Resource demands of object tracking and differential allocation of the resource

Wei-Ying Chen^{1,*}, Piers D. Howe², & Alex O. Holcombe¹

¹School of Psychology, University of Sydney, Brennan MacCallum Bldg (A18), Sydney, NS

W 2006, Australia

²School of Psychological Sciences, University of Melbourne, 12th Floor Redmond Barry

Building, Parkville, VIC 3010, Australia

* Corresponding author: Tel: +61-4-23814138

E-mail address: wingchen0125@gmail.com

Abstract

The attentional processes for tracking moving objects may be largely hemisphere-specific. Indeed, in our first two experiments the maximum object speed (speed limit) for tracking targets in one visual hemifield (left or right) was not significantly affected by a requirement to track additional targets in the other hemifield. When the additional targets instead occupied the same hemifield as the original targets, the speed limit was reduced. At slow target speeds however, adding a second target to the same hemifield had little effect. At high target speeds, the cost of adding a same-hemifield second target was approximately as large as would occur if, observers could only track one of the targets. This shows that performance for a fast-moving target is very sensitive to the amount of resource allocated. In a third experiment, we investigated whether tracking resources can be distributed unequally between two targets. The speed limit for a given target was higher if the second target was slow than if the second target was fast, suggesting that more resource was allocated to the faster of the two targets. This was statistically significant only for targets presented in the same hemifield, consistent with the theory of independent resources in the two hemifields. Some limited evidence was found for resource sharing across hemifields, suggesting that attentional tracking resources may not be entirely hemifield specific. Together, these experiments indicate that the largely hemisphere-specific tracking resource can be differentially allocated to faster targets.

Introduction

Playing a team sport, eluding a group of predators, or taking children to the beach benefits from maintaining attention on multiple moving objects. The multiple object tracking (MOT) task (Pylyshyn & Storm, 1988) has been widely used to study this process (Cavanagh & Alvarez, 2005; Scholl, 2009). In this task, a number of identical objects are presented and a target subset to be tracked is cued for a few seconds. The cues then disappear so that the targets are identical to the non-targets and all objects move about the screen for several seconds. At the end of the trial, all the objects stop moving and observers must indicate which objects were the targets. With commonly used display parameters, people typically succeed at tracking up to four or five targets (Pylyshyn & Storm, 1988; Yantis, 1992; Alvarez & Cavanagh, 2005).

The number of objects that people can track varies with circumstances, and various theories have been proposed to explain the varying limits. Franconeri and his collaborators suggested that spatial interference is the *only* factor that limits the number of targets that can be tracked, writing that “barring object-spacing constraints, people could reliably track an unlimited number of objects as fast as they could track a single object” (p.924, Franconeri et al., 2010). Decrease in tracking performance with additional targets was proposed to be due to the cortical representations of two nearby targets interfering with each other (Franconeri, in press). Specifically, when two targets are close to each other, a suppressive surround of one

target may overlap with the spotlight of attention focused on the other target and vice versa, inducing worse tracking performance in both targets (Franconeri et al, 2008; in press). Spatial interference undoubtedly can contribute to the capacity limit when traditional MOT displays are used, because in those displays objects pass very close to each other, which can cause crowding (Intriligator & Cavanagh, 2001; Pelli & Tillman, 2008).

Evidence from widely-spaced target configurations indicates a second reason for the decrease in performance with increasing target loads. Specifically, that tracking a target requires allocating some process or entity, of which the brain has a finite amount or number. We will use the term “resource theory” broadly to refer to any such theory positing that tracking targets reduces the process available for other targets. This is in contrast to the spatial interference theory of Franconeri and colleagues (2008; 2010).

The spatial interference theory of Franconeri et al. proposes that each target is tracked equally well regardless of the number of targets, providing the targets are all sufficiently well separated. The limitation on number of targets that can be tracked arises due to target-target interactions. These cannot be avoided when there is not enough space in the visual field to keep all the targets widely spaced.

The resource of resource theories may be a finite set of discrete pointers (sometimes called “slots”) that are assigned to the targets (Pylyshyn & Storm, 1988; Horowitz & Cohen, 2010) or a continuous pool of mental resource that is divided among the targets (Alvarez &

Franconeri, 2007). For the continuous models, there are a number of possibilities for how reducing the resource per target (by increasing the number of targets one must track) impairs performance. One possibility is that allocating less resource to a target reduces the temporal precision with which each target is tracked. That theory is consistent with the results of Holcombe & Chen (in press) found that temporal frequency or speed limits on tracking decreased in a situation where the targets never crossed the paths of the other targets. Such evidence rules against a Franconeri-style non-resource theory whereby temporal limits occurred only due to suppression that endures at each location the target visits, impairing the tracking of targets that enter the area subsequently. Instead, the Holcombe and Chen (in press) finding of interference between targets and distractors is more consistent with a resource account.

Another theory that for present purposes is termed a resource theory is the serial switching theory. On this account, only one slot or spotlight is available to track targets, and it must be rapidly switched among targets for tracking to succeed (Tripathy, Ögmen, & Narasimhan, 2011; Tripathy & Howard, 2012; Holcombe & Chen, in press; but see Howe et al., 2010). When more targets are present, then, each target receives proportionally less processing time.

The important point that these resource theories have in common is that adding targets reduces performance per target even if no interactions among targets — spatial or temporal

interference — occurs. In our first two experiments, we validate this prediction of resource theories. We avoid spatial interference among targets by spacing the targets widely, and avoid temporal interference by not allowing the targets to cross paths. We find that performance nonetheless decreases with tracking loads. Evidence for this was already provided by Holcombe & Chen (2012; in press) but here we perform additional tests using a display configuration that allows us to go on and test the possibility of differential allocation of the putative tracking resource.

Under resource theories it is conceivable that different targets, presented simultaneously, may be allocated different amounts of the resource. With the slot theory for example, if the resource comprises four discrete tracking slots and two targets are presented, three slots could be devoted to the more-demanding target, with just one slot devoted to the less-demanding target. If the resource is instead a continuously divisible pool, then the more-demanding target might be allocated 75% of the resource and the less-demanding target only 25%. If the resource is instead a unitary focus of attention time-shared among the targets, it might visit the more-demanding target more often. In summary, a flexible resource may be differentially allocable among targets, rather than split evenly between them. Our third experiment found evidence for this.

Previous evidence for variable resource allocation

Although numerous studies have found evidence for an attentional resource

component to tracking, there is little evidence available that speaks to the possibility of differential resource allocation. For example, Tombu & Seiffert (2008) found evidence that tracking depends on an attentional resource that is also required to perform auditory tone discrimination. They also found evidence that tracking demands more of this resource when the targets are moving quickly. However, they did not investigate whether one can allocate more of the resource to one target (or to one task) than to another.

Indirect support for differential resource allocation was found by Liu et al. (2005). In one of their experiments, half of the targets moved at 1 deg/s and the other half moved at 6 deg/s. They found that tracking accuracy was the same for both kinds of target, even though one would expect that if both received equal resources, then accuracy would be poorer for the faster targets. This was a null result however, and one they did not discuss or follow up.

Experiments by Iordanescu and colleagues (2009) also yielded some data that was interpreted as supporting differential resource allocation. Their observers viewed a number of moving discs and were asked to track a subset of them. All the discs were coloured and each target was a different colour. The objects moved about randomly and at the end of the trial, all the discs disappeared and the observers were asked to indicate the final location of a particular target (e.g. the red one). Observers were more accurate at indicating the final location of a target when the target was located near distractors. Iordanescu et al. (2009) suggested that this occurred because more resource is devoted to targets when they are near

distractors. But because they did not directly manipulate proximity but rather relied on the random movements, there may have been other display characteristics that differed when targets and distractors were nearby. Furthermore, two studies have failed to replicate this finding (Howard & Holcombe, 2008; Howard et al., 2011). The reasons for the failure to replicate are unclear, particularly as the experiments by Howard et al. (2011) used a stimulus paradigm very similar to that employed by Iordanescu et al. (2009).

Howe et al. (2010, Experiment 8) performed a more direct test of whether the tracking resource could be differentially reallocated between targets during tracking. In their displays, each object would repeatedly pause so that it was moving for only half the tracking period. In the simultaneous condition, all the objects moved and paused simultaneously (i.e., synchronously). In the sequential condition, the objects were divided into two groups, each with an equal number of targets, and the two groups moved in alternation. When the objects in one group were moving, those in the other group were stationary. The rationale was that when an object was not moving, it would require less tracking resource (Alvarez & Franconeri, 2007; Bettencourt & Somers, 2009) and more resource could be allocated to the moving objects. Since fewer objects were moving at any one time in the sequential condition than in the simultaneous condition, it was expected that tracking performance would be greater in the sequential condition. In fact, tracking performance was equal for the two conditions, suggesting that the tracking resource could not be dynamically reallocated

between the targets. But to benefit performance in the Howe et al. (2010) study, any unequal distribution of resources would have to be reversed at the rate of the movement alternation.

Perhaps participants *can* allocate resource unequally but cannot change this resource allocation rapidly. Using a constant difference in speed, our Experiment 3 will test for the possibility of allocating more resource to the faster target.

Hemisphere specificity

The factors that limit tracking, such as target speed and number, appear to operate largely independently in the left and right visual fields, suggesting that if a resource does mediate tracking, there are two independent pools, one in each cortical hemisphere. Alvarez and Cavanagh (2005) presented a target in one visual hemifield and tested the effect of adding another target in the same hemifield or in the opposite hemifield. Performance was much poorer when the additional target was presented in the same visual hemifield, but not significantly affected when the second target was presented in the opposite hemifield. This suggests that the resource consumed by additional targets is hemisphere-specific.

A concern with the Alvarez and Cavanagh (2005) finding is that some or all of the decrease in performance with additional targets may have reflected spatial interference rather than resource depletion. In the Alvarez and Cavanagh (2005) study, the objects could pass very close to each other. Since spatial interference (crowding) is known to be greater when objects are presented unilaterally rather than split across hemifields (Liu, Jiang, Sun & He,

2009), the observed decrease in tracking performance when the targets were presented unilaterally might have been due to spatial interference.

Holcombe and Chen (2012) addressed this issue by using a display in which the objects were kept widely separated to avoid crowding. After finding the maximum speed at which observers could track one target, they tested how this speed limit was affected when observers were asked to track an additional target. The speed limit decreased substantially if both targets occupied the same hemifield. If the second target was instead placed in the opposite hemifield, little decrement in the speed limit occurred. This supports the original claim of an attentional resource that is independent in each hemifield. Here we will investigate whether the hemisphere-specific resource can be differentially allocated, by comparing the effect of targets of differing speeds placed in the same hemifield or in opposite hemifields.

The resource vs performance function

If tracking does require attentional resources, how much resource is needed for successful tracking of a single target? Answering this question fully would mean quantifying the function that maps proportion of resource allocated to a target onto percent correct. The resource vs. performance function (Figure 1) relates the proportion of attentional resource devoted to a target to the likelihood of getting it correct (Norman & Bobrow, 1975).

Ideally we would measure the resource vs. performance function by asking

participants on different trials to allocate different proportions of their attention to each of two targets, 90%:10%, 80%:20%, 70%:30%, etc. While that may be a valid method for simple judgments regarding briefly presented stimuli (Bonnell & Miller, 1994; Lee, Koch, & Braun, 1999; Pastukhov, Fischer, & Braun, 2008), we submit that it would be difficult to induce participants to allocate a particular proportion of attention to two targets throughout a tracking trial.

Participants commonly report that they know when they fail to maintain their attention on a particular target for tracking. These reports are validated by the empirical success of the method of adjustment in tracking studies (Verstraten, Cavanagh, & LaBianca, 2000; Vul et al., 2009), which requires that participants recognize when they succeed or fail to track the designated targets. Using this knowledge, during a trial participants may shift the resource formerly used for a lost target to one of the other targets. It seems unlikely that, even if explicitly instructed to allocate (say) 30% of their resources to a target, that upon losing the other target the participant would leave 70% of their resources thereafter unused. In a further complication, towards the beginning of a trial participants may recognize that the targets' speeds are too fast for them all to be tracked and then shift all their resource towards a subset of the targets. These possibilities of strategic allocation and reallocation depending on the characteristics of targets may make it difficult to enforce a particular allocation proportion.

We can nevertheless glean some limited information regarding the resource vs.

performance function. Two points on the function are already specified by the definition of resource theory. When only one target is tracked, 100% of the resource is devoted to it and performance is at maximum. When no resource is available per target (effectively, when a very large number of targets are specified), performance should be at or very near chance. The simplest possible resource vs. performance function would then connect these two points with a straight line, which we will refer to as the linear resource function (see Figure 1). This linear function then predicts that when two targets are tracked (50% resource is allocated to each target), performance will be approximately halfway between the one-target level and chance.

In the literature on simple psychophysical judgments of briefly-presented stimuli, data supporting an approximately linear performance vs. resource function was found for several concurrent discrimination tasks by Braun and colleagues (Lee, Koch, & Braun, 1999; Pastukhov, Fischer, & Braun, 2008).

A different resource function for tracking was proposed by Horowitz & Cohen (2010). Horowitz & Cohen assessed participants' ability to report the direction that tracked targets were moving, after the display stopped. They found evidence that performance matched the resource function predicted from a noisy independent samples model of the resource. Specifically, the theory is that resource improves performance by improving the precision of tracking in the same way that increasing the number (n) of noisy samples taken from a

distribution improves the precision (standard deviation) of the estimate of the mean.

Specifically, the standard deviation improves with the $\sqrt{n}$. Note that the dependent variable in the Horowitz & Cohen (2010) analysis was the standard deviation of the reports of targets' final motion direction.

The relationship between precision and proportion correct in tracking targets is uncertain. If linear, then the noisy independent samples theory predicts that proportion correct will increase with $\sqrt{n}$, as schematized in Figure 2. The corresponding curve (dashed line) in the graph is simply that for which performance grows in proportion to $\sqrt{\text{resource}}$ with the additional constraints that zero resource yields chance performance and 100% resource yields the one-target performance. We don't favor this theory (see General Discussion), but it usefully provides a competing possible resource vs. performance function. Moreover, for a length discrimination task, some evidence for this linear relationship between proportion correct and resource allocated was found by Bonnel & Miller (1994) and it was considered for short-term memory by Zhang & Luck (2008).

In this paper we will only be able to assess the performance vs. resource function empirically at the 50% resource (2 targets) point. We find that at slow speeds, performance exceeds both the linear and independent samples function whereas the function is very different for fast speeds. For fast speeds, performance falls at or below the linear function (as schematized in Figure 2). This implies that participants did no better at fast speeds than if

they had tracked only one target and ignored the other, as explained below.

The capacity-one performance benchmark

The capacity-one performance benchmark, first described by Holcombe and Chen (2012, which they termed the “capacity-one model”), calculates the performance level expected if participants track only one target and completely ignore the second target. This puts any cost of splitting attention in perspective by comparing the cost to what would occur if participants could only track one object. The rationale is that if participants track only one object, then on half of trials their performance will be the same as the one-target condition and on the other half of trials they will guess.

For the present experiments, the capacity-one performance benchmark yields the same performance level as the linear performance vs. resource function. Therefore, when empirical performance falls to this level (as we find for high target speeds), we cannot say whether it is because the linear function is correct or because participants only tracked a single target. It may be that at high speeds the resource requirement for successful tracking is even greater than that indicated by the linear function, and participants switch to tracking only one target as that yields better performance than attempting to track both.

The present experiments

In the first two experiments we sought to confirm that with the present display configuration, a hemifield-specific resource mediated tracking. After finding support for this

hypothesis, we then conducted an experiment to test whether more resource was allocated to the faster of two targets.

As described in the *Hemifield specificity* section above, Holcombe and Chen (2012) found evidence for hemisphere specificity of the resource in their Experiment 3. In that experiment however, the duration of the trials did not differ for different speeds, meaning that for trials testing fast target speeds, the targets traveled much further than on slow-speed trials. Franconeri et al. (2010) pointed out that with this type of design, high-speed trials may be associated with more spatial interference because the targets pass relatively near each other on more occasions. In the Holcombe and Chen (2012) display configuration, the objects were always far from each other, so this explanation seems unlikely, but nevertheless we sought to exclude it here. In Experiments 1 and 2 we therefore equated the distance traveled across speeds. The results supported the hemifield-specific resource theory, so in Experiment 3 we proceeded to investigate the possibility of differential allocation of the resource.

Experiment 1: Testing hemifield specificity of the tracking resource

Holcombe and Chen (2012) in their Experiment 3 found that the maximum target speed that could be tracked (68% threshold) with two targets in a hemifield was no better than expected if participants had ignored one of the targets and simply guessed whenever it was probed (the capacity-one performance benchmark). An alternative theory that makes the same prediction is that the resource vs. performance function is linear. These explanations

suggest that successful tracking (>68% accuracy) at high speeds requires more than 50% of the resource.

An alternative explanation is that performance was much worse for the two-target condition at high speeds because of spatial interference between the targets. The opportunity for spatial interactions may have been higher when the stimulus was presented at high speeds, as it traveled further during the trial (Franconeri et al., 2010). If instead we equate the distance travelled for different speeds, the putative spatial interference should not be higher at faster speeds. According to the spatial interference theory, then, the two-target speed limit cost associated with additional targets should be much reduced or disappear.

Methods

Five participants (four male, aged 24-31) who reported normal or corrected-to-normal vision agreed to participate, following approval of the protocol by the University of Sydney's ethics committee. One was the first author. All had extensive experience fixating in laboratory experiments.

A 120 Hz CRT displayed eight red discs (evoked by the red gun only, with Gaussian intensity profiles; visible diameter 1 deg; peak luminance 20 cd/m²) and a white fixation point against a black background, at a distance of 57 cm from the eyes. The spatial arrangement of the discs is schematized in Figure 3. Four pairs of discs were presented in all conditions. Each pair moved on a circular trajectory centred in one of the four quadrants of

the visual field. The circular trajectories were centred on points 6 deg from the vertical midline and 6 deg from the horizontal midline (representing the centres of the four quadrants). The radius of each trajectory was 2.5 deg of visual angle, and because the two discs on a trajectory were always diametrically opposed, the separation between them was always 5 degrees of visual angle. These distances should be large enough to avoid crowding among the objects (Intriligator & Cavanagh, 2001; Pelli & Tillman, 2008).

To indicate which discs were targets, for the first 0.7s of the motion interval the color of the targets was white instead of red. Following this was the tracking period, during which all the objects were red. To prevent participants from predicting the final target positions from their initial positions and speeds, the discs occasionally reversed direction. Each pair of discs was independently assigned a series of reversal times, which succeeded each other at random intervals between 1.2 and 2 s. This resulted in more reversals for trials that had slow speeds. At the end of the trial, one pair of discs was indicated and the participants used the mouse to indicate which of the two discs was the target.

In the four-targets condition, one disc of each pair was designated as a target to be tracked. In the two-targets bilateral condition, the two target pairs were both above the fixation point in half the trials and both below in the other half. In the two-targets unilateral condition, the two target pairs were either both to the left of the fixation point or both to the right of the fixation point.

To avoid the possibility of more opportunities for spatial interference at higher speeds, the cumulative distance traveled by the discs was the same for all trials. This was achieved by setting the duration of the trial to a different value for each speed condition. All objects revolved about fixation at the same rate. Across trials, five rotation speeds (0.7, 1.0, 1.4, 1.7, and 2.2 revolutions per second) were used, and to achieve a constant distance travelled of 6.6 revolutions, this yielded five corresponding tracking durations (9.4, 6.6, 4.7, 3.9, 3 s).

Each observer participated in 160 trials for each of five rates, yielding 800 experimental trials in total, divided into five sessions. Conditions were mixed, each observer did no more than two sessions a day, and observers had a minimum break between sessions of 5 minutes.

Data analysis

Plots of speed versus proportion correct were fit by a logistic regression that spanned from chance (50% accuracy) to a ceiling level of performance. The ceiling performance corresponds to the lapse rate, which in the fitting procedure was allowed to vary from 1% to 10% to get the best estimate. This estimated lapse rate for each condition is reported in the results section. We refer to the speed at which performance is estimated by the regression to fall to 68% correct as the “speed limit”. The regression was fit separately for each participant and condition to estimate the speed limits.

Holcombe and Chen (2012) calculated the expected effect on speed limits of increasing the number of targets, under the seemingly worst-case assumption that observers could track only one target per hemifield, and had to guess on the trials where they were queried on the untracked target. In fact this capacity-one performance benchmark is not the worst-case scenario, because the resource vs. performance function might fall below the linear function (Figure 2). In that case if participants attempt to track both targets, performance will fall below the capacity-one benchmark.

Calculation of the prediction of the hemisphere-specific capacity-one performance benchmark in the present circumstances proceeded as follows. In each visual hemifield it is assumed that the observer tracks only one target. Thus, in half the trials the observer will by chance have tracked the target in the pair that is queried. For these trials the performance for that speed is provided by the bilateral two-target logistic curve fit as this corresponds to the situation where there is only one target in each hemifield. On the other half of the trials the observer will have been queried about one of the pairs that he had not tracked. Consequently, the observer will be forced to guess which of the two discs is the target and therefore performs at chance (50%). From the resulting psychometric function, we can extract the benchmark's speed limit (68% threshold) for the four-target condition. This benchmark speed limit is shown as the red dashed bar in the lower half of Figure 4.

For slow target speeds, actual performance for tracking two targets is higher than the

capacity-one benchmark, in the present data as well as in that of Holcombe and Chen (2012) and Holcombe & Chen (in press). This shows that at slow speeds, participants can track more than one target in each visual hemifield. At high speeds however, actual performance is similar to the benchmark. Previously assessing the speed at which performance fell to the “speed limit” level of performance (68% correct), Holcombe and Chen (2012) found that the capacity-one benchmark speed limit was not significantly different than that of the humans. This indicates that participants did no better than they would if they tracked just a single target in each visual hemifield.

If anything, the observed speed limits here were even worse than those of the capacity-one benchmark (this statistically non-significant trend was also observed by Holcombe and Chen, 2012, and Holcombe and Chen, in press). This indicates that devoting half of the resource to a target yields only very poor performance in each target. If dividing one’s resource between the targets frequently results in failure to successfully track any of them, one would be better off attempting to track only one, which may be the strategy participants sometimes adopted.

An alternative and unlikely hypothesis, but an instructive one for the contrasting prediction it makes, is that observers track objects independently in the upper and lower visual hemifields and can only track one in each. For this UVF/LVF capacity-one benchmark, in the four-targets condition, in half of trials performance is given by the unilateral two-target

condition, and by the chance level in the other half of trials. This benchmark's speed limit is shown by the blue dashed bar in the lower half of Figure 4.

Results and discussion

The data and fitted curves are shown for each participant in the top panel of Figure 4, with associated speed limits (68% thresholds) shown in the bottom panel. For two targets, consistent with the hemisphere-specific resource theory, the speed limit was better in the bilateral arrangement (2.05 rps, 32.21 deg/s) than in the unilateral arrangement (1.59 rps, 24.98 deg/s). This difference was statistically significant according to a paired t-test, $t(4)=3.557$, $p=0.024$, Cohen's $d=2.572$.

As predicted by hemisphere-specific resource theory, compared to the speed limit for tracking two targets in a single hemifield, adding two more targets in the opposite hemifield had no significant effect on the speed limit (1.59 rps versus 1.50 rps, $t(4)=1.85$, $p=0.138$, Cohen's $d=0.888$).

Compared to having one target in both the left and right hemifield, tracking a second in each hemifield was expected to halve the resource available per target, reducing performance. Consistent with this, the requirement to track an additional target significantly reduced the speed limit. The speed limit decreased from 2.05 rps (32.21 deg/s) to 1.50 rps (23.57 deg/s), $t(4)=3.329$, $p=0.029$, Cohen's $d=1.702$. This cost (0.5 rps, 7.86 deg/s) was significantly larger than the (non-significant) cost of adding targets in the opposite hemifield

described in the previous paragraph (0.09 rps, 1.41 deg/s), as indicated by a paired t-test on the speed limit differences, $t(4)=3.557$, $p=0.024$, Cohen's $d=2.572$.

The non-significant trend present for a poorer speed limit for the four targets condition (1.50 rps) than for the unilateral two-targets condition (1.59 rps) suggests that in the latter condition both the ipsilateral and contralateral hemispheres may have contributed to tracking the targets. A similar non-significant effect was observed in Experiment 3 of Holcombe and Chen (2012). These non-significant trends suggest that the tracking resource may not be 100% hemisphere-specific.

So far we have considered only the simple qualitative prediction of hemisphere-specific resource theory—that adding targets to the opposite hemifield has no effect whereas adding targets to the same hemifield impairs performance. Given that same-hemifield targets appear to load on the same resource, the question arises of how much resource is needed to accurately track a target. This cannot be measured directly, but we can compare performance to the prediction of the linear resource vs. performance function (the linear function in Figure 2). The capacity-one performance benchmark makes the same prediction for the cost of adding a second target, that performance will fall halfway to chance.

At low target speeds, performance is clearly much better than the capacity-one performance benchmark. Indeed at very low speeds participant performance is near 100% correct. The estimated lapse rate for the 4-target condition is .04, suggesting that the

psychometric function saturates at 96%. If participants could only track one object in each hemifield (the capacity-one performance benchmark) performance should never exceed 75% correct for four targets. This indicates that participants are capable of tracking all four objects when they move slowly.

Given that they did not ignore any of the targets, the performance level provides some information about the performance vs. resource function. Reducing the resource available for a target from 100% (one target per hemifield) to 50% (two targets per hemifield) has little effect on performance. The first reaction of many expert readers may be that this is a ceiling effect. That's our point. At slow speeds, tracking is very accurate (near ceiling) whether 50% or 100% resource is used (flat resource vs. performance function in this domain).

For high target speeds, we suspected that tracking would be resource intensive, meaning that adding a target to each hemifield would be costly. To put any speed limit cost in perspective, we calculated the speed limit of the capacity-one performance benchmark. This hemisphere-specific capacity-one benchmark speed limit is shown by the red dotted bar in the bottom of Figure 2. The measured speed limit for tracking four targets (1.50 rps, 23.57 deg/s) was not significantly different from the benchmark (1.57 rps, 24.67 deg/s), as revealed by a paired t -test, $t(4)=0.398$, $p=0.711$, Cohen's $d=-0.182$. Performance at these high speeds then was as bad as if participants simply ignored the second target in each hemifield.

We cannot tell whether participants indeed only attended to a single target in each

hemifield or instead attempted to track both and had a linear resource vs. performance function. But it does suggest that the true resource vs. performance function is linear or lies below the linear function. The reason is that presumably participants would not act against the instructions and ignore the additional target unless tracking was so resource-intensive that devoting only 50% of the resource to a target yielded worse performance than tracking only one.

The large speed limit cost of the additional target in each hemifield does not particularly support the spatial interference theory describe in the Introduction. Because the objects were always widely spaced, it seems that spatial interference theory would predict only a small effect on speed limit, if any.

As further validation of the hemisphere-specific resource theory and the resemblance of the results to the capacity-one performance benchmark, we document how discrepant the results are from the alternative assumption that observers tracked objects independently in the upper and lower hemifields and within each only one target could be tracked. We call this the UVF/LVF resource capacity-one performance benchmark (blue dashed bar in Figure 4). As described in the Methods, this amounted to calculating a benchmark four-targets speed limit using the performance in the two-targets unilateral arrangement and combining it with guessing in half of trials. As shown in Figure 4, this UVF/LVF resource benchmark's speed limit (1.14 rps, 17.91 deg/s) was significantly lower than the measured four-target speed limit

(1.50 rps, 23.57 deg/s), paired $t(4)=8.506$, $p=0.001$, Cohen's $d=3.804$, suggesting that the tracking resources are not specific to the upper and lower visual hemifields.

According to spatial interference theory (Franconeri et al., 2008; 2010), the detrimental effect of additional targets should be equivalent across speeds if the total distance traveled by the objects is constant. Therefore performance should be worse even at slow speeds in the four-target condition compared to the two-target bilateral condition. This would manifest as an increase in the “lapse rate” parameter in our psychometric function fit. This parameter represents the ceiling performance level. If spatial interference among targets impairs tracking, it should reduce accuracy more for conditions with higher target numbers, thus inflating their lapse rates relative to those conditions with fewer targets.

A repeated-measures ANOVA was conducted with conditions and subjects as the independent variables and lapse rate the dependent variable. There were no significant differences among the three conditions: 2-target bilateral (lapse rate = 0.04), 2-target unilateral (lapse rate = 0.05), and 4-target (lapse rate = 0.03), $F(2,8)=0.288$, $p=0.757$, partial $\eta^2=0.067$. These results argue against significantly greater spatial interference when more targets are tracked.

Hemifield independence was also tested by Experiment 3 of Holcombe and Chen (2012), but the lapse rates were not reported as the paper was a brief report. For completeness, we report them here. Contrary to what would be expected from spatial

interference, the lapse rate was higher in one of the one-target conditions than in either of the two-target conditions. A repeated-measures ANOVA finds a significant effect of target number ($F(1,7)=6.462$, $p=0.039$, partial $\eta^2=0.48$), but no significant effect of hemifield arrangement ($F(1,7)=4.036$, $p=0.084$, partial $\eta^2=0.366$) and no interaction between target number and hemifield arrangement ($F(1,7)=1.194$, $p=0.311$, partial $\eta^2=0.146$). The significant effect of target number arose from slightly larger lapse rates in tracking one target than two targets: one-target bilateral (lapse rate = 0.06), one-target unilateral (lapse rate = 0.03), two-target bilateral (lapse rate = 0.03), and two-target unilateral (lapse rate = 0.02).

Experiment 2: Eye tracking and constant number of reversals

This experiment was motivated primarily by a concern with Experiment 1.

Participants may not have maintained accurate fixation on the fixation point. To address this, in the present experiment we recorded eye movements with an eye tracker.

A second point is that in Experiment 1, for trials with lower speeds the number of reversals was greater. Therefore it is uncertain to what extent the detrimental effect of increased speed was due to speed *per se* or to fewer reversals (if reversals might somehow have benefited performance). To resolve this issue, in Experiment 2 we equated the number of reversals across speeds.

Methods

Six participants (four male, aged 22-37) who reported normal or corrected-to-normal

vision agreed to participate, following approval of the protocol by the University of Sydney's ethics committee. Two were authors and three also participated in Experiment 1.

The apparatus, stimulus and procedure used were identical to those of Experiment 1 except for the addition of the eye tracker and the changes in the reversal times. During the 6.6 revolutions of cumulative distance travelled by the blobs after the target-cuing interval, the blobs changed direction at random successive points of between 2.2 and 3 revolutions, resulting in 2 to 3 reversals per trial. The direction changes for each ring were determined randomly and independently of those for other rings. Each observer participated in 48 trials for each of five speeds. This was fewer than Experiment 1, to accommodate the eye tracker calibration and recalibration time. The speeds for individual observers were chosen based on piloting. The tracking durations were set to achieve a constant distance travelled of 6.6 revolutions. Observers were presented with 240 experimental trials in total, divided into two sessions in two separate days.

Eye movements were monitored using an SR Research EyeLink 1000 eye tracker and analyzed with the Eyelink 1000 software, version 1.5.2. At the beginning of each session, the eye-tracking system was calibrated and validated using the standard 5-point calibration. The experimenter monitored the video image of the participant's eye at the beginning of each trial to ensure the participant fixated and the eye tracker continued to report this correctly. The eye tracker was recalibrated if, during the interval before the trial, it registered the participant's

eye location as away from fixation even though the participant reported fixating. If the eye tracker indicated that the participant moved their eye by more than 2 deg of visual angle from the fixation point, the trial was discarded.

Results and discussion

The criterion of eye movement greater than 2 deg from fixation led to exclusion of 8.3% of trials (SD = 3.3% across participants). A repeated-measures ANOVA found no significant difference in number of these eye movements across the 5 speeds, $F(4, 20) = 2.146$, $p = 0.113$, partial $\eta^2 = 0.3$, or 3 conditions (two-targets bilateral, two-targets unilateral, and four targets), $F(2, 10) = 1.677$, $p = 0.235$, partial $\eta^2 = 0.25$. The ANOVA also showed no significant interaction between speed and condition, $F(8, 40) = 0.539$, $p = 0.82$, partial $\eta^2 = 0.097$.

The data and fitted curves are shown for each participant in the top panel of Figure 5, with associated speed limits (68% thresholds) shown in the bottom panel. As for Experiment 1, the speed limit was considerably higher in the bilateral arrangement (1.89 rps, 29.69 deg/s) than in the unilateral arrangement (1.64 rps, 25.76 deg/s) for tracking two targets. This difference was statistically significant according to a paired t-test, $t(5) = 4.233$, $p = 0.008$, Cohen's $d = 2.45$.

Consistent with the hemisphere-specific resource theory, compared to the speed limit for tracking two targets in a single hemifield (two targets unilateral condition), adding two

more targets in the opposite hemifield (four targets condition) had little to no effect on the speed limit (1.64 rps versus 1.72 rps, paired t-test $t(5) = -1.375$, $p = 0.228$, Cohen's $d = -1.17$). But compared to the speed limit in the two-target bilateral condition (1.89 rps, 29.69 deg/s), the speed limit for the four-target condition was significantly poorer (1.72 rps, 27.02 deg/s) paired t-test $t(5) = 8.307$, $p < 0.001$, Cohen's $d = 3.784$.

The capacity-one performance benchmark puts these speed limit differences in perspective by calculating the result that would occur if participants only tracked one target in each hemifield and ignored the other. The ensuing capacity-one speed limit is shown by the red dashed bar in the bottom of Figure 5. As occurred for Experiment 1, the measured speed limit for tracking four targets (1.72 rps, 27.02 deg/s) was not significantly different from that of the capacity-one performance benchmark (1.74 rps, 27.34 deg/s), paired t-test $t(5) = -0.603$, $p = 0.573$, Cohen's $d = -0.266$. This is consistent with the possibility that the tracking resource in each hemisphere was only sufficient to track one fast target in each hemifield.

As we did for Experiment 1, to further validate the hemisphere-specific resource theory, we document how discrepant the results are from the alternative assumption that observers tracked objects independently in the upper and lower hemifields and within each only one target could be tracked (UVF/LVF resource benchmark). Here we see the lone statistical difference from Experiment 1 - the discrepancy of the observed speed limit (1.72 rps, 27.02 deg/s) from the prediction (1.52 rps, 23.88 deg/s) did not reach significance

although the effect was in the expected direction, paired $t(5)=2.119$, $p=0.088$, Cohen's $d=1.414$ (Figure 5). This may reflect the reduced power of this experiment— mainly because of the additional time demands of eye tracking, there were only 30% as many trials per participant as in Experiment 1.

Experiment 3: Resource allocation between two targets of different speeds

Under resource theory, different amounts of resource might be allocated to different targets in a demand-based manner. This possibility is tested in Experiment 3 by comparing the speed limit at which observers could track a particular disc (the “critical target”) under two conditions, in both of which the observer is required to track a total of two targets. In the “other-slow” condition the second target moves at a slow speed of 0.5 rps whereas in the “same-speed” condition the second target moves at the same speed as the critical target.

We reasoned that because in the other-slow condition the second target was slow, less resource would be needed to track it. This should leave more resource for the critical first target, allowing it to be tracked at a faster speed, provided that resource can be allocated unequally. Because the resource pools operate independently for the left and right visual hemifields, we predicted that an improved speed limit for the other-slow condition would only occur for the unilateral arrangement, not the bilateral arrangement.

Methods

Seven participants (aged 27-32, six males) who reported normal or corrected-to-

normal vision agreed to participate in the protocol, which was approved by the University of Sydney's ethics committee. One participant was the first author.

The apparatus and stimulus used were identical to those of Experiment 1 except for a few changes described here. Whereas in Experiment 1 either two or four discs were cued as targets, here in Experiment 3 two discs were always cued, one in each of two quadrants (Figure 6). In the same-speed condition, the two targets moved with equal speeds, which varied across trials from 0.9 to 2.1 rps. In the other-slow condition, the critical target moved at a speed between 0.9 and 2.1 rps and the slow target always moved at 0.5 rps. As in Experiment 1, at the end of the trial, one pair of discs was indicated and the participants used the mouse to indicate which of the two discs was the target. Both targets were equally likely to be queried.

As in Experiment 1, both bilateral and unilateral arrangements were used (Figure 6). In the bilateral arrangement, in half the trials the targets were both above the fixation point and in the other half both targets were below the fixation point. In the unilateral arrangement, in half the trials the targets were both to the left of the fixation point, and in the other half both were to the right of the fixation point.

The sequence of events was identical to that of Experiment 1 but there was a difference in how the duration of the trials was set. The total tracking interval varied randomly between 3.0 and 3.8 s. Because targets during the same trial had different speeds,

their travel distances were necessarily different. At the end of a trial, one pair of discs was indicated with a central arrow and participants were prompted to use the mouse to indicate which was the target. Each observer participated in 128 trials for each of five speeds. Observers were presented with 640 experimental trials in total, divided into four sessions. Each participant did no more than two sessions a day and had a minimum break between sessions of 5 minutes. The data were analyzed as in Experiment 1, with speed limits (68% thresholds) extracted from the psychometric curve fit.

Results and discussion

For each participant, the data and associated psychometric curves are shown in Figure 7 for the same-speed vs. other-slow conditions and the unilateral vs. bilateral arrangements. When the participants tracked two moving targets with equal speeds (same-speed condition), performance for the bilateral arrangement (red curve) was better than for the unilateral arrangement (blue curve). A paired t-test for the same-speed condition found that across the seven subjects, the speed limit of the bilateral arrangement (1.98 rps, 31.11 deg/s) was significantly higher than that of the unilateral arrangement (1.62 rps, 25.45 deg/s), $t(6) = 4.549$, $p = 0.004$, Cohen's $d = 1.897$, bottom right panel of Figure 7, further supporting the hemisphere-specific resource theory.

In the unilateral arrangement, most participants had higher performance for the fast target when the second target moved slowly (other-slow condition) compared to when it

moved at the same speed, paired t-test $t(6)=2.68$, $p=0.037$, Cohen's $d=1.014$. This is the critical finding, supporting the theory that in the unilateral arrangement each target was allocated different portions of the hemifield-specific resource, depending on its speed.

The results of the bilateral condition suggest that tracking resources cannot be shared across the vertical hemifield boundary to the extent that they can within a hemifield—in the bilateral arrangement, there is no significant difference between the speed limit of the faster target for the same-speed condition and the other-slow condition, $t(6)=0.489$, $p=0.642$, Cohen's $d=0.194$. This difference between the unilateral and bilateral arrangements was confirmed by the significant interaction between hemifield arrangement and speed found in a repeated-measures ANOVA, $F(1,6) = 7.034$, $p=0.038$, partial $\eta^2=0.54$.

In order to understand in more detail how tracking performance was affected by the speed difference of the targets, the performance for tracking the slow target (in the other-slow condition) is plotted in the bottom left panel of Figure 7. It shows percent correct for the slow target as a function of speed of the faster target.

For the slow target, a downward trend in tracking accuracy was observed as the speed of the fast target was increased. This was analyzed by linear regression in the unilateral arrangement ($b=-0.106$, $r^2=0.161$, $p=0.017$) and in the bilateral arrangement ($b=-0.058$, $r^2=0.132$, $p=0.032$). The drop was significantly larger for the unilateral arrangement than the bilateral arrangement according to a paired t-test comparing the slopes of the two conditions,

$t(6)=2.512, p=0.046$, Cohen's $d=1.105$. This was not significant however in the alternative analysis of a repeated-measures ANOVA, where it would have manifest as an interaction between speed and hemifield arrangement, $F(4,24)=0.898, p=0.481$, partial $\eta^2=0.13$. The ANOVA did show a significant speed effect, $F(4,24)=3.361, p=0.026$, partial $\eta^2=0.359$, and a marginally significant hemifield arrangement effect, $F(1,6)=5.161, p=0.064$, partial $\eta^2=0.462$. A significant decrease in performance regardless of the hemifield of the speedier target would suggest that the resource is not 100% hemisphere-specific. More data would be needed to be confident of this.

Returning to the main results of speed limits, we have suggested that the lower speed limits in the unilateral arrangement were due to 50% resource per target being insufficient to track accurately at high speeds. The lower speed limit might alternatively reflect greater difficulty regardless of relative disc speed, for example due to greater spatial interference. Such a general difficulty factor would cause the psychometric function to saturate at a lower ceiling in the unilateral condition. This is not apparent in the plots, and we confirmed the lack of a significant effect by examining the lapse rates of the fits. The lapse rate sets the ceiling on performance. A repeated-measures ANOVA showed no significant differences in lapse rates between the unilateral (0.03) and bilateral (0.02) conditions ($F(1,6)=0.948, p=0.368$, partial $\eta^2=0.136$), and between the same-speed (0.03) and other-slow (0.02) conditions ($F(1,6)=0.203, p=0.668$, partial $\eta^2=0.033$), and no significant interaction between the two

hemifield arrangements and speed conditions ($F(1,6)=0.005$, $p=0.944$, partial $\eta^2=0.001$).

Any speed-invariant difficulty difference was non-existent or too small to be detected.

General discussion

These results provide new support for the theory that a largely hemisphere-specific resource mediates tracking, and that this resource can be differentially allocated to targets with different speeds. In the unilateral (same hemifield) arrangement, pairing a target with a slower target rather than one of the same speed yielded better performance for the first target, presumably because slower targets are allocated less of the resource than faster targets.

Hemisphere-specific resource theory

Experiments 1 and 2 showed that the speed limit for tracking four targets was substantially lower than for tracking two targets bilaterally but was similar to that for tracking two targets unilaterally. Previous studies also found that tracking four targets spread across both visual hemifields had little or no additional cost over tracking two targets within a single hemifield (Alvarez & Cavanagh, 2005; Battelli et al, 2009). Our experiment shows that this result holds even when the objects are widely separated to avoid spatial interference and the total distance traveled by the discs is held constant across trials. Holding travel distance constant was done to avoid a possible increase in spatial interactions with speed (Franconeri et al., 2010).

The similarity of the observed speed limits to the capacity-one performance

benchmark is consistent with a linear performance vs. resource function, steeper than the independent samples (square root) function (Horowitz & Cohen, 2010). It is also equivalent to what would occur if participants, realising that 50% of the resource was not enough to track each fast-moving target, gave up on one target in each hemifield and focused all resource on the other target. The consistent although statistically insignificant finding that performance actually falls below the benchmark (here and Holcombe & Chen, 2012; Holcombe & Chen, in press) suggests that the true resource vs. performance function falls below the linear function (e.g., the coarse line in Figure 2). In other words, for fast targets splitting the resource in two may yield performance worse than halfway towards chance from the one-target level. For example, more than half the resource may be required to have any tracking success with fast targets and therefore if the participants try to track both in each hemifield, they will fail and have to guess regarding both, yielding performance even worse than the capacity-one performance benchmark.

Experiment 3 provided evidence that the resource can be differentially allocated to targets differing in speed, and that this ability was much bigger within hemifields than across hemifields, which further supports the hemisphere-specific theory. We found that a target could be tracked at a faster speed when the other target within the same hemifield was moving relatively slowly. Across hemifields, that effect was not significant. For performance tracking the slow target however, there was a significant although small effect of the speed of

the target in the opposite hemifield, suggesting that the resource is largely but not entirely hemisphere-specific.

The conclusion from the differential allocation evidence (Experiment 3) that the resource is largely hemisphere-specific is consistent with our primary measure of the speed limit cost for adding targets to the same or the opposite hemifield. Adding targets to the same hemifield yielded a much larger cost than adding targets to the opposite hemifield. In support of the notion that a small amount of resource can be shared across hemifields, a non-significant trend for a cost of opposite-hemifield targets was found both here in Experiment 1, 2 and in Experiment 3 of Holcombe and Chen (2012) as well as in both relevant experiments of Alvarez and Cavanagh (2005). In the case of an experiment conducted by Hudson et al. (2012), the reduction in accuracy associated with adding targets in the opposite hemifield reached statistical significance (their Experiment 4).

Serial models and unequal time-sharing within a hemifield

Unequal allocation of the tracking resource is compatible with both parallel and serial models. On a parallel account, all the targets' positions are updated simultaneously, but devoting more resource to a target results in positions being updated more accurately. On the serial account, target positions are updated one by one. The more targets there are, the less frequently their positions are updated (Howe et al, 2010; Tripathy et al., 2011; Tripathy & Howard, 2012; Pylyshyn & Storm, 1988; Oksama & Hyona, 2008). At higher speeds, the

targets travel farther between position updates, resulting in a larger speed limit cost for larger tracking loads (Holcombe & Chen, 2012; in press). The findings of Experiments 1 and 2 are thus consistent with the theory that observers serially track multiple objects within one single hemifield but track independently (in parallel) in the two hemifields. However, see Howe et al. (2010) for evidence against the possibility of serial switching within a hemifield.

Regarding the allocation issue, serial models have assumed that the positions of all targets are registered equally frequently. However, a serial model could allow for one target to receive a greater share of the tracking focus' time. This would accommodate our evidence for flexibility of the resource allocation.

Parallel models and differential resource allocation

The first theory ever proposed of MOT (Pylyshyn & Storm, 1988) is a “slots” or discrete model. According to it, targets are tracked in parallel, each by its own mental index (FINST) or slot. Because observers are assumed to have four to five FINSTs, the Pylyshyn & Storm (1988) slots theory explicitly predicts that tracking performance will not vary as the number of targets is increased, up until the number of targets exceeds the number of FINSTs the observer has. Experiments 1 and 2 contradicted this prediction with the finding that speed limit for tracking two target bilaterally was substantially higher than tracking four targets. However, our results might be accommodated with a modified slot theory where each target can be allocated more than one slot (Horowitz & Cohen, 2010). Under such a model, resource

could be differentially allocated—faster targets could be allocated more slots.

Alvarez and Franconeri (2007) and Vul et al. (2009) proposed continuous resource theories to explain tracking limits. The original notion of continuous resource theory is that the resource must be divided among targets, yielding less per target when more targets are tracked, resulting in poorer performance. To explain the results of our Experiment 3, these theories could be elaborated to specify that more resource can be allocated to a fast target than to a slow one.

An alternative to continuous resource theory is the oscillatory, multi-layer neural network model of Kazanovich and Borisyuk (2006). Because in that model each layer is responsible for tracking a single target, the amount of resource devoted to a target is not predicted to increase if a target moves faster than the other targets in the display. As such, this model is not expected to explain our finding that more resource is required to reach a particular performance level with fast targets than with slow targets. However, it could be modified to allow a single target to be tracked by more than one layer and thereby accommodate our results.

Attention and the tracking resource

Theories of the tracking resource have been vague in that they have not specified the resource vs. performance function they entail. Horowitz & Cohen (2010) suggested independent sampling, which involves a square root increase in precision with resource.

Unfortunately it is not clear whether percentage change in tracking accuracy would be linearly related to precision or have some other mapping. If linear, then accuracy should increase with the square root of resource (depicted in Figure 2), but this was ruled out by the present data. We instead found that for high-speed targets, the cost of dividing the resource was similar to the lower linear function. For slow targets, there was little cost of dividing the resource. What explains this change of the resource vs. performance function with speed?

Perhaps the effect of increasing resource devoted to a target is to increase tracking precision. Less temporal precision for instance may be more costly at higher speeds, because at high speeds the object will more often be far away from its last-registered location. Such ideas, however, require theoretical development beyond the scope of this article.

What is the relationship of the tracking resource to the attentional processes required for other tasks? Possibly the hemifield-specific resource documented here is used solely for visual tracking and selection, and thus cannot be shared with other tasks. Alternatively, the resource that was differentially allocated in the experiments here may be a general (albeit hemisphere-specific) attentional resource required for many other tasks. Simultaneously performing visual search, having a telephone conversation, or discriminating auditory tones can reduce one's tracking ability (Alvarez et al., 2005; Kunar et al., 2008; Tombu & Seiffert, 2008). However, whether or not the resource shared with other tasks is hemifield-specific does not appear to have yet been tested. Here we found that the resources that could be

differentially allocated among targets were largely hemifield-specific. More work must be done, especially testing of hemifield specificity of non-tracking tasks, to connect the present findings to other tasks.

This study provides the first evidence for differential allocation of the hemisphere-specific tracking resource between targets. Further work is needed to determine whether the differential allocation is under strategic control, and whether other tasks share this hemisphere-specific resource.

References

- Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological Science*, 16(8), 637-643.
- Alvarez, G. A., & Franconeri, S. L. (2007). How many objects can you track? Evidence for a resource-limited attentive tracking mechanism. *Journal of Vision*, 7(13), 1-10.
- 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.
- Battelli, L., Alvarez, G. A., Carlson, T., & Pascual-Leone, A. (2009). The role of the parietal lobe in visual extinction studied with transcranial magnetic stimulation. *Journal of Cognitive Neuroscience*, 21(10), 1946-1955.

- Bettencourt, K. C., & Somers, D. C. (2009). Effects of target enhancement and distractor suppression on multiple object tracking capacity. *Journal of Vision*, 9(7), 1-11.
- Bonnel, A. M., & Miller, J. (1994). Attentional effects on concurrent psychophysical discriminations: Investigations of a sample-size model. *Perception & Psychophysics*, 55(2), 162–179.
- Cavanagh, P., & Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends in Cognitive Science*, 9, 349-354.
- Franconeri, S. L., Jonathan, S. V., & Scimeca, J. M. (2010). Tracking multiple objects is limited only by object spacing, not by speed, time, or capacity. *Psychological Science*, 21(7), 920-925.
- Franconeri, S. L., Lin, J. Y., Pylyshyn, Z. W., Fisher, B., & Enns, J. T. (2008). Evidence against a speed limit in multiple-object tracking. *Psychonomic Bulletin & Review*, 15(4), 802-808.
- Franconeri, S. L. (in press). The nature and status of visual resources. In D. Reisberg (Ed.), *Oxford Handbook of Cognitive Psychology* (Vol. 8481): Oxford University Press.
- Holcombe, A. O., & Chen, W. Y. (2012). Exhausting attentional tracking resources with a single fast-moving object. *Cognition*, 123(2), 218-228.
- Holcombe, A.O. & Chen, WY. (in press). Splitting attention reduces temporal resolution from 7 Hz for tracking one object to <3 Hz when tracking three. *Journal of Vision*.

- Horowitz, T. S., & Cohen, M. A. (2010). Direction information in multiple object tracking is limited by a graded resource. *Attention, Perception & Psychophysics*, 72(7), 1765–1775.
- Howard, C. J., & Holcombe, A. O. (2008). Tracking the changing features of multiple objects: progressively poorer perceptual precision and progressively greater perceptual lag. *Vision Research*, 48(9), 1164-1180.
- Howard, C. J., Masom, D., & Holcombe, A. O. (2011). Position representations lag behind targets in multiple object tracking. *Vision Research*, 51(17), 1907-1919.
- Howe, P. D., Cohen, M. A., Pinto, Y., & Horowitz, T. S. (2010). Distinguishing between parallel and serial accounts of multiple object tracking. *Journal of Vision*, 10(8), 1-13.
- Hudson, C., Howe, P. D., & Little, D. R. (2012). Hemifield effects in multiple identity tracking. *PLoS One*, 7(8), e43796.
- Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology*, 43(3), 171-216.
- Iordanescu, L., Grabowecky, M., & Suzuki, S. (2009). Demand-based dynamic distribution of attention and monitoring of velocities during multiple-object tracking. *Journal of Vision*, 9(4:1), 1-12.
- Kazanovich, Y., & Borisyuk, R. (2006). An oscillatory neural model of multiple object tracking. *Neural Computation*, 18(6), 1413-1440.

- Kunar, M. A., Carter, R., Cohen, M., & Horowitz, T. S. (2008). Telephone conversation impairs sustained visual attention via a central bottleneck. *Psychonomic Bulletin & Review*, 15(6), 1135-1140.
- Lee, D. K., Koch, C., & Braun, J. (1999). Attentional capacity is undifferentiated: concurrent discrimination of form, color, and motion. *Perception & Psychophysics*, 61(7), 1241–1255.
- Liu, G., Austen, E. L., Booth, K. S., Fisher, B. D., Argue, R., Rempel, M. I., et al. (2005). Multiple-object tracking is based on scene, not retinal, coordinates. *Journal of Experimental Psychology: Human Perception and Performance*, 31(2), 235-247.
- Liu, T., Jiang, Y., Sun, X., & He, S. (2009). Reduction of the crowding effect in spatially adjacent but cortically remote visual stimuli. *Current Biology*, 19(2), 127-132.
- Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive psychology*, 7, 44-64.
- Oksama, L., & Hyona, J. (2008). Dynamic binding of identity and location information: a serial model of multiple identity tracking. *Cognitive Psychology*, 56(4), 237-283.
- Pastukhov, A., Fischer, L., & Braun, J. (2008). Visual attention is a single, integrated resource. *Vision Research*, 49(10), 1166–73.
- Pelli, D. G., & Tillman, K. A. (2008). The uncrowded window of object recognition. *Nature Neuroscience*, 11(10), 1129-1135.

- 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. (2009). What have we learned about attention from multiple object tracking (and vice versa)? In D. Dedrick & L. Trick (Eds.), *Computation, cognition, and Pylyshyn* (pp. 49-78). Cambridge, MA: MIT Press.
- Tombu, M., & Seiffert, A. E. (2008). Attentional costs in multiple-object tracking. *Cognition*, 108(1), 1-25.
- Tripathy, S. P., & Howard, C. J. (2012). Multiple Trajectory Tracking. *Scholarpedia*, 7(4), 11287.
- Tripathy, S. P., Ögmen, H., & Narasimhan, S. (2011). Multiple-object tracking: A serial attentional process? In C. Mole, D. Smithies & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 117-144). Oxford: Oxford University Press.
- Vul, E., Frank, M. C., Tenenbaum, J. B., & Alvarez, G. (2009). Explaining human multiple object tracking as resource-constrained approximate inference in a dynamic probabilistic model. *Advances in Neural Information Processing Systems*, 1-9.
- Yantis, S. (1992). Multielement visual tracking: attention and perceptual organization. *Cognitive Psychology*, 24(3), 295-340.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working

memory. *Nature*, 453(7192), 233-235.

Figure 1. Hypothetical performance vs. tracking resource functions. The dashed line shows a linear function relating proportion of the tracking resource (and corresponding number of targets, top axis) to performance. This is a candidate function based on our results for fast speeds but for slow speeds the function is believed to be more like the solid curve. The relatively flat right portion of the solid curve indicates that when attention is divided among a few targets, there is little cost for performance, as documented by ourselves and others. Only when there are many targets (leftmost, steep part of the curve) do additional targets impose a

significant cost for performance.

Figure 2. Hypothetical performance vs. tracking resource functions for fast object speeds.

Experiments suggest that for fast object speeds there is always a cost of adding additional targets, and the curves shown here are all consistent with that. The dashed line shows a prediction based on an independent samples idea, where performance increases with the square root of the resource applied. Because at fast speeds we find that performance falls at or below the capacity-one benchmark (results here and of Holcombe & Chen, 2012; Holcombe & Chen, in press), the resource required for successful tracking is high- higher than even the linear resource function. The shape of the function is unknown, but one

possibility is shown by the coarse fuzzy line.

Figure 3. Schematic of the displays for Experiments 1 and 2. Four pairs of red discs were presented in each condition. Each pair moved along a circular trajectory (dotted lines) centered in one quadrant and potentially included a target. The targets were presented in one of three conditions: two targets to the left or right side of the vertical meridian (*two-targets unilateral*), above or below the vertical meridian (*two-targets bilateral*), or four targets with one in each quadrant (*four targets*). The targets were initially white before becoming red like the distractors.

Figure 4. Results of Experiment 1. **Top panel.** For each participant, proportion correct is shown for each speed in the two-target bilateral (red), two-target unilateral (blue) and four-targets (green) conditions. Dotted lines show the 68% thresholds (speed limit). **Bottom panel.** Empirical speed limits for tracking two and four objects are shown as solid bars, and capacity-one model predictions for four targets are shown as dashed bars. At slow speeds performance exceeds the 75% correct performance of the hemisphere-specific capacity-one benchmark. At high speeds (the speed limit), the data are consistent with participants being able to track only one target in each hemifield, and having to guess on the other (the

hemisphere-specific capacity-one benchmark, red dashed bars). In contrast, the data are not consistent with the prediction of the UVF/LVF-specific (blue dashed bars) resource capacity-one performance benchmark. Error bars show one standard error across participants.

Figure 5. Results of Experiment 2. **Top panel.** For each participant, proportion correct is shown for the two-target bilateral (red), two-target unilateral (blue) and four-target (green) conditions. Dotted lines show the 68% thresholds (speed limit). **Bottom panel.** Empirical speed limits for tracking two and four objects are shown as solid bars, and capacity-one benchmark limits for four targets are shown as dashed bars. At slow speeds performance exceeds the 75% correct performance of the hemisphere-specific capacity-one benchmark. At high speeds (the speed limit), the data are consistent with participants being able to track only one target in each hemifield, and having to guess on the other (the hemisphere-specific

capacity-one benchmark, red dashed bars). Error bars show one standard error across participants.

Figure 6. In Experiment 3, four pairs of red discs (Gaussian blobs) were presented in each condition. Each pair moved along a circular trajectory (dotted lines) centered in one quadrant. A disc from each of two pairs was designated as targets. For the bilateral arrangement, the targets (cued in white) were both above or both below fixation whereas for the unilateral arrangement, both targets were to the left or to the right. In the other-slow condition, one target always moved slowly (0.5 rps). In the same-speed condition, the two targets moved at the same speed (which varied across trials).

Figure 7. Top. For each participant, proportion correct is shown for each speed in the two-target bilateral (red) and two-target unilateral (blue) arrangements. Top row shows performance in the same-speed condition, averaging across both targets. Bottom row shows the other-slow condition, for the critical faster target. Dotted lines show the speed limits (68% thresholds). **Bottom right.** Red bars show the speed limits for the bilateral arrangements in the same-speed and other-slow (faster target only) conditions. The blue bars show the speed limit for the unilateral arrangements in the same-speed and other-slow (faster target only) conditions. Error bars show one standard error across participants. Red stars indicate statistically significant differences, $p < 0.05$. **Bottom left panel.** Tracking accuracy for the slow target (mean across subjects) is shown for the other-slow condition, as a function of the

speed of the fast target.

Supplementary Files.

Movie S1. Demonstration of a two-targets bilateral trial of Experiment 1. The targets to track are initially white. Many will notice that tracking the targets here is easier than in Movie S2 and S3. Speed is nominally 1 rps, but will be slower on many computers. The animation is only 30 fps rather than the 120 Hz of the experiment display, so the motion may appear jumpy rather than smooth.

Movie S2. Demonstration of a four-targets trial of Experiment 1. The targets to track are initially white. Speed is nominally 1 rps, but will be slower on many computers. The animation is only 30 fps rather than the 120 Hz of the experiment display, so the motion may appear jumpy rather than smooth.

Movie S3. Demonstration of a two-target unilateral trial of Experiment 1 (also equivalent to a same-speed trial of Experiment 2). The targets to track are initially white. Speed is nominally 1 rps, but will be slower on many computers. The animation is only 30 fps rather than the 120 Hz of the experiment display, so the motion may appear jumpy rather than smooth.

Movie S4. Demonstration of an other-slow trial of Experiment 2. The targets to track are initially white. Speed is nominally 1 rps, but will be slower on many computers. The animation is only 30 fps rather than the 120 Hz of the experiment display, so the motion may appear jumpy rather than smooth.