Published in final edited form as:

Vis cogn. 2008 January 1; 16(4): 486–513. doi:10.1080/13506280601087356.

Objects or Locations in Vision for Action? Evidence from the MILO task

Todd S. Horowitz and

Brigham and Women's Hospital, Harvard Medical School

Ian M. Thornton

University of Wales Swansea

Abstract

In the multi-item localization task (MILO, Thornton & Horowitz, 2004), observers are asked to find an ordered sequence of targets. We can measure the influence of both past actions and future plans on search for the current target. Our previous work with static search arrays found evidence for both retrospective and prospective memory. Responding to a target eliminated its influence on subsequent responses, while observers consistently planned ahead at least one item into the future. Here we asked whether these effects were based in location- or object-based reference frames. We used dynamic arrays in which observers had to search for multiple moving targets. Our results suggest that observers can still plan ahead effectively in this dynamic environment, indicating that future target objects can be tracked as they change position. However, memory for previous targets is essentially eliminated, suggesting that locations, not objects, were being tagged in our previous work.

Introduction

We recently introduced a new task for exploring the spatio-temporal context of search (Thornton & Horowitz, 2004). In the multi-item localization (MILO) task, observers are presented with a stimulus array containing multiple target and distractor items (typically alphanumeric characters), and are asked to click on an ordered sequence of targets. The MILO task thus differs from the classic visual search paradigm in two important ways. First, requiring observers to interact with targets, as opposed to a categorical detection or identification response, more closely mimics our natural interactions with the environment. Second, the use of a sequence of targets, as opposed to a single target, makes it possible to explore the influence of both past and future actions within the space of a single trial.

illustrates a typical MILO trial. The observer is first presented with a cue display, indicating the sequence of targets. Once the observer signals that she has memorized the sequence, a search array is presented, consisting of a set of ovals marked with target and distractor characters. The primary dependent variable is the sequential reaction time (SRT), the time to click on target n minus the time to click on target $n-1$.

Two simple manipulations allow us to measure the influence of past and future actions independently. To explore retrospective memory, we compare a condition in which targets vanish from the screen as soon as the observer responds to them to a condition in which the targets remain on the screen. Physically removing the targets from the screen simulates a perfect

retrospective memory. To the extent that performance is inferior when the old targets remain on the screen, we infer that old targets still have some influence on behavior.

Prospective memory can be revealed by altering the target sequence mid-trial. For example, in Thornton and Horowitz (2004), we introduced a manipulation in which each response triggered the shuffling of future target and distractor items. The physical layout of these ovals on the screen was not altered, with the shuffle manipulation simply randomizing which item was in which oval. A key feature of this manipulation was that old targets were left unchanged. This allowed us to keep prospective and retrospective effects separate. The shuffle manipulation prevented observers from gaining any advantage from planning ahead. Aspects of the data that changed when future target locations were shuffled in this fashion could thus be attributed to prospective memory.

Both retrospective and the prospective effects can be seen in , which is redrawn for the data reported in Thornton and Horowitz (2004). This graph shows SRT for a sequence of 4 targets within a field of 4 distractors. In panel A, SRT decreases with each response, a pattern that would be predicted in the case where targets vanish after a response (open symbols), as the physical set-size is being reduced. The fact that a similar pattern is obtained when old targets remain on the screen (filled symbols) illustrates the efficiency of the retrospective memory. The large drop in SRT between items 1 and 2 in both curves illustrates the prospective effect. This gap can be eliminated by shuffling the position of future target and distractor items (panel B). Under these conditions all subsequent SRTs remain at the level of the initial SRT.

All of the experiments reported by Thornton and Horowitz (2004) used static search displays. That is, the physical location of the target and distractor ovals never varied during the course of a trial (though the contents of the ovals were sometimes shuffled, as noted above). The purpose of the current work was to examine MILO performance in situations where search displays were dynamic rather than static, with all ovals constantly in motion.

Our primary purpose in introducing motion to the MILO task was to study the frame of reference in which retrospective and prospective effects take place. Our previous data (Thornton & Horowitz, 2004) are compatible with both location-based and object-based reference frames. That is, observers might have used memory for target locations, or some form of location-invariant representation. With the addition of motion, the ability to use location-based coding should be eliminated or at least severely disrupted. If similar patterns of retrospective and prospective performance were obtained with dynamic displays, this would argue for a location-invariant representation. Is this a plausible outcome?

A substantial literature, beginning with the seminal work of Duncan (1984) supports the notion that attentional selection can act on objects, instead of (or at least in addition to) locations (Abrams & Law, 2000; Davis, Welch, Holmes, & Shepherd, 2001; Schendel, Robertson, & Treisman, 2001; S. E. Watson & Kramer, 1999). A number of selection-related effects, such as negative priming (Shapiro & Loughlin, 1993) and inhibition of return (IOR, McAuliffe, Pratt, & O'Donnell, 2001; Tipper, Jordan, & Weaver, 1999; Tipper, Weaver, Jerreat, & Burak, 1994) have been shown to be coded in both location- and object-centered coordinate systems. Indeed, coding in multiple reference frames seems to be a common representational strategy (Danziger, Kingstone, & Ward, 2001; Ghafouri, Archambault, Adamovich, & Feldman, 2002; Lamy & Tsal, 2000).

We reasoned that both retrospective and prospective memory for targets in the MILO task would likely use multiple coding systems. Search and manipulation of real-world objects would seem to call for an object-centered frame of reference. However, studies demonstrating object-centered effects often use sparse displays with many fewer objects than the MILO paradigm. Object-centered representations might be more computationally expensive, leading to capacity

limitations in the number of objects that can be tracked. Our best guess as to what this capacity might be comes from studies of multiple object tracking (MOT, Pylyshyn & Storm, 1988). Such studies have established that the visual system can track several independently moving items simultaneously; the usual capacity estimate is 3–5 items. (for a review, see Cavanagh & Alvarez, 2005). However, this varies inversely with item speed; at the low velocities we employed in these experiments, observers should be able to track roughly 7–8 items (Alvarez & Franconeri, 2005).

In Experiment 1, we took the basic MILO paradigm and set the ovals in motion at 1.2 degrees visual angle per second ($^{\circ}/s$). The results were broadly similar to those shown in A, except that SRTs declined more steeply for the Vanish condition than for the Remain condition, indicating that retrospective memory was impaired. Experiment 2 replicated this finding, using a larger set size to more clearly expose the impact of motion on retrospective performance. In Experiment 3, we introduced the shuffle manipulation, described above, which is known to disrupt prospective memory in the MILO task. Retrospective memory was again impaired in Experiment 3. More importantly, the shuffle manipulation removed the gap between the first and subsequent items, confirming that, observers were still able to plan ahead in Experiments 1–2. In Experiment 4, we slowed the motion to, 0.9, 0.7, or 0.5 $^{\circ}/s$, but there was no sign that retrospective memory became more effective as the speed was reduced. In Experiment 5, we reduced the velocity even further, to 0.25 $^{\circ}/s$, without changing the results. Finally, in Experiment 6, we introduced rigid motion, so that the configuration of items remained unchanged throughout a trial; this did not alter the results. Together, these results suggest that retrospective memory in the MILO task is location-based, while prospective memory is location-invariant.

General Methods

All the experiments shared the same basic methodology. The method sections for individual experiments will specify only deviations from this procedure. Experiments 1 & 2 were conducted at the Max Planck Institute for Biological Cybernetics, Tübingen, Germany, Experiments 3, 4 & 6 at the Visual Attention Laboratory, Brigham and Women's Hospital, Boston and Cambridge, MA, USA, and Experiment 5 at the Psychology Department, University of Wales Swansea, Wales, UK.

Participants

All observers had normal or corrected to normal vision and were naïve as to the purpose of the research until after the experimental session. Observers provided informed consent, either according to procedures approved by the Max Planck Institute, the Partner's Healthcare Corporation Institutional Review Board, or the University of Wales, depending on where the study was conducted. Observers in Tübingen were paid 8 t/hour for participation, observers in Cambridge and Boston \$10/hour and observers in Swansea £8/hour.

Equipment

Software was custom written using routines based on work by Steinman and Narwot (1992), Pelli and Zhang (1991) and Rensink (1990). Observers were seated at a standard viewing distance of 60 cm in front of the monitor. All experiments were conducted on Macintosh computers and 21" color monitors with a refresh rates of 75 Hz. Screen resolution was either 1152 × 780 (Experiments 1–2 & 5) or 1024 × 640 (Experiment 3–4 & 6). These resolution differences led to slight variation in the size and speed of stimuli, which will be noted in the appropriate experiments¹. In the next section we assume a setting of 1152 × 780.

Stimuli

Stimuli were presented on a uniform, dark gray background. A central viewing area subtending approximately 12.0° was delineated by a black outline square. The search items were light gray oval patches, subtending approximately 1.0° in diameter. The border of the oval was drawn in black and each item contained a single upper case letter, also drawn in black. The bounding box of each letter subtended 0.6° . The letters were always present in the ovals, with the following exception: When an observer clicked on any target, the center of *all* ovals was redrawn in the background gray for 50 ms. This global transient was introduced in order to mask the local transients when letters were shuffled between ovals (see Experiments 3–5 for details). The global transient does not affect any of the experimental manipulations (Thornton & Horowitz, 2004).

In the search display, initial positions of all eight items were randomized within the central viewing area on a trial by trial basis. The only constraint on this position was that each item must be at least 1.0° away from all other items and the edge of the viewing area. As soon as the search display appeared, all items began to move in pseudo-random directions at a fixed velocity of $1.2^\circ/\text{s}$. The animation was achieved by shifting each oval by one pixel after a delay of two screen refresh cycles. Initial path lengths and directions were computed separately for each item. Directions were selected from one of 8 cardinal compass points (i.e., N, NE, E, SE, S, SW, W, NW) and path distance was randomly varied within the range of 2.6° to 3.4° . When an item had traversed its initial path distance, a new direction was randomly selected and a new path length assigned.

When items approached the edge of the central square, the direction of motion was simply reversed, so that the object appeared to bounce away from the edge in a predictable manner. When two items approached each other they did not bounce, but overlapped, with one item passing over and occluding the other. The occlusion precedence (which item was in front) was not predictable on the basis of target sequence order, as items were randomly assigned to ovals. Importantly the current target item was always selectable during such occlusion.

Task

Each trial began with the cue phase, in which the four target items were displayed in the center of the screen, with two items on either side of a central fixation cross (see Figure 1). The target letters always formed an alphabetic sequence (e.g. E, F, G, H) that was randomly chosen on each trial. The sequence could fall anywhere in the Roman alphabet between the letters C and X. Observers were free to view the target sequence as long as they liked and then initiated the search phase of the trial by pressing the spacebar.

The search display contained the four target items plus four distractor items. The distractor items flanked the target set in the alphabet (e.g. C, D, I, J). The observer's task was to locate each of the target letters *in sequence*, beginning with the first letter (i.e., E) and ending with the last letter (i.e., H). Observers responded by clicking anywhere within the target oval using a standard computer mouse. A trial was terminated, and the screen blanked, either when the four-item target sequence had been successfully completed, or when an error occurred. An error could occur either because a mouse click occurred in the wrong oval ("sequence errors"), or a mouse click occurred in the background (i.e. outside of any oval, "click errors").

¹While these changes were not intentional, we have previously varied the size of displays within this range with no influence on performance. In Experiments 4 and 5, we explicitly manipulated speed, again with no obvious effects over a much larger range.

Procedure

There were two main stimulus conditions. In the *Vanish* condition, each target item disappeared after the observer clicked on it. In the *Remain* condition, targets remained on the screen and continued to move after they were clicked on.

Observers completed two blocks of each condition. Each block contained a minimum of 30 trials. Trials in which an error occurred were replaced, so that a block only finished when observers had successfully completed 30 trials. Each observer thus completed a minimum of 120 experimental trials, taking short breaks after each block. Block order was counterbalanced with half of the observers completing the two *Vanish* blocks before moving on to the *Remain* blocks, and the other half seeing the opposite order.

Observers were first shown several example trials. They were then familiarized with the response method, and allowed to practice until they were comfortable using the mouse. The error conditions were illustrated to all observers, and they were informed that a block would only terminate when they had successfully completed 30 trials. Therefore, while speed of response was emphasized, they were also motivated to avoid errors.

Note that while observers were provided with a fixation cross at the start of the trial, the cross disappeared during the search phase. Observers were free to move their eyes, and we did not record eye position.

Data Analysis

Median SRT was the primary dependent variable, defined as the time to click on target n minus the time to click on target $n-1$. Error rates were arcsine transformed before analysis, and we converted the transformed means back into percentages before reporting. We also computed within-subjects confidence intervals (Loftus, 1994).

Experiment 1

The purpose of Experiment 1 was to determine whether the previously reported ability to “ignore” old targets in the MILO task was based on memory for the *location* of an object or for the object itself.

Methods

Twelve observers from the Tübingen community participated in this study. Screen resolution was set to 1152×780 .

Results

Errors are shown in the right-hand panel of Figure 3. There were substantially more click errors than sequence errors. However, neither error type varied with condition (click errors: $F(1, 11) = 1.0$, $p > .10$, $MSE = .02$; sequence errors: $F(1, 11) < 1.0$, $MSE = .01$).

SRTs are plotted in the left-hand panel of Figure 3 as a function of target number. SRTs were an average of 62.9 ms slower in the *Remain* condition than in the *Vanish* condition ($F(1, 11) = 7.3$, $p < .05$, $MSE = 13092$). SRT also declined significantly with target number ($F(3, 33) = 239.2$, $p < .000001$, $MSE = 6719$). There was no interaction ($F(3, 33) < 1.0$, $MSE = 3158$).

Figure 2. Results from the MILO paradigm using static displays (redrawn from Thornton & Horowitz, 2004). SRTs are plotted as a function of target number and condition. Open symbols denote the “vanish” condition, filled symbols the “remain” condition (see text). Panel A depicts the results when stimulus positions remained constant throughout a trial. Panel B depicts the

results when we shuffle the locations of distractors and targets that have not yet been responded to.

Figure 2. Results from the MILO paradigm using static displays (redrawn from Thornton & Horowitz, 2004). SRTs are plotted as a function of target number and condition. Open symbols denote the “vanish” condition, filled symbols the “remain” condition (see text). Panel A depicts the results when stimulus positions remained constant throughout a trial. Panel B depicts the results when we shuffle the locations of distractors and targets that have not yet been responded to.

suggests that the main effect of target number was powerfully driven by the drop from the first target to the second target. However, the analysis did not change substantially when we eliminated the initial SRT. The main effect of condition was still significant ($F(1, 11) = 5.4$, $p < .01$, $MSE = 14056$), as was the main effect of target number ($F(2, 22) = 40.0$, $p < .00001$, $MSE = 4723$). The interaction term increased, but not to the level of significance ($F(2, 22) = 1.3$, $p > .10$, $MSE = 2988$).

We also computed $SRT \times$ target number slopes (based on the last 3 SRTs). The slopes for the Remain condition averaged -60.5 ms/target, compared to -86.0 ms/target for the Vanish condition ($CI \pm 21.0$ ms/target). This difference was marginally significant ($F(1, 11) = 3.6$, $p = 0.08$).

Discussion

An initial comparison between Figure 2 & Figure 3 would suggest that the current data quite closely resemble those we collected with static stimuli (Thornton & Horowitz, 2004). Particularly prominent is the large drop in SRT after the first target, a feature that we have previously attributed to prospective memory; during the search for the first target, observers were able to obtain some information about the location of subsequent targets. This ability would appear to be intact, even when the objects are in motion. We will return to this issue in Experiment 3.

In the retrospective domain, however, motion does appear to have had an effect on behavior. In our previous studies with static stimuli, the Vanish and Remains conditions always yielded virtually identical patterns of results, suggesting highly efficient memory for past targets. In the current experiment, SRTs in the Remain condition were consistently slower than in the Vanish condition. More importantly, the decline in SRT with target order was also generally shallower. Such disruption implies that retrospective performance in the MILO task is sensitive to object location, rather than being based exclusively on object identity.

Motion did have a very clear impact on the error rates. Although there was no difference between the Vanish and Remain conditions, the overall level of errors was considerably larger than that observed with static arrays (Thornton & Horowitz, 2004). In our experiments with static stimuli, overall errors seldom exceeded 5%, compared to 15–17% in the current experiment. However, the bulk of the errors in this experiment were click errors, rather than sequence errors.

A sequence error occurs when a successful motor movement is aimed at the wrong object. Thus, we propose that sequence errors measure cognitive errors, where the observer at some stage selected the wrong object for action. Click errors, on the other hand, represent motor errors in aiming and executing mouse movements. There may be some cases in which conflict at the cognitive level leads to an inaccurate motor movement, or motor errors which by chance land in the wrong object. However, we suspect these cases are rare. Thus, we hold that sequence errors measure cognitive difficulty and click errors motor difficulty.

In this framework, then, it is significant that the bulk of the errors in this experiment were click errors (see the right-hand panel of Figure 3). The rate of sequence errors here (1–2%) is comparable with data from static arrays. However, while click errors in the static arrays were also around 1–2%, here they are 13–16% of trials. We suggest that these errors simply reflect the difficulty of hitting a moving target. In agreement with this idea, subsequent experiments with slower speeds will produce lower click error rates, while sequence errors will remain uniformly low. Importantly, these effects do not interact with the reference frame effects which are the primary focus of this paper.

Overall, the current results suggest that it is the retrospective aspects of performance that are most disrupted by motion. Memory for old items may thus rely on location- rather than object-based representations. With a target sequence of only 4 items, however, it may be premature to exclude the possibility of object-based tagging, especially since the effect was statistically weak. We designed the next experiment to obtain a clearer picture of the relationship between Vanish and Remain conditions in the context of motion.

Experiment 2

One simple way to gain more information about the relationship between retrospective performance in the Vanish and Remains conditions is to increase the number of responses that have to be made. In Experiment 2, we increased the target sequence length from four to six items. If motion does eliminate the ability to retrospectively tag old targets, then SRTs for Vanish and Remain should clearly diverge as a function of position within the target sequence.

Method

Fourteen observers from the Tübingen community were paid for participation in this study. We followed the General Methods except for the following two changes. First, the number of target items was increased from four to six. Second, to maintain the ratio of target to distractor items, the number of distractors was also increased to six, giving an overall display set size of twelve. All other aspects of the display and task were identical to Experiment 1.

Results

Two observers were removed from the analysis, one due to exceptionally slow RTs, and one for very high error rates. Error rates for the remaining twelve observers are shown in the right-hand panel of Figure 4. Neither type of error varied with condition (both $F(1, 11) < 1.0$, $MSE = .02$).

SRTs are plotted in the left-hand panel of Figure 4. As in Experiment 1, SRTs were significantly longer for the Remain condition than for the Vanish condition ($F(1, 11) = 19.1$, $p < .005$, $MSE = 17858$). SRTs also declined with target number ($F(5, 55) = 76.7$, $p < .000001$, $MSE = 14396$), but this decline was steeper for the Vanish data ($F(5, 55) = 6.1$, $p < .0005$, $MSE = 6452$).

As in the previous experiment, the drop from target 1 to target 2 seems to account for much of the effect of target number, so we reanalyzed the data, excluding RTs to the first target. The same effects were obtained: main effects of both condition ($F(1, 11) = 17.8$, $p < .005$, $MSE = 21609$) and target number ($F(4, 44) = 10.7$, $p < .000001$, $MSE = 14872$), as well as an interaction ($F(4, 44) = 7.0$, $p < .0005$, $MSE = 5445$).

In order to explore the interaction, we computed the $SRT \times$ target number slopes for the two conditions (again excluding the initial RT). The slope in the Vanish condition, at -75.2 ms/target, was significantly steeper than the -27.4 ms/target slope ($CI \pm 18.9$ ms/target) in the Remain condition ($F(1, 11) = 15.5$, $p < .005$, $MSE = 887$).

Discussion

With the extended target sequences in Experiment 2, the difference between Vanish and Remain conditions became much clearer. As expected, the Vanish condition SRTs declined sharply as the physical set size was reduced. The corresponding slope for the Remain condition was substantially shallower. This finding indicates that retrospective memory was not as efficient for moving targets as for static targets, where the slope of the two functions has previously been shown to be essentially identical (Thornton & Horowitz, 2004). As in Experiment 1, there was a large drop in SRT after the first response for both Vanish and Remain conditions, suggesting that initial planning was not disrupted by object motion. Error rates were again higher than previously observed with static displays (Thornton & Horowitz, 2004), but did not vary as a function of condition.

Experiment 3

Another way in which to focus more clearly on the retrospective aspects of the MILO task is to eliminate the ability to plan ahead. As noted above, the signature of prospective memory in this paradigm is the large decrease in SRT between the first and second targets. This signature is clearly present in the data from the previous two experiments. We hold this pattern to be the consequence of observers picking up advance information about future targets during search for the first targets. Experiments 1 and 2 thus suggest that the prospective aspects of the MILO task are less tied to location and may indeed operate in an object-based manner. We return to this issue in the General Discussion, but in the current paper will continue to focus on the retrospective aspects of the MILO task.

With static arrays, the ability to plan ahead was easily disrupted by *shuffling* the contents of future target items at the time of a response. In this way information about the position of such items was rendered useless. This manipulation removed the large gap between responses 1 and 2, but did not affect retrospective performance (Thornton & Horowitz, 2004, Experiment 2). Here we replicated that manipulation with dynamic arrays. There were two main goals. First, to reaffirm that gap between the first and second responses does relate to forward planning. Second, to provide a picture of retrospective performance in the absence of such forward planning. We predicted that the SRT gap between the first and second items would be substantially reduced and the difference between Vanish and Remain conditions would be more clearly seen as all four responses would now be of the same order of magnitude.

Method

Sixteen observers were recruited from the Brigham and Women's Hospital Visual Attention Laboratory volunteer panel. The procedure was as described in the General Methods, with the following exceptions. After the observer clicked on a target, the locations of all distractor letters and subsequent target letters were randomly reshuffled among the current stimulus locations. Targets that had already been clicked on were unchanged. It is important to note that the overall layout of the display was not changed, only the contents of the ovals (i.e. the letters) were updated. As previous studies had indicated that the shuffle manipulation is quite demanding for observers we also returned to using four instead of six items.

Screen resolution was set to 1024×768 in this experiment. Thus, the overall size of the stimulus area was 13° and individual targets subtended 1.1° . Items moved at $1.4^\circ/\text{s}$.

Results

Errors are shown in the right-hand panel of Figure 5. Neither click nor sequence errors differed by condition (both $F(1, 15) < 1.0$, $\text{MSE} = .01$). The left-hand panel of Figure 5 plots SRTs as a function of condition and target number. As in previous experiments, SRTs in the Remain

condition were longer than in the Vanish condition ($F(1, 15) = 17.7, p < .001, \text{MSE} = 27257$), and there was a decline in SRT with target number ($F(3, 45) = 30.7, p < .0001, \text{MSE} = 5168$), which was steeper for the Vanish than Remain data ($F(3, 45) = 18.6, p < .0001, \text{MSE} = 4046$).

Unlike in the first two experiments, we did not observe an elevated RT to the first target. This is a typical result from the shuffle manipulation (Thornton & Horowitz, 2004). Therefore, we did not reanalyze data without the first RT. The $\text{SRT} \times \text{target number}$ slope in the Remain condition was -26.0 ms/target , which was significantly shallower than the -51.2 ms/target ($\text{CI} \pm 11.4 \text{ ms/target}$) slope in the Vanish condition ($F(1, 15) = 11.2, p < .005, \text{MSE} = 458$).

Discussion

Experiment 3 yielded two major findings. First, as we predicted, the large drop in SRT from the first to second targets was abolished, and the $\text{SRT} \times \text{target}$ function became quite linear for the Vanish condition. This suggests that prospective memory functions similarly in both static and dynamic displays, and that in Experiments 1 and 2, such forward planning used an object-based coordinate system.

Second, we observed very little in the way of retrospective effects in the current experiment. There was a substantial difference between the Remain and Vanish conditions, both in terms of mean SRT and the slope of the $\text{SRT} \times \text{target}$ function. While the slope in the Vanish condition was somewhat shallower than the -70 to -90 ms/target range seen in the previous two experiments, it was still twice as steep as the Remain slope, which was slightly shallower than that observed in Experiment 2. This suggests that observers were largely failing to keep track of old targets. In the absence of forward planning, this pattern was much easier to see.

Experiment 4

The difference between the Vanish and Remain conditions in Experiments 1–3 suggests that in the context of motion, observers are largely unable to ignore old targets. This suggests that our previous retrospective findings were based on memory for location rather than memory for objects. In the next two experiments we begin to explore the precision involved in such retrospective memory representations.

The object motion we have used so far had a constant velocity of either 1.2 or $1.4^\circ/\text{s}$. If we assume an average response time of around 1200 ms , it is clear that previous targets will have moved completely outside of their initial locations by the time a subsequent response is made. Is a complete lack of overlap with initial position necessary in order to abolish the retrospective effects? Furthermore, average SRTs became longer over the first three experiments, and the retrospective effect correspondingly smaller. These two facts might be related: the longer an observer takes to make a response, the further each old target moves from its remembered location.

In Experiment 4, we tested the hypothesis that the strength of retrospective memory would vary inversely with item speed. We measure the strength of retrospective memory as the inverse of the slope of the $\text{SRT} \times \text{target}$ function in the Remains condition. Our prediction is that this slope will be steeper for slower speeds than for faster speeds, because targets will stay closer to their remembered locations.

Method

Twenty-three observers from the Brigham and Women's Hospital Visual Attention Laboratory volunteer panel participated in this study. Experiment 4 was identical to Experiment 3, except for the following modifications. First, there were two groups. One group of twelve observers ran only the Vanish condition, the other, of eleven observers, ran only the Remain condition.

This modification was made to maintain the overall number of trials performed by each observer within the range used in Experiments 1–3. Second, we varied the speed of object motion in separate blocks. Motion was always slower than that used in previous experiments, and was constant at either 0.5°, 0.7° or 0.9°/second.

Results

Error data are shown in Figure 6C. Neither type of error rate differed between groups (click errors: $F(1, 21) = 2.1, p > 0.10, \text{MSE} = 0.07$; sequence errors: $F(1, 21) < 1.0, \text{MSE} = 0.03$). Click errors increased with speed ($F(2, 42) = 12.8, p < 0.00005, \text{MSE} = 0.01$), more rapidly so for the observers in the Vanish condition than for those in the Remain condition ($F(2, 42) = 4.5, p < .05, \text{MSE} = 0.01$). The simple effect of speed was significant for the Vanish observers ($F(2, 22) = 16.4, p < 0.00005, \text{MSE} = 0.01$), but not for the Remains observers ($F(2, 20) = 1.6, p > .10, \text{MSE} = 0.01$). Sequence errors, however, did not respond to speed ($F(2, 42) < 1.0, \text{MSE} = 0.01$), nor was there a speed $\times$ group interaction ($F(2, 42) = 1.1, p > 0.10, \text{MSE} = 0.01$).

SRTs (Figure 6A & B) were analyzed via a mixed factorial ANOVA with condition as a between subjects variable and speed and target number as repeated measures. The difference between the Remain and Vanish groups was marginal ($F(1, 21) = 3.0, p = .10, \text{MSE} = 217437$). There was a main effect of target number ($F(3, 63) = 19.7, p < .000001, \text{MSE} = 3990$), which interacted with condition ($F(3, 63) = 11.9, p < .000005, \text{MSE} = 3990$). There was no effect of speed ($F(2, 42) < 1, \text{MSE} = 6678$), nor did speed interact with target number ($F(6, 126) < 1, \text{MSE} = 6678$), or condition ($F(2, 42) < 1, \text{MSE} = 3367$). The three-way interaction was not significant ($F(6, 126) = 1.5, p > .10, \text{MSE} = 3367$).

In order to further explore the condition $\times$ target number interaction, we computed the SRT $\times$ target number slopes for the two groups (Figure 6D). Slope differed significantly between the two groups ($F(1, 21) = 23.9, p < .0001, \text{MSE} = 1151$), but did not vary as a function of speed ($F(2, 42) = 1.4, p < 0.10, \text{MSE} = 758$), nor did speed and group interact ($F(2, 42) = 2.0, p > 0.10, \text{MSE} = 758$).

Discussion

The current data appear to replicate Experiment 3 at all three speeds. Notably, our prediction that slope should be steeper for slower speeds in the Remain condition was not borne out. This would have produced a monotonic downward trend in the filled circles in. Instead, we see a nonmonotonic function, with the shallowest slope at the slowest speed. This suggests that either setting the items in motion disrupts retrospective memory altogether, or the spatial representations of old targets has a very tight spatial gradient which has already fallen off substantially within the average response time in this experiment.

Consistent with the idea that click errors primarily reflect the difficulty of clicking on a moving target, click errors increased with speed in this experiment. This was particularly true for the Vanish group. In this experiment, unlike prior (and subsequent) experiments, we observed a difference between the Remain and Vanish conditions in the click error rate. This might reflect a speed/accuracy tradeoff, such that observers in the Remain condition were responding more slowly and cautiously than observers in the Vanish condition.

Sequence errors were again quite low, unaffected by speed, and similar between the two groups, suggesting that the cognitive difficulty of the task was constant across all conditions.

Experiment 5

In Experiment 5, we decided to further explore the spatial gradient idea by conducting one final speed manipulation in which object motion was present but minimized. That is, we moved

the objects so slowly that there would be considerable overlap between original and final positions throughout the course of a trial. If retrospective memory is still disputed under these conditions, it would seem to suggest that any change in position can render past information uninformative.

Method

Twelve students from the University of Wales community participated in this experiment. The design was essentially the same as Experiment 3, with the following exceptions. First, speed was reduced to 0.25°/s. Screen resolution was set to 1152×870 .

Results

Error rates are shown in the right-hand panel of Figure 7. Neither click ($F(1, 8) < 1.0$, $MSE = .01$) nor sequence errors ($F(1, 8) < 1.0$, $MSE = .02$) differed as a function of condition.

The left-hand panel of Figure 7 plots SRTs as a function of condition and target. SRTs declined with target ($F(3, 30) = 4.2$, $p < .05$, $MSE = 5982$), and were faster in the Vanish condition ($F(1, 10) = 8.5$, $p < .05$, $MSE = 7229$). These two factors interacted ($F(3, 30) = 6.2$, $p < .005$, $MSE = 3557$) such that the decrease in SRT with targets was observed primarily in the Vanish data. Analysis of simple effects revealed that SRTs declined in the Vanish condition ($F(3, 30) = 10.3$, $p < .0001$, $MSE = 4406$), but not in the Remain condition ($F(3, 30) < 1$, $MSE = 5133$).

We can also analyze the interaction by looking at the $SRT \times$ target number slopes, which were significantly steeper for the Vanish condition (-44.9 ± 12.2 ms/target) than for the Remain condition (-4.4 ± 12.2 ms/target; $F(1, 10) = 27.4$, $p < .0005$, $MSE = 329$).

Discussion

Even at a very slow speed, when there was considerable overlap between original and final positions, retrospective effects were clearly absent. This suggests that any change of position, however small, can disrupt memory for old targets. Error rates in this experiment did not vary between conditions and were as low as those previously observed with static arrays (Thornton & Horowitz, 2004). We thus obtain the disruption of retrospective effects even in the absence of elevated error rates observed in previous experiments.

Experiment 6

The evidence we have reported so far strongly argues against object-based coding of retrospective memory in the MILO task. However, our method for dissociating objects and locations is slightly different from that used in previous studies. Consider the work of Tipper and colleagues (Tipper, Driver, & Weaver, 1991; Tipper, Weaver, Jerreat, & Burak, 1994) which is perhaps most relevant here as it also used motion to dissociate location and objects. In a typical example of their paradigm, four objects might be presented, distributed symmetrically around fixation. One object is cued, then the objects rotate rigidly around the fixation point for 90°. At this point, one can probe either the originally cued object, or the uncued object which is now in the originally cued location; the standard result is that both objects show some evidence of the effect in question, demonstrating that both object- and space-based reference frames are being used.

In the Tipper paradigm, the overall configuration of the display remains constant throughout the experiment, allowing the experimenter to dissociate *absolute* location from *relative* location. In contrast, in the studies described above, each item was given a unique and random motion path. This non-rigid motion pattern ensured that the overall display configuration was always disrupted.

In the current experiment, rather than having each item move on an independent trajectory, we yoked all items together and shifted them using a single random path. This rigid motion ensured that the display configuration was maintained throughout a trial, allowing observers to rely on relative locations. Does this manipulation allow observers to maintain retrospective memory for targets?

Method

Sixteen observers from the Brigham & Women's Hospital Visual Attention Laboratory volunteer panel participated in this study. Experiment 6 was identical to Experiment 3, except that the motion of all items was slaved to a randomly chosen item and the overall display size was increased 18°. Speed was set to 1.4°/s. Display resolution was set to 1024 × 768.

Results

Errors are shown in the right-hand panel of Figure 8. Neither click ($F(1, 15) < 1.0$, $MSE = .01$) nor sequence errors ($F(1, 15) < 1.0$, $MSE = .02$) differed as a function of condition.

SRTs are shown in the left-hand panel of Figure 1. Example stimulus displays from the MILO task.

Figure 2. Results from the MILO paradigm using static displays (redrawn from Thornton & Horowitz, 2004). SRTs are plotted as a function of target number and condition. Open symbols denote the “vanish” condition, filled symbols the “remain” condition (see text). Panel A depicts the results when stimulus positions remained constant throughout a trial. Panel B depicts the results when we shuffle the locations of distractors and targets that have not yet been responded to.

Figure 3. Left-hand panel plots SRT medians from Experiment 1 as a function of target number. Filled symbols denote the Remain condition, open symbols the Vanish. Right-hand panel plots percent error as a function of error type. Gray bars denote the Remain condition, open bars the Vanish. In this and all subsequent figures, error bars indicate within-subjects 95% confidence intervals (Loftus, 1994).

Figure 4. Left-hand panel plots SRT data from Experiment 2 as function of condition and target number. Filled symbols denote the Remain condition, open symbols the Vanish. Right-hand panel plots percent error as a function of error type. Gray bars denote the Remain condition, open bars the Vanish.

Figure 5. The left-hand panel plots median SRT data from Experiment 3 as a function of condition and target number. Solid symbols denote data from the Remain condition, open symbols data from the Vanish condition. Right hand panel plots percent error for the two error types. Gray bars denote the Remain condition, open bars the Vanish.

Figure 6. Data from Experiment 4. The top two panels plot SRTs as a function of speed and target number, with the Remain condition observers shown in Panel A and the Vanish condition observers shown in Panel B. Panel C depicts % error as a function of speed for the Remain (filled symbols) and Vanish (open symbols) groups. Click errors are plotted as circles, sequence errors as squares. Panel D plots the slope of the SRT × target functions for the Remain (filled symbols) and Vanish (open symbols) groups.

Figure 7. Right-hand panel plots SRT data from Experiment 5 as a function of target number. Filled symbols denote the Remain condition, open symbols denote the Vanish condition. Left-hand panel plots percent error as a function of error type. Gray bars denote the Remain condition, open bars the Vanish.

Figure 8. The Remain condition produced slower SRTs than the Vanish condition ($F(1, 15) = 9.9, p < 0.01, \text{MSE} = 4156$). SRTs decreased with target number ($F(3, 45) = 13.6, p < 0.01, \text{MSE} = 6099$), but more steeply in the Vanish condition ($F(3, 45) = 19.6, p < 0.000005, \text{MSE} = 1599$). Simple effects of target number were significant for the Vanish condition ($F(3, 45) = 11.9, p < 0.00001, \text{MSE} = 3869$) but not for the Remain condition ($F(3, 45) = 1.3, p > .10, \text{MSE} = 3829$). $\text{SRT} \times \text{target}$ slopes were 1.5 ms/target (± 11.7) for the Remain condition and -38.9 ms/target (± 11.7) for the Vanish condition, a significant difference ($F(1, 15) = 27.1, p < .0005$).

Discussion

The current data closely resemble those obtained in Experiments 3 through 5 (compare Figure 7 & Figure 8). The rigid motion of the display seems to have speeded overall responses times somewhat and error rates fall between those seen with static arrays and random motion. However, there is still a very clear difference between the Vanish and Remain conditions. This suggests there is no retrospective advantage to being able to code past targets relative to each other, or to an overall display framework.

General Discussion

We introduced the MILO task in order to study the spatio-temporal context of visual search. In our previous work, using static stimuli, we showed that both past actions and future plans affect responses to the current target (Thornton & Horowitz, 2004). In the experiments reported here, we used motion to explore the frame of reference within which these effects occur. Our results demonstrated a clear dissociation between the retrospective and prospective aspects of the task. Simply setting the search items in motion largely disrupted the utility of retrospective memory, while leaving prospective memory intact. These two aspects of the task will be discussed in turn.

The hallmark of retrospective effects in the MILO task is the near identical performance observed in the Vanish and Remain conditions. With static arrays, a response to an item effectively removes it, improving search efficiency, regardless of that item's continued visibility. In all of the experiments reported here, motion eliminated this identity. That is, conditions in which items remained visible produced slower and generally less efficient search than when the item was physically removed. This was true even for very slow motion, where there was considerable overlap between initial and final locations (Experiment 4–5) and for relative motion, where the overall layout of the display remained constant (Experiment 6). We argue that this disruption demonstrates that the influence of past actions relies, to some extent, on a location-based frame of reference.

Clearly, this argument relies on the use of object motion to separate location-based and object-based reference frames. While this is a well-established strategy (Kahneman, Treisman, & Gibbs, 1992; Tipper, Weaver, Jerreat, & Burak, 1994), it is worth considering whether the addition of motion might also have affected performance in some other way. For example, suppose that motion simply increased the overall task difficulty. Could this have led to the observed divergence between Vanish and Remain conditions? In all of the current experiments the overall level of response times were comparable to those obtained with static arrays (Thornton & Horowitz, 2004). While we did note a slight rise in click-errors, (though not sequence-errors) there was never a substantial difference *between* the Vanish and Remain conditions, arguing against an underlying shift in task difficulty.

Could observers have accomplished performance levels comparable to those obtained with static arrays by strategically shifting limited resources away from memorizing prior target locations? Under this scenario, observers must have had sufficient remaining resources to

support prospective memory in Experiments 1–2. In Experiment 3, however, we prevented observers from using prospective memory. If strategic shifting of resources were a feature of the current experiments, then observers might have been expected to free up resources devoted to prospective memory and shift them to retrospective memory. However, observers showed no more evidence for retrospective memory in Experiment 3 than in previous experiments. Finally, note that varying the speed of motion in Experiments 4 and 5 had no effect on retrospective memory. This lack of modulation is inconsistent with the notion that motion disrupts retrospective memory merely by increasing task difficulty.

Another more plausible alternative explanation for the impact of motion involves the overall density of the displays. Our measure of retrospective memory is the difference between the Remain and Vanish conditions. Since the Vanish condition is defined by a reduction in the number of items in the display over time, any advantage in comparison to the Remain condition might conceivably be explained by reduced crowding effects (Pelli, Palomares, & Majaj, 2004). With static displays, this was not a concern, since the striking finding was that we observed no difference between conditions (Thornton & Horowitz, 2004). In the current experiments, we did observe an advantage for the Vanish condition. Furthermore, while the interstimulus distances at the start of a trial were large enough to make crowding effects unlikely, the motion algorithm allowed objects to come closer, and even overlap.

Could crowding effects explain the difference between Vanish and Remain conditions in these experiments? Two findings suggest otherwise. First, slowing the speed of motion in Experiments 4 and 5, which would reduce the number of close encounters, had no effect on the Vanish/Remain difference. More tellingly, we still observed a difference in Experiment 6, in which the interstimulus distances remained constant. Crowding does not provide a very compelling explanation for the observed Vanish/Remain differences.

The most compelling explanation is that adding motion to the MILO task affects retrospective memory by disrupting location-based information. In the Introduction that we had predicted that both object-based and location-based coding might be used to support performance in the MILO task. Such dual-coding has been reported for a number of other selection-related effects, such as negative priming (Shapiro & Loughlin, 1993) and IOR (McAuliffe, Pratt, & O'Donnell, 2001; Tipper, Jordan, & Weaver, 1999; Tipper, Weaver, Jerreat, & Burak, 1994). Is there any indication that some form of retrospective advantage persists, thus hinting at some role for object-based coding?

In Experiments 4–6, the SRT function for the Remain condition is essentially flat, with each subsequent response requiring an equal amount of time. Such a pattern indicates that the effective set size in the Remain condition is constant and that previously located targets are still interfering with ongoing responses. In these experiments, then, there appears to be no suggestion of residual retrospective coding.

In Experiments 1–3, however, even though there were clear differences between conditions, the Remain slope was not flat. That is, responses occurring later in the sequence were speeded relative to earlier responses, even when items remained visible. Could this be evidence for residual object-based tagging? While we cannot rule this out, there may be another explanation. We noted in our original paper (Thornton & Horowitz, 2004), that some component of the retrospective effect could be explained from a response competition point of view. Imagine that the motor system prepares a response program for each stimulus (Tipper, Howard, & Jackson, 1997; Welsh & Elliott, 2004). At the start of the trial, all of these programs compete to control the response, leading to long latencies. As responses are fired off, fewer programs are competing, and latencies decrease. Such a reduction in competition could thus give rise to non-zero Remain slopes even in the absence of explicit item tagging.

Such an idea might also help explain the shapes of the $SRT \times \text{target}$ functions. Visual search for characters in heterogeneous sets generally produces linear $RT \times \text{set size}$ functions (Schneider & Shiffrin, 1977), so subtracting items should produce linear decreases in latency. However, in number of the current experiments (1,3,5,&6), as well as in our previous work (Thornton & Horowitz, 2004), even the Vanish conditions produce nonlinear decreases. These non-linear patterns may thus indicate the involvement of processes other than search itself. Selective reaching theorists are generally more concerned with trajectories than with latency, so their theories do not directly predict the $SRT \times \text{target}$ functions. Nevertheless, response competition seems more likely to produce nonlinear functions than search.

What is still unclear in the current paper is why the Remain slope is essentially flat in some experiments and not others. This may be evidence for some imperfect object-based inhibitory tagging, or response competition. However, it is difficult to speculate further, since we can find no systematic difference between those conditions which produce some Remain slope and those which produce none. The magnitude of the Remain slope does not correlate with speed, as can be seen in Experiment 4. We also observed steep Remain slopes both with (Experiments 3–4) and without (Experiments 1–2) the shuffle manipulation. Further research will be necessary to resolve this puzzle.

Motion clearly disrupts retrospective memory, and does so regardless of the speed of motion (Experiments 4–5) or the presence or absence of a global context (Experiment 6). Reducing the speed and introducing rigid motion improved performance (in both the RT and accuracy domains), but neither manipulation reinstated the ability to ignore old targets. Thus, the precise *location* of old targets seem to be quite crucial to performance. A straightforward interpretation of these data is that retrospective tagging of old targets occurs relatively early in visual processing, perhaps at the level of the salience map (Itti & Koch, 2001; Li, 2002) or activation map (Wolfe, Cave, & Franzel, 1989) thought to guide the deployment of attention in visual search. These maps are generally assumed to code locations, rather than objects, and retrospective memory could easily be implemented by inhibiting locations on such a map after a successful search (similar to how some models proposed to inhibit rejected distractor locations, e.g. Wolfe, 1994).

Interestingly, previous mechanisms that have been proposed to influence search by inhibiting old items on such maps, for example IOR (Klein, 1988; Klein & MacInnes, 1999; Shore & Klein, 2000) and visual marking (D. G. Watson & Humphreys, 1997), are also known to exploit object-based coding (McAuliffe, Pratt, & O'Donnell, 2001; Tipper, Jordan, & Weaver, 1999; Tipper, Weaver, Jerreat, & Burak, 1994; D. G. Watson & Humphreys, 1998). Here, we found little evidence for such object-based coding. As the MILO task crucially involve both search and response components, this difference could reflect the additional involvement of the motor system. That is, the relevant maps may be more location-based because of the response component. Consistent with this idea, previous research has shown that response programs in cluttered displays are coded exclusively in location-based coordinates (Keulen, Adam, Fischer, Kuipers, & Jolles, 2002).

We turn now to the prospective aspects of the data. The most obvious indication of such planning in the MILO task is the large difference between the speed of response to an initial item compared to later items. Such a pattern is clearly visible in the data of Experiments 1–2 of the current paper and, as with static arrays, can be eliminated by introducing a shuffle manipulation (Experiment 3). Clearly then, motion does not eliminate prospective planning in the same way that it disrupts retrospective effects. This suggests that in terms of future planning items may be tagged and tracked within an object-based, rather than location-based frame of reference. For example, an observer may come across target $n+1$ during the search for target n . Target $n+1$ is recognized as a future target, but the observer cannot respond to it out of

sequence. The object is therefore tagged as a future target so it can be easily recovered. As we suggested in the introduction, such tagging might be accomplished via a spatiotemporal indexing mechanism such as Pylyshyn's (1989) FINSTs. Such indexes are intrinsically object-based and would provide a natural substrate for this prospective memory.

In the current paper, we have focused more on the retrospective aspects of the MILO task. Clearly, an important direction for future research will be to more fully investigate the prospective aspects of the task. The sparing of prospective memory in dynamic arrays suggests a number of interesting research questions. With static arrays, we have shown planning effects in the MILO task that extend beyond the next item. We did this by leaving item $n+1$ or $n+2$ etc., stable, and shuffling further and further in advance. How far ahead might observers be able track targets in the context of motion? Also, spatiotemporal indexing is known to be a capacity-limited mechanism (Trick & Pylyshyn, 1993). If we systematically vary the target set size, will we also see capacity limits on prospective memory?

In the typical MOT experiment, observers need to distinguish between two classes of moving objects: targets and non-targets. They usually do not need to discriminate among members of the target class. This task is thus more analogous to retrospective memory in our experiments, where the visual system needs to segregate old targets from other items. However, in prospective memory, observers presumably need to keep track of which upcoming target is which, since their responses must be ordered. Recent work suggests that capacity limits for this sort of task are more severe than on the segregation task.

Pylyshyn (2004) asked observers to track 4 targets out of 8 moving objects. In one experiment, targets started off in different corners of the display, in another they had digits presented briefly at the start of the trial. In both cases observers could correctly discriminate targets from non-targets, without knowing which target was which. Horowitz et al. (in press) quantified this discrepancy. They asked observers to track 4 of 8 unique cartoon animals. At the end of the trial, the animals were masked by cartoon cactuses, and observers were either asked to indicate the locations of all the targets, or locate a particular target (e.g. the zebra). Observers could not adopt different strategies for different questions, since they did not know which question they would be asked until the end of the trial. Nevertheless, capacity was between three and four objects when observers were asked to segregate targets from non-targets, and one to two objects when asked to discriminate among targets.

Converging evidence comes from work with the multiple object permanence tracking paradigm (Saiki, 2003a, 2003b) and the object file reviewing paradigm (Henderson, 1994; Henderson & Anes, 1994). In multiple object permanence tracking (Saiki, 2002), a set of colored shapes rotate behind a windmill-shaped occluder. The observer's task is to detect two objects exchanging colors. The change itself is obscured behind the occluder, so the observer must rely on object feature binding. Saiki finds that observers can maintain only one or two bound objects (Saiki, 2003a, 2003b). In object file reviewing (Kahneman, Treisman, & Gibbs, 1992) experiments, observers see two displays, a preview display and a target display. Each display consists of a number of objects, such as boxes, containing letters. Typically, the preview and target displays are linked in such a way as to ensure that objects in the target display are seen as continuations of objects in the preview display. Observers are required to name the letters in the target display. Nonspecific priming occurs when letters from the preview show up anywhere in the target display, while object-specific priming occurs when a letter is presented in the same object in both displays. Henderson (1994; Henderson & Anes, 1994) showed evidence for both types of priming, and argued that they reflected two different representational systems, *object types* and *object files*, latter being a limited capacity system capable of handling approximately two objects. Together, these results suggest that prospective memory could be quite limited.

Another important question for future research concerns response modality. We asked observers to move a mouse pointer. The results might have been different had we used saccadic responses. Fischer, Pratt, and Neggers (2003) have demonstrated that motor IOR is limited to the oculomotor domain, and not observed with pointing movements. This suggests that it may be useful to study the MILO task with real pointing responses and with oculomotor responses.

In any event, the apparent dissociation between prospective and retrospective memory suggested by the current experiments may provide a very powerful tool. That is, we can use motion to block retrospective effects, sparing prospective effects, and the shuffle manipulation to block prospective effects, sparing retrospective performance. Such a dissociation may allow us to shed further light on the nature of the mechanisms underlying these two aspects of the task using further behavioural or neuroimaging techniques.

References

- Abrams RA, Law MB. Object-based visual attention with endogenous orienting. *Perception & Psychophysics* 2000;62(4):818–833. [PubMed: 10883587]
- Alvarez GA, Franconeri SL. How many objects can you track? Evidence for a flexible tracking resource. *Journal of Vision* 2005;5(8):641.[Abstract]
- Cavanagh P, Alvarez GA. Tracking multiple targets with multifocal attention. *Trends Cogn Sci* 2005;9(7):349–354. [PubMed: 15953754]
- Danziger S, Kingstone A, Ward R. Environmentally defined frames of reference: their time course and sensitivity to spatial cues and attention. *Journal of Experimental Psychology: Human Perception & Performance* 2001;27(2):494–503. [PubMed: 11318063]
- Davis G, Welch VL, Holmes A, Shepherd A. Can attention select only a fixed number of objects at a time? *Perception* 2001;30(10):1227–1248. [PubMed: 11721824]
- Duncan J. Selective attention and the organization of visual information. *Journal of Experimental Psychology: General* 1984;113(4):501–517. [PubMed: 6240521]
- Fischer MH, Pratt J, Neggers SF. Inhibition of return and manual pointing movements. *Perception & Psychophysics* 2003;65(3):379–387. [PubMed: 12785068]
- Ghafari M, Archambault PS, Adamovich SV, Feldman AG. Pointing movements may be produced in different frames of reference depending on the task demand. *Brain Research* 2002;929(1):117–128. [PubMed: 11852038]
- Henderson JM. Two representational systems in dynamic visual identification. *Journal of Experimental Psychology: General* 1994;123(4):410–426. [PubMed: 7996123]
- Henderson JM, Anes MD. Roles of object-file review and type priming in visual identification within and across eye fixations. *Journal of Experimental Psychology: Human Perception & Performance* 1994;20(4):826–839. [PubMed: 8083637]
- Horowitz TS, Klieger SB, Fencsik DE, Yang KK, Alvarez GA, Wolfe JM. Tracking unique objects. *Perception & Psychophysics*. (in press)
- Itti L, Koch C. Computational modelling of visual attention. *Nature Reviews Neuroscience* 2001;2(3):194–203.
- Kahneman D, Treisman A, Gibbs BJ. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 1992;24(2):175–219. [PubMed: 1582172]
- Keulen RF, Adam JJ, Fischer MH, Kuipers H, Jolles J. Selective reaching: evidence for multiple frames of reference. *Journal of Experimental Psychology: Human Perception & Performance* 2002;28(3):515–526. [PubMed: 12075885]
- Klein RM. Inhibitory tagging system facilitates visual search. *Nature* 1988;334:430–431. [PubMed: 3405288]
- Klein RM, MacInnes WJ. Inhibition of return is a foraging facilitator in visual search. *Psychological Science* 1999;10(4):346–352.

- Lamy D, Tsal Y. Object features, object locations, and object files: which does selective attention activate and when? *Journal of Experimental Psychology: Human Perception & Performance* 2000;26(4):1387–1400. [PubMed: 10946721]
- Li Z. A saliency map in primary visual cortex. *Trends in Cognitive Sciences* 2002;6(1):9–16. [PubMed: 11849610]
- Lofthus GRM, Michael EJ. Using confidence intervals in within-subject designs. *Psychological Bulletin & Review* 1994;1(4):476–490.
- McAuliffe J, Pratt J, O'Donnell C. Examining location-based and object-based components of inhibition of return in static displays. *Perception & Psychophysics* 2001;63(6):1072–1082. [PubMed: 11578051]
- Pelli DG, Palomares M, Majaj NJ. Crowding is unlike ordinary masking: distinguishing feature integration from detection. *J Vis* 2004;4(12):1136–1169. [PubMed: 15669917]
- Pelli DG, Zhang L. Accurate control of contrast on microcomputer displays. *Vision Research* 1991;31(7–8):1337–1350. [PubMed: 1891822]
- Pylyshyn ZW. The role of location indexes in spatial perception: a sketch of the FINST spatial-index model. *Cognition* 1989;32(1):65–97. [PubMed: 2752706]
- Pylyshyn ZW. Some puzzling findings in multiple object tracking (MOT): I. Tracking without keeping track of object identities. *Visual Cognition* 2004;11(7):801–822.
- Pylyshyn ZW, Storm RW. Tracking multiple independent targets: evidence for a parallel tracking mechanism. *Spatial Vision* 1988;3(3):179–197. [PubMed: 3153671]
- Rensink RA. Toolbox-based routines for Macintosh timing and display. *Behavior Research Methods, Instruments, & Computers* 1990;22(2):105–117.
- Saiki J. Multiple-object permanence tracking: limitation in maintenance and transformation of perceptual objects. *Progress in Brain Research* 2002;140:133–148. [PubMed: 12508587]
- Saiki J. Feature binding in object-file representations of multiple moving items. *J Vis* 2003a;3(1):6–21. [PubMed: 12678621]
- Saiki J. Spatiotemporal characteristics of dynamic feature binding in visual working memory. *Vision Res* 2003b;43(20):2107–2123. [PubMed: 12855247]
- Schendel KL, Robertson LC, Treisman A. Objects and their locations in exogenous cuing. *Perception & Psychophysics* 2001;63(4):577–594. [PubMed: 11436731]
- Schneider W, Shiffrin RM. Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review* 1977;84(1):1–66.
- Shapiro KL, Loughlin C. The locus of inhibition in the priming of static objects: object token versus location. *Journal of Experimental Psychology: Human Perception & Performance* 1993;19(2):352–363. [PubMed: 8473844]
- Shore DI, Klein RM. On the manifestations of memory in visual search. *Spatial Vision* 2000;14:59–75. [PubMed: 11334182]
- Steinman SB, Nawrot M. Real-time color-frame animation for visual psychophysics on the Macintosh computer. *Behavior Research Methods, Instruments, & Computers* 1992;24(3):439–452.
- Thornton IM, Horowitz TS. The Multi-Item Localization (MILO) task: Measuring the spatio-temporal context of vision for action. *Perception & Psychophysics* 2004;66(1):38–50. [PubMed: 15095938]
- Tipper SP, Driver J, Weaver B. Object-centred inhibition of return of visual attention. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology* 1991;43A(2):289–298.
- Tipper SP, Howard LA, Jackson SR. Selective reaching to grasp: Evidence for distractor interference effects. *Visual Cognition* 1997;4(1):1–38.
- Tipper SP, Jordan H, Weaver B. Scene-based and object-centered inhibition of return: evidence for dual orienting mechanisms. *Perception & Psychophysics* 1999;61(1):50–60. [PubMed: 10070199]
- Tipper SP, Weaver B, Jerreat LM, Burak AL. Object-based and environment-based inhibition of return of visual attention. *Journal of Experimental Psychology: Human Perception & Performance* 1994;20(3):478–499. [PubMed: 8027711]
- Trick LM, Pylyshyn ZW. What enumeration studies can show us about spatial attention: evidence for limited capacity preattentive processing. *Journal of Experimental Psychology: Human Perception & Performance* 1993;19(2):331–351. [PubMed: 8473843]

- Watson DG, Humphreys GW. Visual marking: Prioritizing selection for new objects by top-down attentional inhibition of old objects. *Psychological Review* 1997;104(1):90–122. [PubMed: 9009881]
- Watson DG, Humphreys GW. Visual marking of moving objects: A role for top-down feature-based inhibition in selection. *Journal of Experimental Psychology: Human Perception and Performance* 1998;24(3):946–962. [PubMed: 9627427]
- Watson SE, Kramer AF. Object-based visual selective attention and perceptual organization. *Perception & Psychophysics* 1999;61(1):31–49. [PubMed: 10070198]
- Welsh T, Elliott D. Movement trajectories in the presence of a distracting stimulus: evidence for a response activation model of selective reaching. *Quarterly Journal of Experimental Psychology A* 2004;57(6):1031–1057.
- Wolfe JM. Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review* 1994;1(2):202–238.
- Wolfe JM, Cave KR, Franzel SL. Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception & Performance* 1989;15(3):419–433. [PubMed: 2527952]

Acknowledgments

We would like to thank Randall S. Birnkrant, Sarah B. Klieger, Andries Hof and Jessica Wilner-Reid for assistance with data collection and analysis. TSH was supported by NIH grant MH65576, and IMT by the Max Planck Institute for Biological Cybernetics, Tübingen, Germany. We would also like to thank Steve Franconeri and two anonymous reviewers for helpful comments on the manuscript.

Figure 1.
Example stimulus displays from the MILO task.

Figure 2.

Results from the MILO paradigm using static displays (redrawn from Thornton & Horowitz, 2004). SRTs are plotted as a function of target number and condition. Open symbols denote the “vanish” condition, filled symbols the “remain” condition (see text). Panel A depicts the results when stimulus positions remained constant throughout a trial. Panel B depicts the results when we shuffle the locations of distractors and targets that have not yet been responded to.

Figure 3.

Left-hand panel plots SRT medians from Experiment 1 as a function of target number. Filled symbols denote the Remain condition, open symbols the Vanish. Right-hand panel plots percent error as a function of error type. Gray bars denote the Remain condition, open bars the Vanish. In this and all subsequent figures, error bars indicate within-subjects 95% confidence intervals (Loftus, 1994).

Figure 4. Left-hand panel plots SRT data from Experiment 2 as function of condition and target number. Filled symbols denote the Remain condition, open symbols the Vanish. Right-hand panel plots percent error as a function of error type. Gray bars denote the Remain condition, open bars the Vanish.

Figure 5.

The left-hand panel plots median SRT data from Experiment 3 as a function of condition and target number. Solid symbols denote data from the Remain condition, open symbols data from the Vanish condition. Right hand panel plots percent error for the two error types. Gray bars denote the Remain condition, open bars the Vanish.

Figure 6.

Data from Experiment 4. The top two panels plot SRTs as a function of speed and target number, with the Remain condition observers shown in Panel A and the Vanish condition observers shown in Panel B. Panel C depicts % error as a function of speed for the Remain (filled symbols) and Vanish (open symbols) groups. Click errors are plotted as circles, sequence errors as squares. Panel D plots the slope of the SRT $\times$ target functions for the Remain (filled symbols) and Vanish (open symbols) groups.

Figure 7.

Right-hand panel plots SRT data from Experiment 5 as a function of target number. Filled symbols denote the Remain condition, open symbols denote the Vanish condition. Left-hand panel plots percent error as a function of error type. Gray bars denote the Remain condition, open bars the Vanish.

Figure 8.

Right-hand panel plots SRT data from Experiment 6 as a function of condition and target number. Filled symbols denote the Remain condition, open symbols denote the Vanish condition. Left-hand panel shows percent error as a function of error type and condition. Gray bars denote the Remain condition, open bars the Vanish.