinker, S. (1985). The spatial structure of visual attention. In M. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 171–187). Hillsdale, NJ: Erlbaum.

Fazl, A., & Mingolla, E. (2007). Consistency of eye movements during multiple object tracking. *Poster presented at the annual meeting of the Vision Sciences Society*, 5/5/07, Sarasota, FL.

Fehd, H. M., & Seiffert, A. E. (in press). Eye-movements during multiple object tracking: Where do participants look? *Cognition*.

Flombaum, J. I., Kundey, S., Santos, L. R., & Scholl, B. J. (2004). Dynamic object individuation in rhesus macaques: A study of the tunnel effect. *Psychological Science, 15*, 795–800.

Flombaum, J. I., & Scholl, B. J. (2006). A temporal same-object advantage in the tunnel effect: Facilitated change detection for persisting objects. *Journal of Experimental Psychology: Human Perception & Performance, 32*, 840–853.

Flombaum, J. I., & Scholl, B. J. (2007). Attending to moving vs. static stimuli: A surprising dissociation in multiple object tracking. *Poster presented at the annual meeting of the Vision Sciences Society*, 5/5/07, Sarasota, FL.

Flombaum, J. I., & Scholl, B. J. (submitted for publication). The influence of motion on attention..

Flombaum, J. I., Scholl, B. J., & Pylyshyn, Z. W. (2006). Attentional high-beams in tracking through occlusion [Abstract]. *Journal of Vision, 6*, 765a. Available from http://journalofvision.org/6/6/765/.

Haimson, C., & Behrmann, M. (2001). Cued visual attention does not distinguish between occluded and occluding objects. *Psychonomic Bulletin & Review, 8*, 496–503.

He, Z. J., & Nakayama, K. (1992). Surfaces versus features in visual search. *Nature, 359*, 231–233.

He, Z. J., & Nakayama, K. (1994). Apparent motion determined by surface layout not by disparity or 3-dimensional distance. *Nature, 367*, 163–165.

Horowitz, T. S., Birnkrant, R. S., Fencsik, D. E., Tran, L., & Wolfe, J. M. (2006). How do we track invisible objects? *Psychonomic Bulletin & Review, 8*, 516–523.

Humphreys, G. W., Stalmann, B. J., & Olivers, C. (2004). An analysis of the time course of attention in preview search. *Perception & Psychophysics, 66*, 713–730.

Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology, 43*, 171–216.

Jeffries, L. N., Wright, R. D., & Di Lollo, V. (2005). Inhibition of return to an occluded object depends on expectation. *Journal of Experimental Psychology: Human Perception & Performance, 31*, 1224–1233.

Johnson, S. P., Amso, D., & Slemmer, J. A. (2003). Development of object concepts in infancy: Evidence for early learning in an eye tracking paradigm. *Proceedings of the National Academy of Sciences of the United States of America, 100*, 10568–10573.

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

Kaufman, J., Csibra, G., & Johnson, M. H. (2005). Oscillatory activity in the infant brain reflects object maintenance. *Proceedings of the National Academy of Sciences of the United States of America, 102*, 15271–15274.

Kunar, M. A., Humphreys, G. W., Watson, D. G., & Smith, K. J. (2003). When a reappearance is old news: Visual marking survives occlusion. *Journal of Experimental Psychology: Human Perception & Performance, 29*, 185–198.

Liu, G., Austen, E. L., Booth, K. S., Fisher, B. D., Argue, R., Rempel, M. I., & Enns, J. T. (2005). Multiple-object tracking is based on scene, not retinal, coordinates. *Journal of Experimental Psychology: Human Perception & Performance, 31*, 235–247.

Marr, D. (1982). *Vision*. New York: W.H. Freeman.

Michotte, A., Thinès, G., & Crabbé, G. (1964/1991). Les complements amodaux des structures perceptives. In *Studia Psychologica*. Louvain: Publications Universitaires. Reprinted and translated as: Michotte, A., Thinès, G., & Crabbé, G. (1991). Amodal completion of perceptual structures. In G. Thines, A. Costall, & G. Butterworth (Eds.), *Michotte's experimental phenomenology of perception* (pp. 140–167). Hillsdale, NJ: Erlbaum.

Moore, C. M., & Fulton, C. (2005). The spread of attention to hidden portions of occluded surfaces. *Psychonomic Bulletin & Review, 12*, 301–306.

Moore, C. M., Yantis, S., & Vaughan, B. (1998). Object-based visual selection: Evidence from perceptual completion. *Psychological Science, 9*, 104–110.

Nakayama, K., He, Z. J., & Shimojo, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In S. M. Kosslyn & D. N. Osherson (Eds.), *An invitation to cognitive science (2nd ed.): Visual cognition* (pp. 1–70). Cambridge, MA: MIT Press.

Olson, I. R., Gatenby, J. C., Leung, H., Skudlarski, P., & Gore, J. C. (2003). Neuronal representation of occluded objects in the human brain. *Neuropsychologia, 42*, 95–104.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*, 2–25.

Pylyshyn, Z. W. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial index model. *Cognition, 32*, 65–97.

Pylyshyn, Z. W. (1994). Some primitive mechanisms of spatial attention. *Cognition, 50*, 363–384.

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition, 80*, 127–158.

Pylyshyn, Z. W. (2003). *Seeing and visualizing: It's not what you think*. Cambridge, MA: MIT Press.

Pylyshyn, Z. W. (2006). Some puzzling findings in multiple object tracking (MOT): II. Inhibition of moving nontargets. *Visual Cognition, 14*, 175–198.

Pylyshyn, Z. W. (2007). *Things and places: How the mind connects with the perceptual world (2004 Jean Nicod Lectures)*. Cambridge, MA: MIT Press.

Pylyshyn, Z. W., Haladjian, H. H., King, C. E., & Reilly, J. E. (submitted for publication). Selective nontarget inhibition in multiple object tracking (MOT).

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

Scholl, B. J. (2001). Spatiotemporal priority and object identity. *Cahiers de Psychologie Cognitive, 20*(5), 359–371.

Scholl, B. J. (in press). What have we learned about attention from multiple-object tracking (and vice versa)? In D. Dedrick & L. M. Trick (Eds.), *Computation, cognition and pylyshyn*. Cambridge, MA: MIT Press.

Scholl, B. J., & Feigenson, L. (2004). When out of sight is out of mind: Perceiving object persistence through occlusion vs. implosion [Abstract]. *Journal of Vision, 4*, 26a. Available from http://www.journalofvision.org/4/8/26/.

Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology, 38*, 259–290.

Scholl, B. J., Pylyshyn, Z. W., & Feldman, J. (2001). What is a visual object? Evidence from target merging in multiple-object tracking. *Cognition, 80*, 159–177.

Sears, C. R., & Pylyshyn, Z. W. (2000). Multiple object tracking and attentional processing. *Canadian Journal of Experimental Psychology, 54*, 1–14.

vanMarle, K., & Scholl, B. J. (2003). Attentive tracking of objects vs substances. *Psychological Science, 14*, 498–504.

Viswanathan, L., & Mingolla, E. (2002). Dynamics of attention in depth: Evidence from multi-element tracking. *Perception, 31*, 1415–1437.

Watson, D. G., & Humphreys, G. W. (2000). Visual marking: Evidence for inhibition using a probe-dot detection paradigm. *Perception & Psychophysics, 62*, 471–481.

Yantis, S. (1992). Multi-element visual tracking: Attention and perceptual organization. *Cognitive Psychology, 24*, 295–340.

Yi, D.-J., Kim, M.-S., & Chun, M. (2003). Inhibition of return to occluded objects. *Perception & Psychophysics, 65*, 1222–1230.

Zelinsky, G., Neider, M., & Todor, A., (2007). Multi-object tracking in a realistic 3-D environment. *Poster presented at the annual meeting of the Vision Sciences Society*, 5/5/2007, Sarasota, FL.