

How do we track invisible objects?

Todd S. Horowitz^{ab}, Randall S. Birnkrant^a, David E. Fencsik^{ab}, Linda Tran^c, Jeremy M. Wolfe^{ab}

^aBrigham and Women's Hospital; ^bHarvard Medical School; ^cCharlestown High School

Abstract

We previously demonstrated that observers in multiple object tracking experiments can successfully track targets when all objects simultaneously vanish for periods lasting several hundred ms (Alvarez, Horowitz, Arsenio, Dimase, & Wolfe, 2005). How do observers do this? Since observers can track objects that move behind occluders (e.g. Scholl & Pylyshyn, 1999), they might treat a temporal gap as a case of complete occlusion. If so, performance should improve if occlusion cues (deletion and accretion) are provided and items disappeared and reappeared one by one (asynchronously) rather than simultaneously. However, we found better performance with simultaneous than with asynchronous disappearance (Experiment 1), while occlusion cues were detrimental (Experiment 2). We propose that observers tolerate a gap in tracking by storing the current task state when objects vanish and resuming tracking based on that memory when objects reappear (a task-switching account).

Corresponding author: Todd S. Horowitz, Brigham and Women's Hospital, Visual Attention Laboratory, 64 Sidney Street, Suite 170, Cambridge, MA 02139. Tel: 617-768-8813; Fax: 617-768-8816. Email: toddh@search.bwh.harvard.edu

Introduction

While research in visual attention often involves static stimuli or abrupt changes, the environment we live in features changes that evolve smoothly over time. The multiple object tracking (MOT) task, devised by Pylyshyn & Storm (1988) is a tool for studying attention under these circumstances. From a display of identical objects, a subset are cued as targets. All objects move independently for several seconds, then observers classify objects as targets or non-targets. The standard finding is that humans are able to track 3-5 objects (Bahrami, 2003).

While MOT is a demanding task thought to require sustained attention, we have recently demonstrated that observers can maintain high tracking performance while simultaneously performing a difficult visual search through a briefly presented array (Alvarez et al., 2005). Dual-task costs were small, comparable to those obtained when MOT was paired with an auditory tone-monitoring task. While this finding might be taken to demonstrate two independent resources, specialized respectively for MOT and visual search (Leonard & Pylyshyn, 2003), we argued that it is more likely that observers are able to remove a single attentional mechanism from the tracking task, search for the target, then return attention to MOT.

Several experiments supported this "switching hypothesis". We devised a search task where the target could appear at any time, so observers had to search continuously throughout a trial. If observers were performing the two tasks with independent resources, then we would have expected to see no change in the reaction time \times set size function. A change in slope would have indicated that search efficiency had been reduced, while a change in intercept would be consistent with a single resource swapping between tasks. Adding a concurrent MOT task increased intercept, but not slope, consistent with the switching hypothesis.

This switching hypothesis implies that observers could remove attention from MOT, do something else, then successfully recover targets. If so, it should be possible to simply remove the stimuli for some time without disrupting tracking. Accordingly, we had observers track objects across a gap when all the stimuli disappeared. The computer would continue to update object locations during this gap so that they reappeared at predictable new locations. Surprisingly, observers performed rather well on this task, even at gap durations of 300-500 ms. Similar results involving longer durations have been obtained by Keane & Pylyshyn (submitted). Adding a difficult visual search task during the gap did not reduce tracking performance.

Successful tracking under these circumstances cannot be attributed to apparent motion. The temporal parameters of tracking over a gap would place it in the realm of "long-range" or "second-order" motion (Bex & Baker, 1999; Petersik, 1989). This type of apparent motion must be constructed by attention, probably one object at a time (Ashida, Seiffert, & Osaka, 2001; Horowitz & Treisman, 1994). It is unlikely that long-range apparent motion could survive the deployment of attention to an intervening search task as demonstrated in Alvarez et al. (2005). Therefore, in order to explain tracking of invisible objects, we need to appeal to some other mechanism.

The attention-switching explanation argues that simultaneous disappearance of all objects serves as a cue to store information for subsequent target recovery, thereby freeing attention. Therefore, the track across the gap paradigm can be regarded as a task-switching design without a secondary task. Since targets are not actively tracked during the gap under this hypothesis, we term it the *offline* hypothesis.

The alternative is that tracking across a gap is not qualitatively different from tracking continuously visible objects. According to this *online* hypothesis, the visual system treats the disappearance of stimuli as the simultaneous occlusion of all objects, and uses amodal perception mechanisms to continue to track the invisible objects.

The experiments presented here sought to distinguish between these two hypotheses by manipulating occlusion cues and the synchrony of object disappearance.

In standard MOT experiments, the visual system can take advantage of amodal perception during occlusion events (Scholl & Pylyshyn, 1999). The online hypothesis assumes that the visual system treats the instantaneous disappearance of all items as a massive occlusion. Scholl (Scholl & Feigenson, 2004; Scholl & Pylyshyn, 1999) has shown that MOT is better when occlusion is signaled by gradual deletion of items and reappearance is accompanied by gradual accretion, in comparison to simple disappearance and reappearance. Therefore, the online hypothesis predicts that adding such cues should improve performance in our tracking across the gap paradigm. Under the offline hypothesis, occlusion cues should provide no benefit; in fact, they might

actually be counterproductive, since the offline store might not be invoked by objects that are merely seen as passing behind something else.

We also reasoned that, when tracking online, it should be easier to track objects which disappear separately (asynchronously, as in Scholl & Pylyshyn, 1999), rather than simultaneously. All else being equal, simultaneous disappearance makes the correspondence problem more difficult to solve. According to the offline hypothesis, however, all else is not equal. Different processes would compensate for simultaneous and asynchronous disappearance. Offline processes would preserve target information in the simultaneous condition, while online processes would track items behind virtual occluders in the asynchronous condition. This should lead to an interaction between synchrony and the presence of occlusion cues. In the simultaneous condition, occlusion cues should be, at best, useless for the offline process. In the asynchronous condition, occlusion cues should assist the online tracking process.

To summarize, the online hypothesis predicts superior tracking with occlusion cues and asynchronous disappearance. The offline hypothesis predicts an interaction where occlusion cues should only be useful with asynchronous disappearance. With simultaneous disappearance, occlusion cues might even be harmful. In Experiment 1, we varied both factors and found no benefit for occlusion cues, and a disadvantage for asynchronous presentation, with a trend towards the interaction predicted by the offline hypothesis. In Experiment 2, we used only simultaneous disappearance but increased tracking load in order to increase the chance of finding an occlusion advantage; instead, we found that occlusion cues actually impaired performance. On balance, these results support the offline storage hypothesis over an online tracking account.

Experiment 1: Simultaneous vs. asynchronous disappearance

In Experiment 1 we manipulated three variables: synchronicity of object disappearance and reappearance, presence or absence of accretion and deletion cues, and gap duration. Figure 1 illustrates the four stimulus conditions. Objects could disappear suddenly (*vanish* conditions, Figure 1b,d), or be deleted gradually along

their leading contours (*occlusion* conditions, Figure 1c,e). Objects could disappear *simultaneously* (Figure 1b,c) or *asynchronously* (Fig. 1d,e). In *asynchronous* conditions, the gap duration for *each* object was exactly the same as the gap duration for *all* objects in the *simultaneous* conditions. Each condition was tested at three gap durations (107, 307, or 507 ms), for a total of twelve cells.

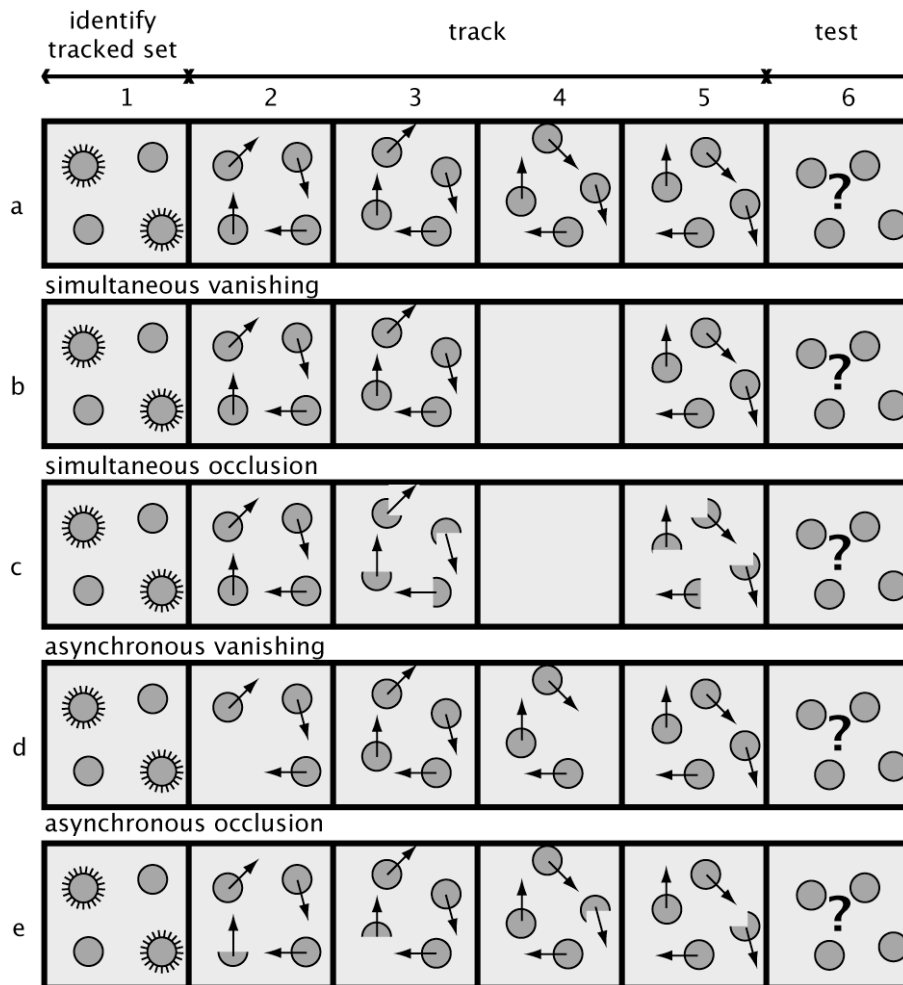


Figure 1. Schematic depiction of the conditions in Experiment 1. Row *a* illustrates the standard MOT paradigm used in training. Target disks are highlighted in panel 1 (in the actual experiment, this was done by blinking the targets on and off). Disks move independently in panels 2-5 (arrows indicate direction of motion and were not present in the actual display). This phase of the trial lasted 5000 ms. In panel 6, disks stop moving and the observer has to click on all target disks. Row *b* illustrates the *simultaneous vanishing* condition (corresponding to the tracking across the gap paradigm of Alvarez et al., 2005). The display is identical to row *a* except that all of the objects simultaneously vanish and continue to move while invisible (panel 4). After a fixed gap duration (107, 307, or 507 ms), all objects simultaneously reappear and continue to move (panel 5). Row *c* illustrates the *simultaneous occlusion* condition. In panel 3, each moving disk begins to move behind its own private, invisible occluder. Invisible motion continues in panel 4 as above. Then in panel 5, each disk moves out from behind its occluder. Row *d* illustrates the *asynchronous vanishing* condition. In panel 2, one disk abruptly disappears, moves invisibly for a fixed gap duration as above, then reappears in panel 3. Another disk vanishes in panel 4, then reappears after a fixed gap duration in panel 5. This process would be repeated for all disks, and the time at which each disk disappeared would be selected randomly. Row *e* illustrates the *asynchronous occlusion* condition. Each target disappears separately and reappears after a fixed gap duration, as in row *d*. However, each disappearance is preceded by the disk moving behind a private, invisible occluder (panels 2 & 4) as in row *c*, and each disk reappears by moving out from behind its occluder (panels 3 & 5)

Method

Observers

Eighteen observers were recruited from the Brigham and Women's Hospital Visual Attention Laboratory volunteer pool. Observers' ages ranged from 18 to 54 years (mean = 28.1, s.d. = 9.7). All observers had visual acuity of at least 20/25 when corrected, passed the Ishihara color screen, gave IRB-approved informed consent, and were compensated \$10/hr for participation.

Apparatus & stimuli

Stimuli were presented on 21" color CRT monitors controlled by PowerMacintosh G4 computers (Mac OS 9.2.2) running Matlab 5.2.1 using Psychophysics Toolbox routines (Brainard, 1997). Monitor resolution was 1024 × 768 pixels (36.1° × 27.1° at 57 cm viewing distance) at a 75 Hz refresh rate.

Tracking stimuli were dark gray disks with a black border on a mid-gray background. Disks subtended 1.8° in diameter; the border was 0.1° wide. Disks were initially positioned randomly on a 5 × 6 grid separated by at least 3.5° (center to center; 1.7° edge to edge). Disks moved at 6°/s (or 0.08°/refresh), changing direction only when bouncing off the 2.8° buffer zone (1.9° edge to edge) at the display edges or the 2.5° zone (0.7° edge to edge) around other disks.

Procedure

Observers were first trained on standard MOT without a gap. Training trials began with presentation of ten disks. Four disks blinked four times at 1 Hz. All disks then moved for 5000 ms. When the disks stopped, observers were asked to identify all targets by clicking on four disks with the mouse. Disks under the cursor were highlighted with a light gray square. The selected disk turned green after correct responses, red after incorrect responses. Missed disks blinked yellow. Observers received fifty training trials. Observers who did not correctly identify all targets on at least 80% of training trials were given a second session. If second session performance did not reach 80% they were excused.

The main experiment introduced the gap. There were three gap durations, 107, 307, and 507 ms, presented in separate blocks of trials. In the simultaneous vanishing condition, the gap began at a randomly selected time > 2000 ms after the start of the trial and > 1000 ms before the end of the trial. At the start of the gap, disks vanished from the screen. They reappeared after the gap as though they had continued to move. In the simultaneous occlusion condition, invisible occluders provided deletion cues for 200 ms prior to the gap and accretion cues for 293 ms after reappearance (see Figure 2).




Figure 2. Occlusion and disocclusion sequence. In both simultaneous and asynchronous occlusion conditions, stimuli disappeared gradually by moving behind invisible occluders (square patches the same size as the disks and the same color as the background), then gradually appeared from behind similar occluders after the gap. The gap duration, defined as the time of total invisibility, was the same as for the vanishing conditions. Here we present a sampling of frames from the deletion and accretion phases. Spatial relationships are to scale, but we are only presenting a few frames from each phase. (a) Disk one frame (13.3 ms) prior to occlusion. (b) Deletion phase. Disk gradually moves behind fixed invisible occluder over 200 ms (15 frames). (c) Gap. Disk is invisible for 107, 307, or 507 ms (8, 23, or 38 frames); scale of the figure assumes a 307 ms gap. (d) Accretion phase. Disk gradually moves out from behind fixed invisible occluder over 293 ms (22 frames). (e) Disk is fully visible.

Occluders initially appeared slightly over their corresponding disks during deletion, and disappeared before the disks had completely cleared the occluders. We observed in pilot work that the perceptual smoothness of occlusion and disocclusion appeared to differ, so the duration of the disocclusion phase was extended slightly to compensate.

In asynchronous conditions, each disk was assigned a random disappearance time following the same constraints as the simultaneous conditions. Disks reappeared after the appropriate gap duration. Asynchronous conditions were otherwise identical to simultaneous conditions.

Each block consisted of 10 practice trials and 40 experimental trials. Disappearance type was counterbalanced across observers. The order of combinations of gap duration and synchronicity was randomized for each observer. The experiment required 3 - 4 hours to complete.

Data analysis

The dependent variable was number of items tracked, corrected for guessing using a formula adapted from Scholl, Pylyshyn, & Feldman (2001). Data were submitted to a 2 (disappearance type: occlusion vs. vanishing) \times 2 (synchronicity: simultaneous vs. asynchronous) \times 3 (gap duration: 107, 307, or 507 ms) repeated-measures analysis of variance (ANOVA), using SPSS 11 for MacOS X. We applied Mauchly's Test of Sphericity, and applied the Huynh-Feldt correction when violations were observed. We report effect sizes using partial eta-squared ($\hat{\eta}^2$).

Results

Two observers failed to reach 80% performance after two training sessions. Mean training performance for the remaining sixteen observers was 87.3% (s.e.m. 1.2 %), or 3.2 items.

Results are shown in Figure 3. Performance declined with gap duration ($F(1.5, 22.2) = 152.8, p < .00001, \hat{\eta}^2 = .91$). There was a decrement for asynchronous disappearance relative to simultaneous ($F(1, 15) = 16.2, p < .005, \hat{\eta}^2 = .52$), but no effect of disappearance type ($F(1, 15) < 1, \hat{\eta}^2 = .01$). Performance declined more steeply with gap

duration for asynchronous conditions ($F(2, 30) = 6.1, p < .05, \hat{\eta}^2 = .29$). The disappearance type \times synchronicity interaction was not significant ($F(1, 15) = 2.6, p = 0.13, \hat{\eta}^2 = .15$), though the three-way interaction was marginal ($F(2, 30) = 2.8, p = .08, \hat{\eta}^2 = .16$). With longer gap durations, the data more closely approximated the predictions of the offline model.




Figure 3. Number of targets tracked as a function of gap duration, synchronicity, and disappearance type. Data from the simultaneous conditions are plotted in the left-hand panel, data from the asynchronous conditions in the right-hand panel. Squares denote the vanishing conditions, triangles the occlusion conditions. Error bars indicate the standard error of the mean

Discussion

These data do not support the online hypothesis. Adding occlusion cues provided no advantage. Furthermore, when disks disappeared asynchronously instead of simultaneously, performance declined, rather than improving, as would be expected if a single mechanism were handling both conditions. The data are more consistent with the offline hypothesis. With asynchronous disappearance, it is better to have occlusion cues (at least for the longer durations), but with simultaneous disappearance, occlusion cues are of no benefit.

Experiment 2: Varying tracking load

The lack of an advantage for occlusion cues favors the offline hypotheses. However, it remains a null effect; perhaps Experiment 1 lacked sensitivity. This argument is bolstered by the fact that the occlusion benefit did not reach statistical significance even in the asynchronous conditions, where both hypotheses predicted it.

Perhaps occlusion cues only help when observers are operating at capacity and we failed to place observers under sufficient load. In Experiment 2, we used only simultaneous disappearance, since the two hypotheses make different predictions for this case. We then varied the number of targets to track, in hopes of seeing a difference between occlusion and disappearance conditions at high load.

Method

Observers

Eighteen observers were recruited from the Brigham and Women's Hospital Visual Attention Laboratory volunteer pool. Observers' ages ranged from 19 to 50 years (mean = 28.6, s.d. = 7.5).

Apparatus & stimuli

Apparatus and stimuli were identical to Experiment 1.

Procedure

We used only simultaneous disappearance, and varied the number of targets from four to six in blocks. Since the number of non-targets must be at least equal to the number of targets (otherwise observers might employ a strategy of tracking non-targets instead of targets), we used twelve total disks. The longest gap duration was 520 ms. Disappearance type was counterbalanced across observers. The order of combinations of gap duration and target number was randomized for each observer. The experiment took roughly six hours to complete.

Results

Four observers failed to reach the 80% criterion after two training sessions. Two other observers failed to complete the experiment. Mean training performance for the remaining twelve observers was 87.8% (1.8%)¹.

Results are plotted in Figure 4. A 2 (disappearance type) \times 3 (number of targets: 4, 5, or 6) \times 3 (gap duration) ANOVA showed that performance declined with gap duration ($F(2, 22) = 197.1, p < .00001, \hat{\eta}^2 = .95$), but increased with number of targets ($F(1.2, 13.5) = 8.5, p < .01, \hat{\eta}^2 = .44$). The target number effect was greatest at long gap durations ($F(4,44) = 10.9, p < .005, \hat{\eta}^2 = .50$). There was a substantial disadvantage for occlusion cues ($F(1, 11) = 18.6, p < .005, \hat{\eta}^2 = .63$), which increased with number of targets ($F(2,22) = 5.0, p < .05, \hat{\eta}^2 = .31$).

¹ One observer ran 50 experimental trials instead of 40. Removing that observer did not change the outcome of the analyses.

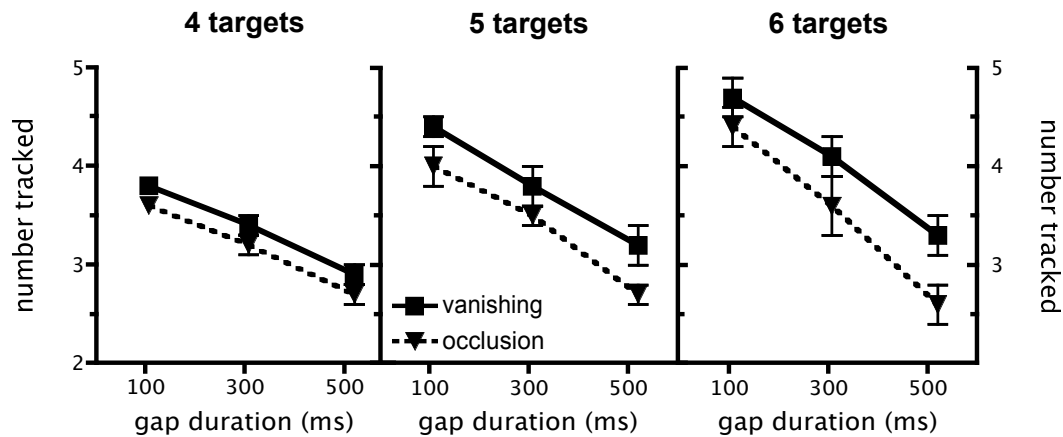


Figure 4. Number of targets tracked as a function of gap duration and disappearance type. Number of targets increases from four to five to six as panels move from left to right. Squares denote the vanishing conditions, triangles the occlusion conditions. Error bars indicate the standard error of the mean.

Discussion

Increasing tracking load in the simultaneous disappearance condition, unmasked a disappearance type effect in the direction opposite to that typically seen with asynchronous disappearance. When all items disappeared simultaneously, occlusion cues impaired performance. Under asynchronous conditions, occlusion cues usually improve tracking performance (Scholl & Feigenson, 2004; Scholl & Pylyshyn, 1999 see also Experiment 1). This is further evidence that simultaneous disappearance invokes a different mechanism from normal tracking.

In contrast to Experiment 1, this occlusion deficit was significant when only four targets ($F(1,11) = 18.9, p < .0005, \hat{\eta}^2 = .63$) were tracked. Apparently the increase in overall display size from 8 to 12 disks was sufficient to unmask the occlusion deficit (Pylyshyn, 1994).

General Discussion

These experiments provide further replication of our findings that observers can successfully track multiple moving objects that become invisible for several hundred ms (Alvarez et al., 2005; Keane & Pylyshyn, submitted; see also Yin & Thornton, 1999). Intuitively, this is rather surprising. The phenomenal experience of MOT is of continuous attentional effort by the observer.

It does not seem as if one could let go of the task for 500 ms and successfully recover tracked items.

How can we explain this remarkable ability? Here we contrasted two possible explanations. The online tracking hypothesis argues that the disappearance of the disks is treated as a degraded occlusion stimulus. The offline hypothesis assumes that simultaneous disappearance of the disks triggers storage of information necessary to recover targets. On the online view, adding occlusion cues and desynchronizing object disappearances should have improved performance. In fact, we observed the opposite. In Experiment 1, performance was worse in the asynchronous condition. In Experiment 2, performance was worse when accretion/deletion cues were present.

Of course, our occluders were invisible, appeared suddenly, and the only plausible object that could produce such occlusion events would be an elongated rectangle along each disk's trajectory. However, Scholl and Pylyshyn (1999) convincingly demonstrated that strictly local information is used to handle occlusion during MOT. In their experiments, each object had an invisible "virtual private occluder", which only occluded one particular object. There was no evidence for the existence of these occluders until the

moment of occlusion. Yet tracking under these conditions was not noticeably impaired. So it is unlikely that the peculiarities of our occluders played a role.

We tested a specific online hypothesis, based on the work of Scholl & Pylyshyn (1999). Yantis (1992) has offered an alternative view, that MOT performance is mediated by chunking targets into a single perceptual object via grouping. A purely online tracking mechanism, based on Yantis' views, might explain our results as follows. As occluders deleted different segments of the disks, they reduced perceptual similarity between targets and thus interfered with grouping. Similarly, asynchrony may also interfere with perceptual grouping (Palmer, 2002). We consider this explanation unlikely for two reasons.

First, a disk moving behind an invisible occluder is still perceived as a disk, rather than as an object changing shape and shrinking as it moves (Michotte, Thines, & Crabbe, 1964). It is this amodal representation that grouping (Palmer, Neff, & Beck, 1996), motion correspondence (He & Nakayama, 1993) and attention (Davis & Driver, 1998; He & Nakayama, 1992) are known to operate on. Therefore, perceptual grouping of targets should not have been affected by our invisible occluders.

Furthermore, data from our laboratory suggests that tracking of objects is affected by disruption of grouping only when targets and distractors are regrouped to appear within the same groups, not when all items are made less similar to one another (Horowitz et al., submitted; Klieger, Horowitz, & Wolfe, 2004). Even if grouping processes were affected by the changing retinal shapes of the items, the invisible occluders would have reduced overall inter-item similarity without promoting any particular grouping of targets with non-targets. Thus, the invisible occluders are unlikely to have impaired performance by their effects on perceptual grouping.

We suggest that the experiments presented here, combined with our previous results (Alvarez et al., 2005), reveal a flexible attention system capable of time-sharing between two or more demanding visual tasks. Such time-sharing requires an equally flexible memory store that allows critical information from one incomplete task to be maintained while attentional resources are devoted to another task. Here, the simultaneous disappearance of all tracked

items invokes this offline memory store, allowing targets to be recovered more easily. The presence of occlusion cues discourages this process by indicating that the objects are still present and still the focus of attention. Asynchronous disappearance also would tend to act in the same way; if most of the objects are still present and being tracked, then it is unlikely that attention has switched away from tracking.

What can we say about the nature of this "flexible memory store"? In the case of MOT, the representation at a minimum must include the position of the targets at the time of disappearance. Keane & Pylyshyn (submitted) have presented convincing evidence that only location information is stored. We have replicated their results, though our experiments suggest that direction information may also be represented (Fencsik, Horowitz, Klieger, & Wolfe, 2004). Given that non-targets seem to be processed in MOT (Ogawa & Yagi, 2003; Pylyshyn & Leonard, 2003), some information about non-targets may also be stored.

More generally, and more speculatively, we suggest that this offline memory stores the current state of the system as a whole. For example, in visual search, Lleras, Rensink, & Enns (2005) have recently demonstrated a phenomenon they call "rapid resumption". When observers are shown a search display in brief episodes alternating with blank intervals (or another interpolated search task), they are able to pick up search where they left off with no appreciable setup time. This strikes us as quite similar to what we have reported here in MOT. Rather than a new memory store, we might think of this mechanism as a mode of memory, in which the processing state of the system is saved while a given task is briefly put on hold.

Why do we propose a speculative new form of memory, instead of assuming that spatial working memory is involved? We suspect that spatial working memory is indeed involved in MOT, but in the online, rather than offline mode. Awh has argued that there is little distinction between spatial working memory and spatial attention (Awh & Jonides, 1998; Awh, Jonides, & Reuter-Lorenz, 1998); holding a location in memory is equivalent to attending to it. Performing an MOT task may comprise rehearsing target objects in spatial working memory. However, attending to something outside

spatial working memory will disrupt that memory (Awh et al., 1998). In particular, there is strong interference between spatial working memory and visual search (Oh & Kim, 2004; Woodman & Luck, 2004), suggesting overlap in resources. Since we see little interference between visual search and tracking across a gap (Alvarez et al., 2005), it is unlikely that spatial working memory is involved in the latter, though of course the definitive experiment has not yet been performed.

Tracking simultaneously vanishing objects appears to be qualitatively different from tracking occluded objects. We suggest that the ability to recover such objects serves the ecological goal of allowing us to time-share among multiple attention-demanding tasks. It would be dangerous if sustained tracking of behaviorally relevant moving objects (e.g. prey or automobiles) could be completely disrupted by momentary diversion of attention (e.g. avoiding obstacles, dialing a phone).

References

- Alvarez, G. A., Horowitz, T. S., Arsenio, H. C., Dimase, J. S., & Wolfe, J. M. (2005). Do multielement visual tracking and visual search draw continuously on the same visual attention resources? *Journal of Experimental Psychology-Human Perception and Performance*, 31(4), 643-667.
- Ashida, H., Seiffert, A. E., & Osaka, N. (2001). Inefficient visual search for second-order motion. *Journal of the Optical Society of America. A, Optics, Image Science, & Vision*, 18(9), 2255-2266.
- Awh, E., & Jonides, J. (1998). Spatial working memory and spatial selective attention. In R. Parasuraman (Ed.), *The attentive brain*. (pp. 353-380). Cambridge, MA, USA: The Mit Press.
- Awh, E., Jonides, J., & Reuter-Lorenz, P. A. (1998). Rehearsal in spatial working memory. *Journal of Experimental Psychology-Human Perception and Performance*, 24(3), 780-790.
- Bahrami, B. (2003). Object property encoding and change blindness in multiple object tracking. *Visual Cognition*, 10(8), 949-953.
- Bex, P. J., & Baker, C. L., Jr. (1999). Motion perception over long interstimulus intervals. *Perception & Psychophysics*, 61(6), 1066-1074.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 443-446.
- Davis, G., & Driver, J. (1998). Kanizsa subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology-Human Perception and Performance*, 24(1), 169-184.
- Fencsik, D. E., Horowitz, T. S., Klieger, S. B., & Wolfe, J. M. (2004). Target reacquisition strategies in multiple object tracking. *Vision Sciences Society* (Vol. 4, pp. 77).
- He, Z. J., & Nakayama, K. (1992). Surfaces versus features in visual search. *Nature*, 359(6392), 231-233.
- He, Z. J., & Nakayama, K. (1993). Perceived surface shape not features determines correspondence strength in apparent motion. *Vision Research*, 34(16), 2125-2135.
- Horowitz, T., & Treisman, A. (1994). Attention and apparent motion. *Spatial Vision*, 8(2), 193-219.
- Horowitz, T. S., Klieger, S. B., Fencsik, D. E., Yang, K. K., Alvarez, G. A., & Wolfe, J. M. (submitted). Tracking unique objects.
- Keane, B. P., & Pylyshyn, Z. W. (submitted). Can multiple objects be tracked predictively? Clues to enduring visual objecthood.
- Klieger, S. B., Horowitz, T. S., & Wolfe, J. M. (2004). Is multiple object tracking colorblind?, *Vision Sciences Society* (Vol. 4, pp. 75).
- Leonard, C., & Pylyshyn, Z. W. (2003). Measuring the attentional demand of multiple object tracking (MOT), *Journal of Vision* (Vol. 3, pp. 582a).
- Lleras, A., Rensink, R. A., & Enns, J. T. (2005). Rapid resumption of interrupted visual search. *Psychological Science*, 16(9), 684-688.
- Michotte, A., Thines, G., & Crabbe, G. (1964). *Les compliments amodaux des structures perceptives [Amodal complements of perceptual structures]*. Louvain, Belgium: Institut de psychologie de l'Université de Louvain.
- Ogawa, H., & Yagi, A. (2003). Priming effects in multiple object tracking: An implicit encoding based on global

- spatiotemporal information. *Journal of Vision*, 3(9), 339a.
- Oh, S. H., & Kim, M. S. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic bulletin & review*, 11(2), 275-281.
- Palmer, S. E. (2002). Perceptual grouping: It's later than you think. *Current Directions in Psychological Science*, 11(3), 101-106.
- Palmer, S. E., Neff, J., & Beck, D. (1996). Late influences on perceptual grouping: Amodal completion. *Psychonomic Bulletin & Review*, 3(1), 75-80.
- Petersik, J. T. (1989). The Two-Process Distinction in Apparent Motion. *Psychological Bulletin*, 106(1), 107-127.
- Pylyshyn, Z. (1994). Some primitive mechanisms of spatial attention. *Cognition*, 50(1-3), 363-384.
- Pylyshyn, Z. W., & Leonard, C. (2003). Inhibition of nontargets during multiple object tracking (MOT). *Journal of Vision* (Vol. 3, pp. 585a).
- Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: evidence for a parallel tracking mechanism. *Spatial Vision*, 3(3), 179-197.
- Scholl, B. J., & Feigenson, L. (2004). When out of sight is out of mind: Perceiving object persistence through occlusion vs. implosion. *Journal of Vision*, 4(8), 26a.
- Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: clues to visual objecthood. *Cognitive Psychology*, 38(2), 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(1-2), 159-177.
- Woodman, G. F., & Luck, S. J. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin & Review*, 11(2), 269-274.
- Yantis, S. (1992). Multielement visual tracking: Attention and perceptual organization. *Cognitive Psychology*, 24, 295-340.
- Yin, C., & Thornton, I. M. (1999, May 7-8, 1999). *Visual tracking across multiple streams of information*. Paper presented at the 3rd Annual Vision Research Conference, Ft. Lauderdale, FL.

Author's Note

This research was supported by NIMH R01 65576 to TSH. We would like to thank Edward Awh, Anne Hillstrom, Shlomo Sher, Zenon Pylyshyn, Steven Yantis, and an anonymous reviewer for constructive criticism and useful suggestions.