

The Spatial Resolution of Visual Attention

James Intriligator and Patrick Cavanagh

Harvard University

Two tasks were used to evaluate the grain of visual attention, the minimum spacing at which attention can select individual items. First, observers performed a tracking task at many viewing distances. When the display subtended less than 1° in size, tracking was no longer possible even though observers could resolve the items and their motions: The items were visible but could not be individuated one from the other. The limiting size for selection was roughly the same whether tracking one or three targets, suggesting that the resolution limit acts independently of the capacity limit of attention. Second, the closest spacing that still allowed individuation of single items in dense, static displays was examined. This critical spacing was about 50% coarser in the radial direction compared to the tangential direction and was coarser in the upper as opposed to the lower visual field. The results suggest that no more than about 60 items can be arrayed in the central 30° of the visual field while still allowing attentional access to each individually. Our data show that selection has a coarse grain, much coarser than visual resolution. These measures of the resolution of attention are based solely on the selection of location and are not confounded with preattentive feature interactions that may contribute to measures from flanker and crowding tasks. The results suggest that the parietal area is the most likely locus of this selection mechanism and that it acts by pointing to the spatial coordinates (or cortical coordinates) of items of interest rather than by holding a representation of the items themselves. © 2001 Academic Press

INTRODUCTION

For many people, having their vision tested simply involves reading letters on an eye chart. However, resolving items is not all that there is to seeing them. In order to scrutinize a particular item, for example, to report its identity and pick out its features and actions, it is first necessary to individuate it. This “individuation” is one of the core operations of attention, and it plays many roles in perception. For example, individuation of an item is necessary in order to encode its location and track its position over time.

Support came from Grant EY09258 to PC and from a National Science Foundation fellowship for JI. The authors thank the reviewers for their invaluable comments.

Address correspondence and reprint requests to Patrick Cavanagh, Department of Psychology, Harvard University, 33 Kirkland Street, Cambridge, MA 02138. E-mail: patrick@wjh.harvard.edu.





FIG. 1. A simple demonstration which shows the difference between visual acuity and attention. While fixating the cross, the lines to the right are easily seen—they are thin, vertical, parallel, evenly spaced, black, and all about the same height. However, while still fixating the cross, it is difficult or impossible to attend to an individual line in the middle of the group, say, the fourth line from fixation.

Individuation is also necessary when counting sets of more than four items (Warren, 1897).

The image in Fig. 1 illustrates the distinction between resolving and individuating items. Notice that it is quite easy to resolve all the items in this display. While fixating the cross, the vertical lines in the patch on the right are clearly seen—they are all vertical, parallel, thin, and black. However, while fixating the cross, attempt to attend to the fourth line from fixation. You will probably find it difficult to pick out this item with attention alone. This simple demonstration is related to the “crowding” effect (Bouma, 1970, 1973; Toet & Levi, 1992; Townsend, Taylor, & Brown, 1971) and it shows that even when items are easy to resolve *visually*, there are additional spatial constraints that may limit our ability to *select and scrutinize* individual items. The goal of this article is to measure the spatial limit, or resolution, of attention. We also demonstrate that this spatial limit is separate from the capacity limit of attention.

The spatial limit on selection is an empirical property that can be measured independently of the nature of the mechanism (or mechanisms) of selection. For example, whether attention is conceived of as space-based or object-based, as a window of selection, a resource, or a filter, it will still have limits on the finest spacing of either the locations or objects on which it can operate. Many findings suggest that attention is location-specific (Broadbent, 1982; Downing & Pinker, 1985; Eriksen & St. James, 1986; Posner, 1980; Tsal, 1983; Tsal & Lavie, 1988). William James suggested in 1890 that attention operates like a “spotlight” that can be shifted around the visual field, offering improved processing of items falling within its “beam.” Posner, Eriksen, and others (e.g., Eriksen & Hoffman, 1973; Posner, 1980; Posner & Petersen, 1990; Posner, Snyder, & Davidson, 1980) have explored many aspects of attentional processes within this metaphor.

However, several studies suggest that it may be possible to attend to a group of similar objects based on attributes other than location, such as color, motion, shape, and surface (e.g., Driver & Baylis, 1989). These studies suggest that visual attention may not simply spread across visual space, but that it may also spread across feature spaces. Other studies indicate that attention

is bound to objects or to surfaces rather than locations (Duncan, 1984). In particular, attention appears to speed processing of objects that are part of the same surface regardless of their absolute spatial location (He & Nakayama, 1994; Nakayama, He, & Shimojo, 1995). Many psychophysical studies (Baylis & Driver, 1992; Duncan, 1984, 1993; Lavie & Driver, 1996; Kahneman, Treisman, & Gibbs, 1992; Kramer & Jacobson, 1991) and neuropsychological reports (Behrmann & Moscovitch, 1994; Driver & Halligan, 1991; Humphreys & Riddoch, 1993) support an object-based view.

Given the range of views concerning the mechanisms of attention, it is hard to describe the region over which attention operates in terms compatible with all the alternatives—space- or object-based, selection, resource allocation, or filters. For simplicity, we use the term “selection” to describe the operation of attention and “region of selection” to describe the area over which it operates but we do not favor one model over another. Our measures of attentional resolution, even though they are presented in terms of selection, are not dependent on the assumptions of any particular mechanism.

Spatial Distribution of Attention

Four types of studies have addressed the size of the region selected by attention: cueing, flanker interference, crowding, and counting tasks. The results in these four paradigms vary widely so we review them briefly and outline the likely reasons for the inconsistent outcomes. These points lead us to a specific definition for the spatial resolution of attention that avoids these problems and to a set of appropriate tasks to measure it.

In cueing tasks, a briefly presented cue is assumed to draw attention to its location. Following some delay, a probe is presented at various spatial offsets from the cue. The profile of performance as a function of distance from the cue gives an idea of the spatial distribution of attention around the cue. Some cueing studies have examined how the size of the focus of attention changes with retinal position. Posner (1978) claimed that there was no variation with eccentricity but others report a tighter focus at the fovea than in the periphery. For example, Sagi and Julesz (1986) report that the focus doubles in size from 2° to 4° eccentricity. Even for cues within a few degrees of the fovea, however, estimates of the size of the attentional focus range widely, from 5 min of arc or less (Nakayama & Mackeben, 1989, Experiment 5, 8 arcmin display, although this study used flankers as well as cueing), to 1° (Posner, Snyder, & Davidson, 1980), to about 1.5° (at 2° eccentricity, Sagi & Julesz, 1986) to 10° to 12° (Steinman, Steinman, & Lehmkuhle, 1995, using a motion test) to entire hemifields (Hughes & Zimba, 1987).

How can these estimates vary over almost 3 orders of magnitude? As many point out (e.g., Eriksen & St. James, 1986), the region of selection can undoubtedly scale up and down to suit the demands of the task and there is little if anything in many of the cueing tasks to constrain the size of the attentional focus. Following the offset of the cue, the subject is typically

left attending to blank space until the probe appears, a poor framework for anchoring a selection mechanism, especially one that may deal with objects, not space. Despite these limitations, some authors claim that the cueing paradigm measures the "resolution of spatial attention" (Shulman, Wilson, & Sheehy, 1985) or the shape of the "searchlight of attention" (Sagi & Julesz, 1986). In reality, these studies address the spread of attention under conditions that do little to constrain or define how it should be spreading.

In the flanker paradigm (see Eriksen, 1995), observers respond to a test item that is surrounded by distractors. A central feature of this paradigm is the compatibility of the distractors and the test. A typical experiment might have one response for an A or a U as the test and a different response for, say, an H or an M. If the test item is an A, then a compatible distractor would be a U and might speed the response, whereas an incompatible distractor would be either an H or an M and might slow the response. The measure of interest in the task is the spacing (from test to distractor) over which compatibility effects can be registered. This spacing is assumed to represent the area of selection, but it characterizes primarily its outermost reach. Estimates of the area of selection also vary with the flanker task but not over as large a range as with the cueing paradigm. Using the flanker task, Eriksen and coworkers (Eriksen & Hoffman, 1973; Eriksen & Eriksen, 1974) first suggested a fixed size of approximately 1° of visual angle. However, when their model was updated to the "zoom lens" version (Eriksen & St. James, 1986), the attentional focus was allowed to change its size, and no specific test of the smallest available size was made. Goolkasian (1999) estimated a selection area of 10° or so around a central test but about 1° for a peripheral test with a foveal distractor. Stroop versions of this flanker paradigm have shown compatibility effects reaching as far as 5° from a central target (Gatti & Egeth, 1978). Finally, Lavie (1995) did not establish a specific size of the region of selection but she did demonstrate that it was influenced by the attentional load of the task. Specifically, adjacent distractors interfered with a central task only when it was an easy one, suggesting that resources not tied down by the central task were available to pick up information from adjacent regions of space. In difficult tasks, no surplus resources were available so the region of selection was at its smallest size, effectively excluding the influence of the flankers.

In the crowding paradigm, observers view an array of regularly spaced items. A test item near the center is more difficult to report than a test item at either end of the array (Bouma, 1970, 1973; Andriessen & Bouma, 1976; Westheimer, Shimamura, & McKee, 1976). The measure of interest is the spacing of the flankers at which identification of the test is at threshold. The effect of crowding cannot be overcome by additional viewing time (Townsend et al., 1971), and it increases with the similarity of the target and its distractors (Andriessen & Bouma, 1976; Kooi, Toet, Tripathy, & Levi, 1994; Banks & Prinzmetal, 1976). Performance at different eccentricities and along

different meridians of the visual field has been measured in more detail in this task than in the others (e.g., Toet & Levi, 1992). Outside the fovea, the critical spacing is surprisingly large, typically equal to about $\frac{1}{3}$ the eccentricity, but at the fovea, estimates of critical spacing are quite small, $\frac{1}{10}$ th of a degree or less (Wolford & Chambers, 1984; Toet & Levi, 1992). One reason for difference between estimates from flanker and crowding studies is the different definitions typically used for critical spacing in the two tasks (although these differences are also found between some studies within the same task). In the flanker task, the critical spacing is generally the widest spacing beyond which the distractors have no differential effect on reaction time, whereas in the crowding task, the critical spacing is the closest spacing at which sufficient information remains to identify the test with, say, 75% accuracy.

The flanker and crowding paradigms provide well-defined measures of the spacing at which adjacent items interfere with the test; however, there are a number of factors that may influence this critical spacing. Some researchers have explained the flanker and crowding results by appealing to a process of lateral interaction. In this case, features of adjacent items suffer mutual interference or exchange (Wolford, 1975; Treisman & Schmidt, 1982), degrading the representation of the test item. For example, if features of nearby distractors were mixed somehow with those of the test, the test would become degraded—doubly so if the distractor features, rather than being neutral concerning the task response, favor a different response (e.g., Eriksen & Hoffman, 1973). Others have attributed the results to considerations of spatial uncertainty (Levi, Klein, & Yap, 1987). Some have suggested that attention is fundamentally involved (Walley & Weiden, 1973; Wolford & Chambers, 1983; He, Cavanagh, & Intriligator, 1996) while others claim that attention plays no part in the matter (Banks, Larson, & Prinzmetal, 1979). We are interested in the possibility that the critical spacing may be set by attentional resolution (He et al., 1996). In this case, the observers cannot report the test item's identity because they cannot access the item. However, it is equally possible that the item is selected perfectly well by attention but the item's representation is no longer recognizable. In this case, prior to selection, the representation has been distorted or degraded by interference from neighboring items. Until there is a consensus concerning the reasons for interference caused by nearby distractors, we cannot know whether flanker and crowding studies offer useful measures of attentional resolution.

Finally, a counting task may sidestep these problems and offer a pure measure of attentional resolution. When observers are asked to enumerate more than four finely spaced items, individual items must be selected and, once counted, mentally tagged to avoid recounting them. As with the flanker and crowding paradigms, there is a clear definition of the spatial bounds within which attention must select items. However, all items are now identi-

cal so feature exchange between adjacent items cannot alter the item's identity. More to the point, the observer does not have to identify the items. It is sufficient merely to note their presence, and as long as anything spatially distinct remains visible, counting can proceed. Unfortunately, there have been very few studies of the density limits on counting (Landolt, 1891; Kowler & Steinman, 1977, 1979), none of which investigated the change of performance with eccentricity. Nevertheless, the results have been interesting. In the earliest study, Landolt (1891) had observers count finely spaced dots or stripes directly at the center of gaze. He reported that his observers could not count the stripes or dots if their spacing was less than about 5 arcmin, even though they remained perfectly and distinctly visible: "on arrive à un point où l'on ne peut plus les compter d'aucune façon, alors qu'ils demeurent encore parfaitement et distinctement visibles" (p. 385). Landolt allowed eye movements during counting so his results do not necessarily reflect limits of attention alone. However, Kowler and Steinman (1977, 1979) repeated this counting task with and without eye movements and found that the eye movements led to slightly *worse* performance for regularly spaced arrays like those that Landolt had used. As mentioned above, we assume that counting requires each bar to be individuated in turn. Consequently, Landolt's results suggest that the finest spacing which supports individuation in the fovea is about 5 arcmin when eye movements are allowed. Kowler and Steinman (1977) found only 40 to 50% accurate counting even with spacings of 7 to 14 arcmin in the fovea, and they found that this performance was similar with or without eye movements. These studies imply that the finest spacing at which attention operates may be on the order of 5 to 10 arcmin at the fovea—a value substantially coarser than the finest spacing that can be visually resolved at the center of gaze, about 1 arcmin (Campbell & Robson, 1968).

In this article, we verify and extend these initial estimates using two different methods that have the advantages we noted for the counting task: The spatial bounds on selection are well defined by the spacing to nearby distractors; feature interactions are simplified by making all items identical; and the effects of feature interactions are minimized by requiring only target selection, and not identification. These advantages combine to allow, as much as possible, uncontaminated measures of the spatial limits of attention's access to visual targets. The two tasks are multiple item tracking in dynamic displays and attention stepping through static displays.

Definition of the Resolution of Spatial Attention

Whether attention operates on locations or objects, there will be a finest scale at which it can operate. Objects spaced more finely than this limit are beyond the limit of *attentional* resolution and thus cannot be selected individually for further processing based only on their location (as indicated by a nearby cue or by serial position, say, fourth from the left). The finest scale

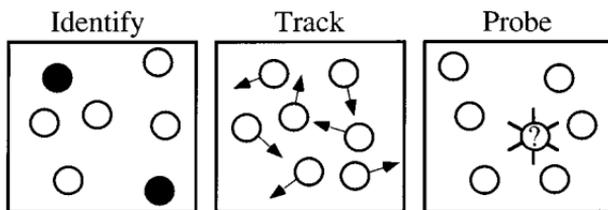


FIG. 2. Multiple-object tracking paradigm. A subset of objects is identified (left). Next the objects become the same color as the others and all objects undergo a period of random motion. Observers attempt to track the previously identified subset (middle). An item is probed and the accuracy of reporting whether it had been tracked is used to assess performance (right).

of *attentional* resolution will necessarily be limited by the scale of available *visual* resolution, but as noted above, it might be substantially coarser than this.

The Spatial Resolution of Attention in a Tracking Task

Our first experiment used dynamic tracking displays in which the targets must be individuated in order to keep track of their constantly changing locations. In this procedure, developed by Pylyshyn and his colleagues (Pylyshyn, 1989; Pylyshyn & Storm, 1988; Sears & Pylyshyn, 2000), observers view a display containing, say, seven identical items (Fig. 2). The items are initially stationary and a subset of them (two in this example) is identified by a color change. The subset then returns to the original color so that the attended objects are visually indistinguishable from the other displayed items. All the items then move for some time. Observers are asked to keep track of the items that were identified at the beginning of the trial. Note that this type of tracking can be done without eye movements and therefore has been called “attentive” or “covert” tracking. A recent study by Scholl and Pylyshyn (1999) demonstrated that the tracking can be maintained even if the targets are briefly occluded—leaving no physical stimulus on the display.

Using this procedure, Pylyshyn and his colleagues demonstrated that most observers are capable of tracking up to four objects at the same time. What interests us is not the surprising level of performance, but the limits on that performance. Specifically, a loss of tracking can occur when a tracked item comes so close to other untracked items that it can no longer be resolved by attention—it is no longer individuated and therefore information regarding its location is lost. After such a close encounter, the observer no longer knows which of the items involved in the encounter was the target. Although there are several other possible sources of error in this task, we manipulate this factor of close encounters (by changing viewing distance) in order to estimate the smallest region that can be attended.

Pylyshyn and his colleagues hypothesized that multiple-object tracking was carried out via a set of FINSTs (an acronym that stands for “fingers of

instantiation'') that index the locations of the tracked objects. However, Yantis (1992) showed that Pylyshyn's model (based on a list of locations) could not explain many aspects of tracking performance. In particular, Yantis showed that anything that facilitates grouping of the targets also improves tracking performance. For example, he showed that if the elements formed a convex polygon, had similar velocities, or were vertices of a rigidly rotating 3D object, performance was enhanced. None of these grouping effects is easily explained by appealing to the location list of Pylyshyn's FINST model. Clearly, a set of targets that is easily grouped is also easily tracked, implying, as Yantis (1992) suggested, that attention is focused not as separate tags on unrelated individuals, but on a spatially organized group whose structure changes over time. Although these grouping factors strongly influence tracking performance, our display manipulation keeps grouping strength constant across displays. We therefore were able to study the specific effects of item density independent of other factors that may influence tracking performance.

In this first experiment, our goal was to measure the spatial resolution of attention during tracking. Observers' overall tracking performance was measured as a function of the visual angle of the entire stimulus display. Specifically, viewing distance was varied over more than two orders of magnitude so that the display size ranged from extremely large (80° of visual angle) to extremely small (0.625° of visual angle). No other factors were varied in the display (of course, velocity and size on the retina varied with distance). Separate control tasks were conducted to examine whether low-level limits imposed by spatial or motion acuity determined the limits of the tracking performance. We found that, well before the items or their motions became hard to resolve, tracking performance declined dramatically. This result indicates that the spatial resolution of attention is considerably coarser than would be predicted based on low-level spatial and motion considerations.

We also examined whether this resolution limit was independent of attentional load. Perhaps, when tracking only a single item, the extra attentional resources would allow a further tightening of the attentional field (i.e., a finer resolution of attention). However, we found that performance with a single target showed little improvement in the finest spacings at which tracking was possible.

Attentional Resolution in Static Displays

The multiple-object tracking task gave us an estimate of resolution, but it did not allow us to evaluate the stimulus configurations that actually caused errors. In particular, because the accuracy of tracking was tested only at the end of a trial, the location and spacing of items at the moment of a tracking failure were not available. In our second experiment, we obtained better control of spacing and eccentricity by using static arrays of uniformly spaced,

identical items. As a test of how well observers could individuate, they were asked to “attentionally step” from item to item following verbal instructions to move, say, “right,” then “left,” then “right,” then “right,” and so on. At the end of several steps, we evaluated whether the observer was attending to the correct item. Since the judgment was made at the end of several steps, we still do not know when or where any errors occurred in the sequence of steps. However, we do know the spacing and eccentricity for each error, as these were held constant in each trial (and varied across trials). This task differed from the first tracking task in two ways: (1) rather than tracking several items, only a single item was individuated at a time; and (2) in contrast to the randomly changing densities and eccentricities of the tracking task, in the stepping task the eccentricity of the target and the spacing of the flankers remained constant within each trial. In situations where the items were too densely packed for the target to be individuated, observers were unable to attentionally select and step through individual items.

Given the nature of our two tasks, we are measuring only the limits on attention’s fundamental access to an item’s location. The results provide data regarding the resolution of attention, how it varies across the visual field, and how it is influenced by stimulus configurations. Finally, based on these data, we generate a display that allows observers to visualize their attentional resolution as a function of location in the visual field.

EXPERIMENT 1: TRACKING AT DIFFERENT SCALES

In this first experiment, attentive tracking was performed over a range of display sizes and performance at each scale was measured. The smallest scale at which attentive tracking was possible was shown to be coarser than the smallest scale at which ordinary visual and motion acuity operate effectively.

Previous research by Pylyshyn and colleagues has shown that, using attention alone, it is possible to track up to four independently moving targets. Subsequent studies (some involving simulations) have suggested that the tracking of the objects takes place in parallel (Pylyshyn & Storm, 1988), can continue through object occlusion (Scholl & Pylyshyn, 1999), and that the underlying processes may be related to memory (Kahneman et al., 1992) or to grouping (Yantis, 1992). In this first experiment, observers viewed a set of objects moving with quasirandom trajectories. Observers were asked to use attention to track four of these objects. To establish the range of scales over which attentive tracking can operate, this task was performed at many viewing distances. Note that changing viewing distance modified both the scale of the display and the retinal velocities of the stimuli: smaller displays had correspondingly slower retinal velocities.

Additionally, in order to evaluate possible interactions between the spatial resolution of attention and capacity limits of attention, the tracking task was

also run with one and three targets while maintaining fixation. A subset of the viewing distances of the main task was used to probe tracking performance at the smaller display sizes. The fixation point was added so that the tracked disk would sample the same range of retinal eccentricities when the observer was tracking one disk as when the observer was tracking three or four disks. Without the fixation point, the observer could use eye movements to track the single target and keep it and its neighboring distractors all within the centermost portion of the fovea. This overt tracking would give an artificial advantage to tracking a single disk compared to tracking three disks (where even if one target is kept at the fovea, the other two must have greater eccentricities). In effect, a fixation point is equivalent to adding an additional target (albeit one which never moves), so that the condition of tracking three targets with a fixation should be about as difficult as tracking four targets with free viewing.

Finally, spatial and motion control tasks were also performed at the closest and farthest viewing distances. These control tasks established the limits of spatial and motion acuity in conditions where these factors were most likely to limit performance. The condition of particular interest is when two disks have a near approach in the most peripheral regions of the display. The spatial task was designed to measure the smallest center-to-center spacing that still allowed observers to distinguish a gap between two disks (i.e., the closest spacing at which items were still resolved as separated). The motion control task was used to measure the smallest separation for which observers could still identify the motion direction (clockwise vs counterclockwise) as the two disks rotated around a common center. A simulation of the displays in the main (tracking) task—via a Monte Carlo procedure—was used to determine two separations: the closest approach that ever occurred and the closest approach that occurred at least once on every trial for every target. These separations were compared to the spatial and motion limits measured in the control tasks.

Methods and Procedures

Observers

Four observers (two males and two females) were tested in the main experiment (tracking four targets) and the two control experiments. Two of the observers were naive. The observers ranged in age from 22 to 35 years. Four different observers participated in the one- and three-target tracking sessions. They ranged in age from 24 to 32 years. All observers were recruited from the Harvard University community and participated in at least one session lasting about 2 h. All observers had normal or corrected-to-normal visual acuity.

Apparatus

The experiments were conducted on an Apple Macintosh computer. All software was written in Think C™ using the VisionShell™ libraries created by Raynald Comtois (1999). A 14" color monitor was used for all experiments. To insure the accuracy and constancy of all color

and luminances used in the tasks, the monitors were calibrated for luminance linearity before the experiments were conducted.

Tracking task, four targets. The display was a dark square 17 cm on a side. Within this region, nine identical green (20 cd/m^2) disks (0.86 cm diameter) moved in a semirandom fashion. The velocity of the disks was 10 cm/s. Although all motions were approximately linear, very small random variations in the path of each disk were introduced every 45 ms.

When the disks' movements brought them into contact with the imagined "walls" of the 17-cm square they would "bounce" off of them. In addition, to ensure that disks did not overlap in the display, each disk was programmed to avoid coming within approximately 1.36 cm of any other disk, center-to-center (0.5 cm between the edges of the disks). This buffer zone was not absolute. Disks occasionally came closer than this, but a Monte Carlo simulation (described later) showed that they would never approach closer than 1.24 cm center-to-center (0.38 cm, edge-to-edge). This "repulsion" was accomplished by simulating a repulsion field between each disk and every other disk. Furthermore, this repulsion acted even when the disks were quite far away so that the disks had a subtle avoidance behavior and never underwent sudden shifts in their motions. At the beginning of each trial the color of the "target" disks changed to red (for identification). In addition, during a later test phase, all disks stopped moving and a small circular response frame (roughly twice the size of the disks) appeared at the center of the screen.

Tracking task, one and three targets. In these conditions, all aspects of the displays were identical to those of the main task, with two exceptions. Only one or three targets were initially highlighted in red and a central fixation point was present throughout.

Spatial and motion control tasks. The display used for the spatial and motion control tasks is illustrated in Fig. 3. This display was similar to the display in the main task—however, instead of having many disks moving in unknown directions, only two disks were used. These two disks, identical to those used in the main task, were displayed approximately 7 cm from a screen-centered fixation point.

The two disks rotated around an imagined center (between the two disks) giving the disks

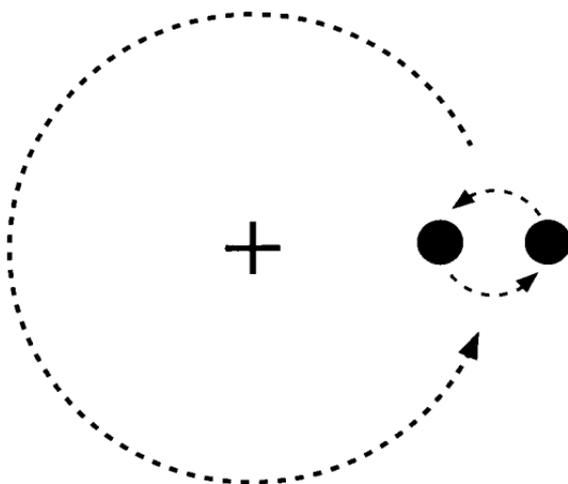


FIG. 3. Stimuli for the spatial and motion control tasks: two disks rotated around a point near the outer edge of the display. This center-of-rotation also rotated around the fixation point. In the spatial task, subjects reported whether they saw a gap between the two disks. In the motion task, they reported whether the disks rotated in the clockwise or counterclockwise direction. The distance between the two disks was varied from trial to trial.

the appearance of two moons orbiting an imaginary planet. In addition, the system of two rotating disks rotated, as a whole, around the screen-centered fixation point. The distance between the two disks was varied—however, the disks always maintained a constant linear velocity which was equal to the velocity of motion used in the main experiment.

Procedures

All experiments were conducted in a dimly lit room. A chin and headrest was used in all conditions.

Tracking task, four targets. Performance was assessed at eight different viewing distances yielding display visual angles of 80°, 40°, 20°, 10°, 5°, 2.5°, 1.25°, and 0.625° (covering a range of viewing distances from 10 cm to 15.58 m). At each viewing distance, observers tracked a subset of four disks from the field of nine identical disks. On each trial, the sequence of events was as follows: First, nine green disks appeared in randomly chosen locations and moved with random directions (all moving at the same speed); second, the to-be-tracked disks were identified: four of the nine disks turned red for the next 5 s and observers were instructed to track these disks with attention alone; third, after this identification phase, the color of the target disks returned to green and thus became visually indistinguishable from the distractors; fourth, during a 5-s tracking interval, observers attempted to keep track of the four target disks; finally, at the end of this tracking interval, all disks stopped moving and a small, circular frame appeared at the center of the screen. Observers used the mouse to move this frame to encircle each disk that they had tracked and clicked the mouse button to record their selections. Disks were selected one at a time. If observers wished to change their decision, items could be unselected. After observers felt that they had correctly identified the four tracked disks, they pressed a key to begin the next trial sequence. Observers performed 12 trials at each viewing distance.

Tracking task, one and three targets. The procedure was identical to that of the main task except that observers were instructed to fixate the central, static fixation point throughout each trial; only one or three targets were highlighted on each trial (the one- and three-target conditions were run in separate sessions) and performance was assessed at only six different viewing distances yielding display visual angles of 20°, 10°, 5°, 2.5°, 1.25°, and 0.625° (corresponding to viewing distances from 48.2 cm to 15.58 m). In each session, observers performed 12 trials at each viewing distance. Observers ran one session in the single target conditions and three sessions in the three-target condition.

Spatial and motion control tasks. The spatial and motion control tasks were only conducted at the nearest and farthest viewing distances (10 cm and 15.58 m, respectively). Each of these tasks began with the observers fixating a screen-centered fixation point. The rotating pair of disks then appeared in a position that varied randomly across trials (but always at a constant distance from the fixation point). The disks then moved around each other and around the fixation point in a circular arc for 1200 ms. Even though the close encounters in the main experiment would probably be much briefer than this, the items themselves are present continuously over a much longer period (5 s). A briefly presented control test would not reflect the advantages of the maintained attention to the target items in the main task. Given that lateral masking does not seem to be much affected by even unlimited viewing time (Townsend et al., 1971) we felt that this test duration would be suitable for estimating the critical spacing for visual resolution. For each trial, the interdisk spacing was chosen from a set of either five or six possible values selected separately for each task and viewing distance. For the motion control task, the six spacings ranged from about 0.2 to 2.2 cm (center-to-center) and thus included some stimulus conditions in which the disks were substantially overlapping (they just touch at a center-to-center spacing of 0.86 cm). For the spatial control task, the five spacings ranged from about 0.6 to 1.4 cm (center to center) and thus included some stimulus conditions in which the items were clearly overlapping and others in which a separation was present. In the motion control task, observers had to identify (with a keypress) which direction

(clockwise/counterclockwise) the pair of disks rotated. In the spatial control task, observers were asked to indicate (by keypress) whether they saw a gap between the two disks. After a response was made, and a brief interstimulus interval was followed by the next trial. Six responses were obtained for each spacing.

Results

Tracking Task, Four Targets

Performance with four targets was quite accurate at the larger viewing angles, better than that reported in Yantis (1992) for comparable display size (10°). We are not sure what factors led to the better performance seen with our observers. Most important, however, is that their performance became dramatically worse at small viewing angles. These results can be seen in Fig. 4, where, for the four observers, percentage correct is graphed as a function of the display's visual angle. To be analogous to a typical contrast sensitivity function, the scale of the horizontal axis runs from large to small viewing angles. The horizontal line at 44% correct indicates chance performance, i.e., if, on each trial, an observer were randomly choosing four disks from the nine possible, then their performance would approach this value. The vertical dashed line at the far right indicates the limit predicted by standard

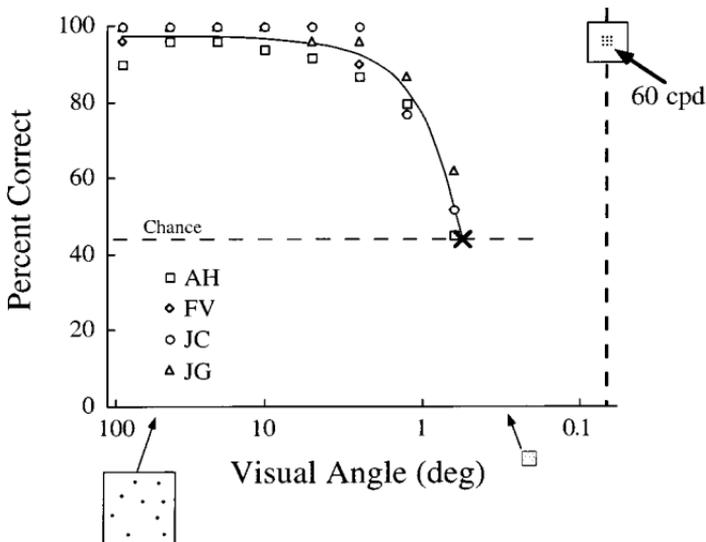


FIG. 4. The performance of four observers on the four-target tracking task of Experiment 1 is shown as a function of visual angle of the display. There were four responses on each trial and the vertical axis shows the percentage of these which were correct identifications of the original targets. The horizontal dashed line at 44% indicates chance levels of performance (choose four randomly of nine). The vertical dashed line (far right) indicates the limit predicted by spatial resolution. The fit of a simple model (see text) to the average data of the four observers is shown as a solid line. The "X" marks the limiting display size at which the model predicts that performance reaches chance level (0.57° or 34 arcmin).

measures of spatial resolution. At this size, when they are packed at their closest spacing, the disks form an array spaced at 60 disks per degree.

From these data we would like to derive a critical radius which bounds the smallest region of attentional selection in this tracking task. At intermediate levels of tracking, say 75% correct, performance is probably determined by a mixture of factors, including both attentional load and attentional resolution. We therefore looked to the absolute limit of performance which we believe will be set only by resolution. Specifically, the level of chance performance indicates the density beyond which no selection at all is possible. For less dense displays, some tracking is possible and the final performance may then also be modulated by other factors such as overall attentional load.

Model of the Region of Attentional Selection

The data of Fig. 4 are clearly descending along a steep, fairly linear gradient toward chance as the display size drops below 1° . This descent could be extrapolated until it reaches the level of chance (44%) without invoking any model, but we offer one in order to make a reasoned estimate that we can compare to estimates from the data in our second experiment. The particular model we have chosen is overly simple but it does capture underlying properties of the data and can be extended to the data of the second experiment. To determine the display size at which accuracy is at chance levels, we used the model for the spatial profile of the region of attentional selection shown in Fig. 5. We assume a model with a central region, where selection is certain,

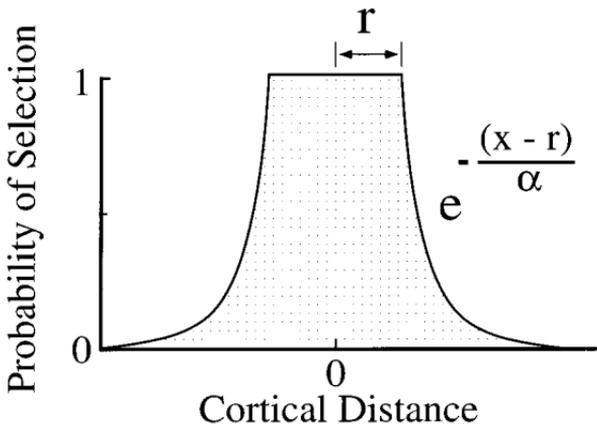


FIG. 5. The model of the region of selection. The vertical axis gives the probability that an item will be selected and the horizontal axis depicts the distance from the center of the region. All items within the central region of radius r are selected with 100% certainty. Items falling in the immediate surround are selected with an exponentially decreasing probability where e to the power of $-(x - r)/\alpha$ determines the rate of drop-off. In this model, if two items are selected by this selection operator then no further individuation of the two is possible. They are treated as one entity.

surrounded by a gradient of selection (e.g., Downing, 1988), where the probability of selection drops as a negative exponential function of distance. The results of Pylyshyn's experiments (Pylyshyn, 1989; Pylyshyn & Storm, 1988) suggest that attention may be selecting items using as many as four of these regions, but we concentrate on a single selection region for the moment. For attention to successfully select and individuate a given item, the item must be alone within the central region. If two or more items are within the central region, they are all selected and no individuation is possible. If one item is in the central region and another in the surround, then there is some possibility that more than one item is selected and individuation is lost.

For an object located at distance x (on the cortical representation) from the center of the selection region, the probability of selection $S(x)$ is given by the following:

$$\begin{aligned} S(x) &= 1 && x < r \\ S(x) &= e^{-(x-r)/\alpha} && x \geq r, \end{aligned} \quad (1)$$

where r is the critical radius and α governs the rate of drop-off beyond the central region. Errors are predicted by assuming that tracking processes maintain a to-be-tracked target at the center of the selection region and that additional items flank the tracked item at a spacing x . The value of $S(x)$ in Eq. (1) gives the probability that a flanking item will be selected in addition to the central item. When this occurs, the target is no longer uniquely identified.

We make two assumptions to simplify our modeling of tracking: one concerning the scaling of resolution with eccentricity and one concerning the loss of tracking.

First, we assume that attentional resolution becomes coarser as items move into the periphery. However, we consider the selection region in terms of its size *on the cortex* and here we assigned it a constant width that does not change with eccentricity. Because of the mapping between the retina and the cortex (e.g., Schwartz, 1980; Johnston, 1986), this fixed size is equivalent to an increase in selection area on the retina with increasing eccentricity. Analogous claims of a fixed processing scale across the visual cortex have been made previously for visual resolution (Virsu & Rovamo, 1979) and motion processing (Johnston & Wright, 1986). In our tracking task, the disks are distributed randomly on the display with no systematic variation in size or spacing with eccentricity. However, on the cortex, the farther the disks are from the fovea, the more densely they are packed together. Consequently, we apply our model to the spacing of the disks at the edge of the display because that is where they are most likely to get closer than attentional resolution limits and trigger tracking errors. The randomly moving disks actually spend a fair amount of time at the outermost parts of the display simply

because this region accounts for a significant portion of the display area. For example, the outermost 30% of the eccentricities covered by the display account for more than half of its area.

What happens when observer moves closer to the display, increasing its size on the retina? Increasing the size of the display increases the disk size and spacings uniformly on the retina. The disk spacings also increase on the cortex, although by a smaller proportion (e.g., Johnston, 1986). Resolution and tracking should therefore become easier at closer viewing distances.

Second, we assume that the probability of losing the target is governed by its single closest approach during tracking and that if it is lost, the observer abandons the target without picking up a new one. Subjective reports suggest that when a target is lost, observers are often acutely aware of the failure. They can attempt to replace the lost target with one of the items in the region, but most often, they simply give up on that item and concentrate on the remaining items. Consequently, we have assumed that once a target is no longer uniquely individuated, the response for the lost target is based only on guessing.

With these two assumption, the values of $S(x)$ in Eq. (1) give the probability of tracking loss as a function of cortical spacing. We also allow the best performance to be asymptotic to a value less than 100% to reflect response mishaps and memory loss during the response period. The three parameters were set to minimize the root-mean-squared deviation between the predictions and the data.

We took the critical item spacing during tracking (the closest spacing which occurs at least once for every disk, see the Monte Carlo simulation below) and converted it to the corresponding cortical spacing near the edge of the display. To map the critical spacing along the radial direction from the fovea into cortical coordinates we used the function adapted from Schwartz (1980) as follows:

$$x = k \ln(E + E_2) \quad (2)$$

where k is a constant of proportionality, x is the position on the cortex; E is eccentricity on the retina; and E_2 , a constant, is the eccentricity at which the resolution halves (0.18° along the radial direction; from the crowding study of Toet and Levi, 1992, Table 1). This allowed us to generate a prediction of the probability of correct tracking at each display size and to extrapolate this fit to predict the spacing for chance performance.

As many others have noted, this mapping function predicts that resolution (and in our case attentional resolution) should scale linearly with eccentricity. The derivative of Eq. (2) with respect to E is called the magnification factor: the size of the step on the cortex that corresponds to a fixed step on the retina at a given eccentricity (typically in millimeters of cortex per degree visual angle). The inverse of the magnification factor gives us the relation we want:

R , the spacing on the retina at each eccentricity, E , that maps to the same fixed radius, r , of the attentional selection region at all eccentricities on the cortex (the formula is approximate until r approaches 0) as follows:

$$R(E) = \rho(E + E_2),$$

where ρ is a constant function of r . Or, rearranging terms, we have the following:

$$R(E) = \rho E_2(1 + E/E_2). \quad (3)$$

This form of linear scaling is widely reported for many types of resolution measures, including visual resolution, vernier acuity, and crowding (Levi, Klein, & Aitsebaomo, 1985; Toet & Levi, 1992). Data from our second experiment will confirm this relationship in the case of attentional resolution, supporting our use of this mapping function here.

In Fig. 4, we see that the model (the black curve) predicts that performance reaches chance levels (44% correct) at a display size of about 34 arcmin. The fit of the model is reasonable. However, this is to be expected, as there are three free parameters in the model and only three or four points that actually matter to the fit of the model. Many alternative characterizations of the region of selection and the tracking process could give as good or perhaps better fits to our data and we have explored several of these alternatives. There are insufficient data in our experiments to make meaningful comparisons between models, however, and we do not propose ours as the sole or best explanation of the results. We use it as a plausible mechanism for extrapolating performance to chance levels and one that can be extended to the data of the rest of the experiments in this article. Significantly, the predicted display size for chance performance remained virtually constant across the models that we examined. Clearly, the data at the smallest display size we tested are already so close to chance levels that all models, including our own, are strongly constrained to predict chance performance for display sizes only slightly smaller—at or near 34 arcmin.

Finally, although our model has its limitations, our assumption that errors are determined only by the single closest approach is appropriate at the critical level of chance performance (the level that we will be comparing across conditions and experiments). Even several approaches each with some possibility of escape do not add up to chance performance; only an approach from which there is never an escape does this. Given this restriction, all less catastrophic approaches can be ignored in predicting the spacing that leads to chance performance. This property is captured by our analysis of the single closest approach.

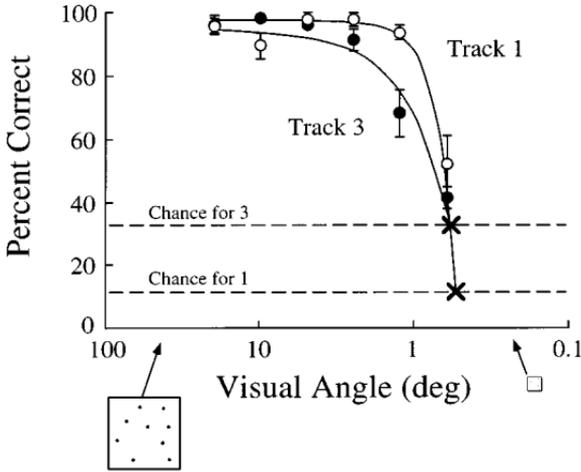


FIG. 6. The performance in tracking one or three disks while fixating the center of the display. The data are averaged over four observers and standard errors (± 1 SE) are shown as vertical bars where they are larger than the data symbols. The fit of the model to each set of data is shown as a solid line. The horizontal dashed lines at 11 and 33% indicate chance levels of performance for the two conditions and the Xs mark the limiting display size at which the model predicts that performance reaches chance level.

Tracking Task, One and Three Targets

Our choice of chance level as the critical measure was based on our conjecture that the chance level might be limited only by resolution, independently of attentional load. This conjecture is tested in our data, averaged over four subjects, on tracking one versus three targets (Fig. 6). As expected, performance for three targets with fixation is reasonably similar to that with four targets in free viewing. More important is the performance with a single target, which remains above that for three targets at the midrange of display sizes but drops more rapidly toward approximately the same limit at the smallest sizes. Clearly, in the midrange there is an advantage to tracking only one item. With more resources available to keep track of the one target, difficult close encounters are probably dealt with more successfully.

Ultimately, however, as the display size approaches the smallest size we tested, performance with one target drops rapidly, suggesting that even the extra attentional resources available when tracking only a single target cannot overcome the spatial limits. The solid curves (fits of the model) intersect chance performance when the display size reaches 35 arcmin while tracking three targets and 32 arcmin when tracking a single target. All of the critical sizes for chance performance (including the 34 arcmin for tracking four targets, Fig. 4) are quite similar despite variations in attentional load. This is evidence for independent spatial and capacity limits to attention. Even with all the available attentional resources applied to tracking a single target, accu-

rate tracking does not approach the limit of visual resolution, which is still finer by a factor of 4 or 5.

Spatial and Motion Control Tasks

These two control tasks were conducted to ensure that the limits found in the main tracking task were not due to a difficulty in visually resolving the items or their motions. In particular, these tasks were used to estimate the minimum center-to-center spacings required for accurate visual (as opposed to attentional) resolution of the positions and motions of just two disks. The results of the two tasks are shown in Figure 7.

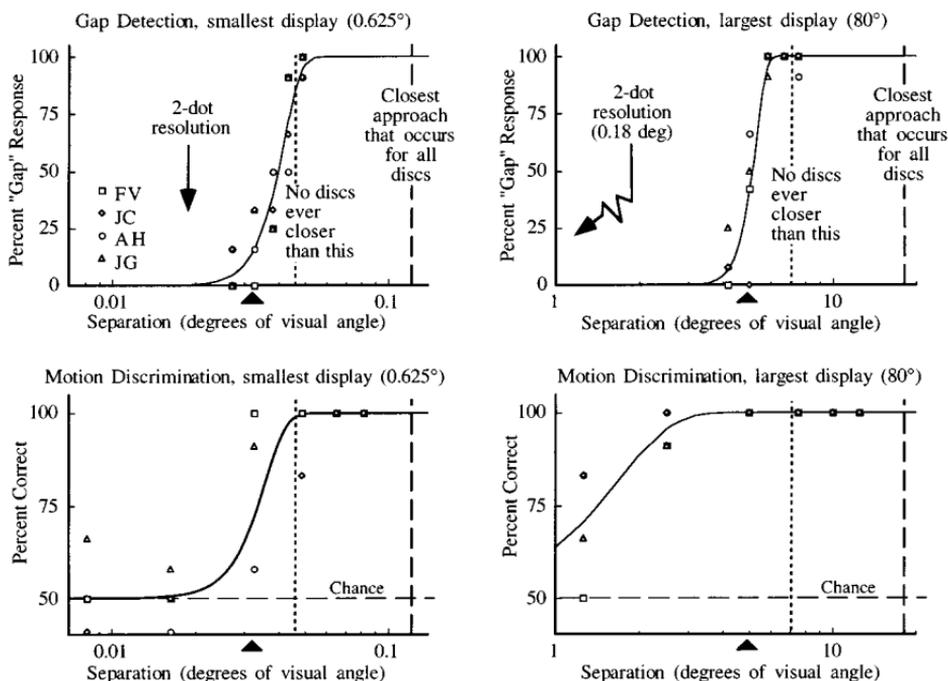


FIG. 7. The performance of four observers on the spatial and motion control tasks. The percent of "gap" responses in the gap detection task are plotted as a function of center-to-center separation for the four observers in the upper two panels. The percentages of correct responses in the motion discrimination task are plotted in the bottom two panels. The smooth curves are fits of a Weibull function to the data averaged across the four observers. In each panel, the rightmost vertical line indicates the closest spacing that occurred at least once for each item in each trial in the 180 Monte Carlo runs. This spacing is therefore the smallest spacing that is frequent enough to be responsible for losing all tracked balls on every trial (chance performance). The vertical line to the left of this indicates the absolute minimum spacing: no two disks in the tracking task ever got this close or closer on any of the 180 Monte Carlo runs. The center-to-center spacing at which the two disks just touched is shown by the \blacktriangle on each horizontal axis. Two-dot resolution thresholds for the eccentricities of the two displays are shown on the Gap Detection graphs as derived from data (shown in Fig. 12) from Yap, Levi, and Klein (1989) and Wertheim (1887). Note that the two-dot resolution for the large displays is 0.18° of visual angle, off the scale to the left.

To estimate the frequency with which different spacings actually occurred in the main tracking task, we ran a Monte Carlo simulation of 180 tracking runs (each of which was similar to one of the 5-s tracking intervals in the main task). In each run, all the distances between all disks were collected and analyzed to find the closest approach that occurred for each disk. The simulation showed that it was *never* the case that even one of the disks approached closer than 1.24 cm (center-to-center) to any of the other disks. More important, the closest approach that occurred at least once for every disk on 99.8% of the trials was 3.25 cm (center-to-center). This is a rough estimate of the critical approach radius that would occur often enough to lead to the loss of all four tracked disks on every trial (and thus result in chance performance). These two values of critical spacing, 1.24 and 3.25 cm, are shown as vertical lines on the graphs in Fig. 7 (in degrees of visual angle at each viewing distance) along with the observers' performance data.

The top two panels of Fig. 7 show that all observers could resolve the gap between the two disks with 90% accuracy or better at the closest center-to-center spacing that ever occurred on any trial. If the observers can detect the gap between the disks, they can undoubtedly resolve the two disks as separate items. Our gap detection task is a very conservative estimate of the ability to resolve two disks, so it is quite likely that even in the few cases where observers might fail to detect the gap at spacings that occur in the tracking task, they can still determine that there are two disks. Our gap detection task is limited in that it cannot measure visual resolution any finer than the spacing at which the disks touch because at that point there is no longer a gap to detect (the spacing at which the disks just touch is marked with a triangle on the horizontal axes in Fig. 7). Studies of two-dot resolution (as opposed to gap detection) show a lower threshold, and values derived from Yap, Levi, and Klein (1989) and Wertheim (1887) are shown in Fig. 7. Based on these data and our own, we conclude that the disks were always resolved at all spacings that occurred in the tracking task.

In addition to testing gap detection, we also tested the resolution of motion for the two disks. The rationale was that in conditions where positions cannot be resolved, motion might carry sufficient trajectory information to support continued tracking. In particular, a target might be tracked through a close encounter with another disk even if position information itself was insufficient. The data (bottom two panels of Fig. 7) show that, with the exception of one datum point for one observer, the direction of motion of the two disks could be discriminated with 100% accuracy at all spacings that occurred in the tracking trials. The critical finding of the two control tasks is that both the relative position and motion of any pair of disks in the tracking experiments were always within the range that can be visually resolved by the observers.

Considerations of the spatial frequencies and velocities of the display also

support the contention that all stimuli were easily resolvable. The disks used in these experiments had a minimum center-to-center spacing that was never less than 1.24 cm and traveled at a speed of 10 cm/s. In the largest display, this disk spacing corresponds to a spatial frequency of approximately 0.14 cycles/degree and the speed of each disk is 45°/s. Both these parameters lie well within perceptible limits (Burr & Ross, 1982). Indeed, in the smallest display (0.625°) the disks, at their closest typical spacing, correspond to a spatial frequency of approximately 22 cycles/degree—a value that is about 3 times lower than the maximum spatial frequencies that are resolvable in the fovea (Campbell & Robson, 1968).

Discussion

When tracking four disks, performance drops as the visual angle subtended by the display decreases. This result clearly demonstrates that attentive tracking is strongly influenced by display size and is poorest at small viewing angles. This is true even though observers can resolve the items and their motions, as shown by our control tasks. Therefore, the poor tracking performance that observers exhibit at small viewing angles seems not to be a function of limits imposed by the early processing of visual information. Instead, we claim that these limits are a property of attentive mechanisms.

Two aspects of attentive mechanisms could have contributed to the tracking failures for closely spaced targets. On one hand, as we suggest, the center-to-center spacing could be finer than the resolution of attention so that targets can no longer be individually selected and tracked. On the other hand, the processing overhead required to keep track of the multiple targets and resolve close encounters may exceed the limited capacity of attention. To examine the contribution of these two factors we ran the tracking task comparing performance for one and three targets. Clearly the overhead for attention was greater for three targets and, as the display got smaller, performance dropped first for the three-target condition while staying quite accurate in the single-target case. However, at still smaller sizes, performance with a single target plunged rapidly so that both conditions approached chance performance at approximately the same display size, 33 arcmin. This finding suggests that the ultimate limit on performance is imposed by the resolution of attention, a limit which exists independently of the capacity limits of attention.

To summarize the results of this first experiment, the failure of attention to adequately track items appears to be due to an inability of attention to grab onto or individuate the separate items of the display when they are spaced too closely. This is true even though the items can be visually resolved as multiple disks. These results imply that attention has a grain that is coarser than that of spatial vision.

EXPERIMENT 2: INDIVIDUATION IN STATIC DISPLAYS

In the tracking experiments, items move randomly over a range of retinal locations. Because tracking accuracy is evaluated only at the end of a trial, it is hard to know the location and configuration of items that led to each tracking failure. The next experiment measures attentive performance for specific retinal locations and local item densities. Moreover, because the task involves attending to only a single item at a time, the role played by the limited capacity of attention has been reduced.

One simple test of individuation is to ask observers to count the number of items in a display without making any eye movements. An accurate count might require that each element be individuated in turn and, to avoid re-counting items, that some trace of the counting path be kept. However, after some pilot work, we felt it was important to avoid counting as the task. There were two main problems. First, with small numbers of items, it was possible for observers to simply apprehend the number of items without counting (i.e., they could subitize the items). Furthermore, even when the number of items exceeded the limits of subitizing, it was possible for observers to view the density of the items and to estimate the number of items in the display fairly accurately. In our pilot tests, it was very difficult for observers to avoid these alternate strategies and the resulting counting performance was not a robust measure of individuation. As an alternative, we developed a "stepping" technique. In one condition, for example, disks were arrayed around the circumference of a circle and one disk was identified (by a color change) as the starting point. The target disk then reverted to its original color and a series of auditory instructions was given by the computer. The instructions told observers to select (i.e., shift their attention to) the adjacent disk on one side or the other of the currently attended item (a sequence might be, for example, "left," "left," "right," "left," "right"). A test disk was then highlighted and the observer had to indicate whether this was the disk on which the stepping sequence had ended. This method makes it easier for the observer to be aware of and avoid alternative strategies.

In our stepping technique, a static display was used to measure the individuating ability of observers at different locations within the visual field. Observers' performance as a function of density was measured separately in different parts of the visual field (upper/lower) and at different eccentricities. Additionally, because preliminary studies indicated that individuating abilities may differ depending on stimulus configurations, we used both a radial and a tangential configuration of items.

Methods and Procedures

Observers. Five observers were volunteers with normal or corrected-to-normal visual acuity and color vision. The observers were recruited from the Harvard University community and participated in eight 45-min sessions. Four males and one female participated and all (except

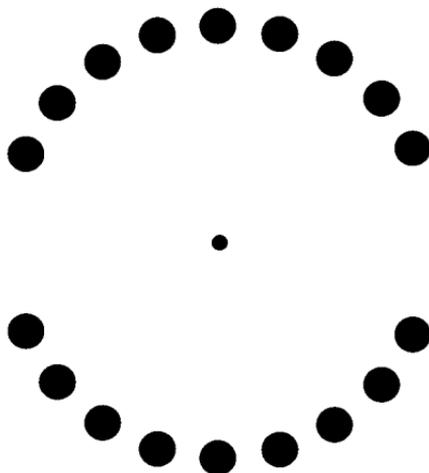


FIG. 8. A schematic drawing of the tangential stimulus used in Experiment 2. This trial has nine disks arrayed in the 140° arc in the upper field and nine arrayed in the 140° arc in the lower field. On a single trial, the attentional “steps” are constrained to always remain entirely in either the upper or the lower visual field.

J1) were naive to the purposes of the experiment. The observers ranged in age from 22 to 35 years.

Apparatus. The hardware was the same as that used in the earlier experiments. Two different displays were used: a tangential display and a radial display. A schematic of the tangential display is shown in Fig. 8. In all displays a screen-centered fixation point appeared on a gray field.

Disks were displayed around an imaginary circle centered at the vertical midline. A gap was always present near the horizontal midline so that the upper and lower arc each covered 140° . An equal number of white disks were spaced regularly within the upper and lower arcs. On each trial there were either 5, 7, 9, 11, or 13 disks in each of the arcs—Figure 8 illustrates a tangential display containing 9 disks in each arc. The size of the imaginary circle (and thus the individual item eccentricities) was varied between blocks and had a diameter of 1° , 7° , or 30° of visual angle—in these conditions the items had eccentricities of 0.5° , 3.5° , and 15° , respectively. The individual items were scaled directly with eccentricity such that they had radii of 0.05° , 0.35° , and 1.5° , respectively.

Figure 9 illustrates the radial display. In this display a variable number of disks (8, 12, or 16) appeared along an imaginary line radiating from the screen-centered fixation point, divided equally on either side of the fixation. The disks were scaled in size such that the disk closest to fixation (0.7° out) was the smallest (0.2°) and the disk farthest from fixation (7.0°) was largest (1.5°). The angle of the imaginary radial line was varied such that the disks either crossed from the top-right, to the bottom-left, or from the top-left to the bottom-right of the display. Although the disks extended in both directions from fixation, only one half, from fixation outward, was tested in a given trial.

Procedures. All experiments were conducted in a dimly lit room. At the beginning of each trial a single disk changed color from white to red for 1 s. The color change indicated that the observer should attend to that disk. The disk then changed color back to white. Subsequently, the observer heard a series of computer-issued verbal commands, one command every 1.5 s. In the tangential condition the series consisted of five to seven words, each of which was either “left” or “right.” In the radial condition the series consisted of five to seven words, each of which was either “in” or “out.”

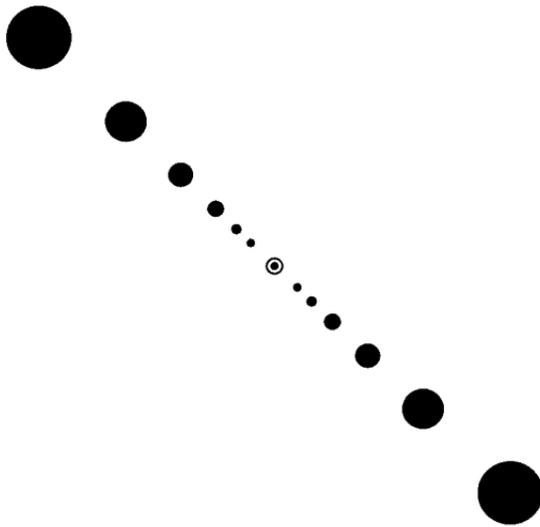


FIG. 9. A schematic diagram of the radial stimulus used in Experiment 2. This trial had six disks on each side of fixation. On a single trial, the attentional “steps” are constrained to always remain entirely in either the upper or the lower radial arm.

Observers were asked to shift their attention (based on the verbal command) to an item adjacent to the currently attended item. For example, assume the observer began by attending to the innermost disk in Fig. 9 and heard the commands “out, out, in, out, in, out, out.” After shifting their attention through this series they should have been attending to the fourth disk from the center. After the observers mentally “stepped” as instructed, a single disk would again change to red. On half of the trials, this “probe disk” would be on the location where the observer’s walk should have ended, whereas on the other trials, it would be one position off the final disk (in either direction). Observers then indicated (with a keypress) whether this probe disk accurately identified the ending position of their walk.

In the tangential walks, observers were never required to move their attention over the horizontal midline—half the walks remained in the upper visual field and the other half remained within the lower visual field. Similarly, in the radial condition, all walks remained on one side of the fixation. All trials within a condition were presented in a random order. The radial condition and the three tangential conditions were all run in separate blocks. Within each condition (tangential/radial), for each density, and within each visual field (tangential) or quadrant (radial) observers took 32 “walks.” Observers were encouraged to carefully focus their attention on each item that they “walked” over and to avoid using any alternative strategies.

Results

The data of each observer was smoothed and interpolated to extract the density leading to 75% performance. These individual thresholds were used in a two-way, repeated-measures ANOVA to examine the effects of eccentricity and upper versus lower field. The percentage correct in each condition was also averaged over the five observers and fit using the model as described below to predict densities for other performance criteria.

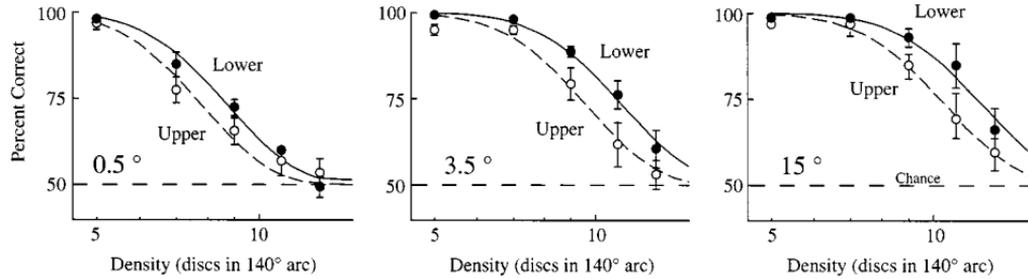


FIG. 10. Average performance for five observers on the tangential individuation task at 0.5° (left), 3.5° (center), and 15° (right) eccentricity. The dashed line at 50% indicates chance levels of performance. Filled symbols show results for tests in the upper field and outline symbols for the lower field. The smooth curves show the fit of the model to the data (see text). Standard errors ($\pm 1 SE$) are shown as vertical bars where they are larger than the symbols.

Tangential conditions. Figure 10 shows the average results of the individuation task for the five observers in the tangential conditions. This graph illustrates percentage correct as a function of the number of disks arrayed in each 140° arc. Notice that there is a visual field asymmetry such that performance is somewhat better in the lower visual field. Also, performance for a given disk density drops as the circle's radius decreases.

Figure 11 gives the density of disks that allowed 75% performance (aver-

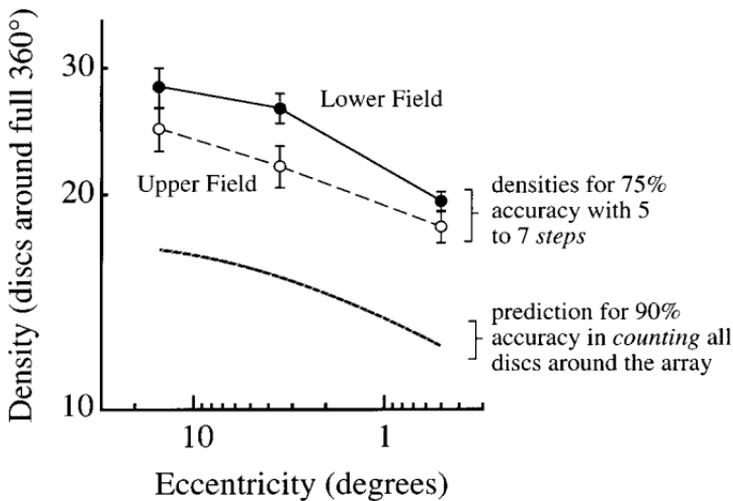


FIG. 11. The density of items leading to 75% correct performance in the stepping task is plotted as a function of eccentricity for upper and lower visual fields separately. Note that density is now given in terms of the number of disks that would fit around a full 360° using the spacing of the test disks around the 140° arc. The bottom trace shows the predicted density of disks that would support 90% accuracy of counting of a set of disks that are uniformly spaced around a 360° display. The prediction takes into account the effect of eccentricity and the differences between selection in the upper and lower fields.

aged over five observers) following the five to seven steps of the tasks. Results are plotted against eccentricity for the upper and lower fields separately. In order to achieve this 75% level of final performance, each of the five to seven steps in the "walk" needs to succeed about 95% of the time. Note that density has been converted from the number of disks displayed in the original 140° test arc to the number of disks at this same spacing which can be placed around a full 360° array. The threshold density dropped significantly as eccentricity decreased, $F(2, 8) = 21.33$, $p < .001$. There was also a small but significant advantage for the lower field, $F(1, 8) = 7.966$, $p < .05$. The interaction between eccentricity and visual field was not significant, $F(2, 8) = 1.69$, *ns*.

We extended our model of the region of attentional selection [Eq. (1)] to this task to extract predictions for other performance criteria. The fits of the model to the psychometric functions are shown as solid and dashed lines in Fig. 10. In order to adapt the model to the stepping task, we had to take account of the cumulative effects of five to seven sequential selections of adjacent items. We continued with our assumption that the target was tracked until selection was lost at which point tracking was abandoned and the observer merely guessed the final response. In this simple case, the probability of being correct at the end of several steps is the product of the probability of being correct on each step added to the probability of losing the target but making a correct guess. We also examined a model in which tracking continued after an error and could, with an appropriate sequence of errors, regain the original target. This model's predictions did not differ in any important way from our simpler version. The multiplication of probabilities across steps changes the form of the model's predicted curve from the final step decrease that we saw in Experiment 1 to a sigmoidal shape. The parameters were fit by minimizing the root-mean-squared deviation.

Since each condition was run at a fixed eccentricity and spacing varied only tangentially, the relation between physical spacing and cortical spacing is constant within each condition (ignoring meridional variations). We could therefore set the model parameters in physical units, bypassing the scaling from display spacing to cortical spacing. The scaling of attentional selection area with eccentricity that we find can then be compared to the linear increase with eccentricity expected from the retina-to-cortex mapping function we used in Experiment 1. The size of the central selection area (r in Fig. 5) was allowed to vary for each condition. One additional parameter then set a common "shape" for the selection function for all the conditions. This shape is set by the ratio between the central area of complete selection and the rate of drop-off of the surrounding gradient (r/α). By holding this ratio constant, we ensure that if the central selection area doubles in size, so will the size of the "skirt" of graded selection.

Based on the performance for the five to seven steps of the task, we used our model to predict the density at which all the disks around a full 360°

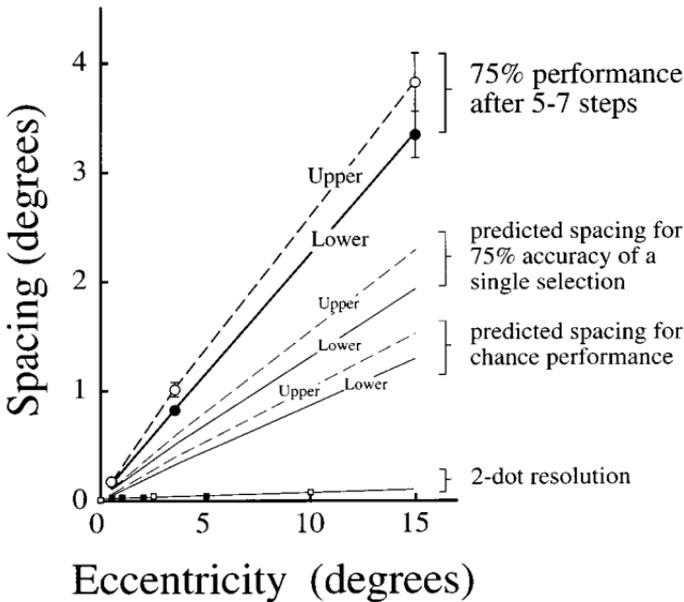


FIG. 12. Center-to-center separation (degrees of visual angle) leading to 75% correct as a function of eccentricity (○, upper field; ●, lower field). The predicted separation of disks which would support 75% accuracy in each selection step and then chance performance in each selection is shown in the middle sets of lines (see text for description of model). For comparison purposes, the solid line at the bottom shows the spacing required for two-dot resolution, fit to the data of (■) Wertheim (1887) and (□) Yap, Levi, and Klein (1989).

display can be counted with a final accuracy of 90%. This prediction is shown as the bottom curve in Fig. 11 and suggests that fewer than 20 disks can be counted with 90% final accuracy.

The threshold data from all three eccentricities are replotted in Fig. 12 in terms of the center-to-center *separation* in degrees of visual angle required for 75% correct performance (solid lines). We also plot the model's predictions for the separations which lead to 75% performance in each individual step of the sequential task (middle two curves). The performance for a single selection is probably the most appropriate to compare to the two-dot resolution values from the literature, which are shown as the bottom trace on Fig. 12. Clearly, the difference between visual resolution and attentional resolution increases dramatically with increased eccentricity. Attentional resolution is coarser than visual resolution by a factor of about 5 at 0.5° eccentricity but by a factor of about 20 at 15°. Finally, Fig. 12 also shows the model's predictions for the separations that lead to chance performance in the stepping task. Even these spacings are up to an order of magnitude coarser than the finest spacing that can be visually resolved.

The format of Fig. 12 is a frequently used convention for plotting the effects of eccentricity on resolution and for recovering two standard param-

ters that characterize resolution: T_0 , the estimated threshold extrapolated to the fovea; and E_2 , the distance from the fovea at which threshold spacing doubles from its foveal value. We fit our data with the standard regression equation [Eq. (4)] and recovered these parameters for both the upper and lower fields as follows:

$$T(E) = T_0(1 + E/E_2), \quad (4)$$

where $T(E)$ is the threshold spacing at eccentricity E .

Note that the form of Eq. (4) follows closely that of Eq. (3), which was derived to predict the scaling of resolution with eccentricity based on the mapping function of Schwartz (1980). The data of Fig. 12 do show a reasonably linear relation to eccentricity, lending support to our use of the mapping function in the analysis of Experiment 1.

From Eq. (4), the estimate of foveal threshold spacing in our stepping task is 5.2 arcmin for the upper field and 3.5 arc min for the lower field. Of course, the notion of upper and lower field is not relevant right at the fovea and even if it were, multiple items cannot fit in the fovea at nonzero spacings. Perhaps more appropriate is the threshold for 75% accuracy in each selection. This level of accuracy would lead to very low performance after the five to seven selection steps of our task but it does match the typical threshold criterion of a single resolution judgment. We used our model to derive these values, and they are shown in Fig. 12, middle traces. Entering these individual threshold spacings into Eq. (4), we get critical spacings in the fovea of 3.0 and 2.6 arcmin for the upper and lower field data, respectively. Two-dot resolution at the fovea is about 1 arc min, which is finer than this estimate of foveal attentional resolution by a factor of 2 or 3. Our value is quite similar to the critical spacing obtained from a "flanking" task reported by Toet and Levi (1992). They presented three Ts in a row and asked observers to report the orientation of the center T (each T was oriented up, down, left, or right). They measured the separation between the Ts that permitted 75% accuracy. At the fovea they found an average threshold spacing of 3.6 arcmin for reporting the center T.

Of equal interest are the values for E_2 , the eccentricity at which thresholds in the tangential direction double from their values at the fovea. The value for the upper field is 0.34° and for the lower field 0.27° for the 75% thresholds in the stepping task. For the computed 75% accuracy of each individual selection, the values of E_2 are only slightly different: 0.33 and 0.35° for upper and lower fields, respectively. These values indicate that the resolution of attention decreases extremely rapidly as stimuli move into the periphery. For comparison, the drop-off with eccentricity in resolution required to discriminate one from two dots, to judge the orientation of a single letter, or to discriminate the orientation of a grating is much more gradual. These resolution functions do not drop by a factor of 2 until reaching 2° to 4° of eccentricity

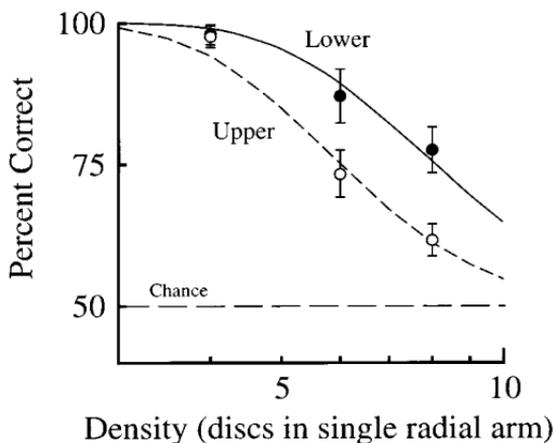


FIG. 13. Average performance (± 1 SE) for five observers on the radial stepping task. Both upper (\circ) and lower (\bullet) visual-field data are presented. The line at 50% indicates chance levels of performance. The smooth curves indicate the fits of the model.

(Johnston, 1987; Toet & Levi, 1992; Virsu, Naesaenen, & Osmoviita, 1987; Yap, Levi, & Klein, 1989; Wertheim, 1887). Nevertheless, our values for E_2 are again quite similar to those obtained from the crowding task reported by Toet and Levi (1992). In the tangential direction tested by our displays, Toet and Levi found an E_2 value of about 0.34° for the lower visual field.

Radial task. Figure 13 shows the results of the radial stepping task as a function of the number of items in the tested side of the display. In this radial condition, as in the tangential condition, there is a lower visual field advantage: The average 75% threshold is 5.95 disks in the upper field and 8.30 in the lower field, $t(4) = 8.53$, $p = .001$. To fit our model of the region of selection to this task, we considered our constant ratio spacing (Fig. 9) as converted to a relatively uniform spacing on the cortical surface. [Specifically, with an E_2 of 0.18° in the radial direction (Toet & Levi, 1992) and with 0.7° the smallest eccentricity of a disk in our display, a simple log mapping where constant ratios on the retina become uniform spacing on the cortex is a reasonable approximation to the mapping function of Eq. (2)]. We then used the constant radius of selection for all eccentricities covered by the display, as before. With only three data points per curve, there are few constraints on the model, so the fits over the densities shown on the graph are reasonable.

How do the critical spacings in the tangential displays compare to those for the radial displays? The critical ratio for spacing found in the radial displays does specify a spacing at each eccentricity from 0.7° to 7° but the data are probably most representative of the spacing for the midrange of the eccentricities spanned. The mean eccentricity for the radial display is similar to that of the 3.5° tangential condition. Computing the density at the midpoint

of the radial display gives spacings for 75% correct of 1.69° and 1.12° for the upper and lower fields, respectively, compared to 1.01° and 0.83° for the tangential case at 3.5° of eccentricity. This indicates that the resolution is better in the tangential direction than in the radial direction by about 50% [averaged over both fields, $t(4) = 4.97$, $p = .0076$]. A 100% advantage, also in the tangential direction, was reported by Toet and Levi (1992) in their flanker interference task.

Discussion

Effects of eccentricity, visual field, and axis. Our experiment shows that the resolution of attention is not uniform over the visual field, but worsens dramatically as eccentricity increases. The selection process underlying our task loses half its resolution in the first $\frac{1}{3}$ of a degree from the fovea. Selection is not homogeneous over the visual field but neither does it follow the properties of low-level, visual resolution. In particular, the resolution of dot pairs, individual letters, or gratings falls off about 6 to 12 times less rapidly than our individuation task (Johnston, 1987; Toet & Levi, 1992; Virsu et al., 1987; Yap et al., 1989). Our results are consistent with other studies that have examined selection in crowded displays (e.g., Toet & Levi, 1992; Wolford & Chambers, 1984): When multiple items are present, selection is coarse and drops off steeply with eccentricity.

The results of Experiment 2 and the study by He et al. (1996) also indicate that attention has a finer resolution in the lower visual field than in the upper field. We found an advantage for the lower field of about 17% in the tangential conditions and 50% in the radial case. Furthermore, we find that individuation is more difficult in the radial dimension than in the tangential dimension. The radial/tangential asymmetry has been observed in other experiments, most strongly in experiments involving adjacent distractors (e.g., Chambers & Wolford, 1983; Toet & Levi, 1992). Tasks assessing low-level visual and motion resolution also show a radial/tangential asymmetry favoring tangentially arrayed patterns. However, the advantage only becomes significant at 20° of eccentricity or more and is absent at the 3.5° of eccentricity that we have examined here (Rovamo, Virsu, Laurinen, & Hyvarinen, 1982).

These results demonstrate that selection is not homogeneous across the visual field but recall that our model specifically proposed a constant region of selection at all eccentricities. In the model, however, the region of selection was a constant size in the *cortical* representation of visual space, not in visual space itself. This means that we have assigned the measured inhomogeneities in resolution to the mapping of visual space onto the cortex where selection is occurring. For example, there should be a larger cortical representation of the lower than the upper field and the representation should be stretched in the tangential direction (perpendicular to the meridia) relative to the radial direction (along the meridia). If our assumption of constant

regions of selection is correct, then these inhomogeneities in representation become criteria for identifying the specific cortical area on which selection operates.

Comparing critical densities from Experiments 1 and 2. How does the estimate of minimum spacing from this experiment compare to that found in Experiment 1? Although the displays and the tasks are quite different, a rough comparison is possible. Our model predicted chance-level performance in Experiment 1 for a display size of slightly more than 0.5° (32 to 35 arcmin). The Monte Carlo simulation showed that 3.24 cm was the closest spacing that occurred at least once for all disks and therefore often enough to cause all targets to be lost on every trial, producing chance performance. For the viewing distance at which the 17-cm display subtends 33 arcmin, 3.24 cm corresponds to 6.3 arcmin. In the stepping experiment, the spacing for chance performance can be recovered from our model's fit to the data at the three eccentricities (Fig. 12) and extrapolated to 0.25° , the eccentricity of the edge of the critical display of Experiment 1 where we expect most tracking errors to arise. The predicted spacing for chance performance for tangentially aligned disks is 3.5 arcmin and 3.0 arcmin for the upper and lower field data, respectively. Due to the tangential/radial asymmetry, however, the critical spacing for disks aligned radially should be larger: 5.9 and 4.0 arcmin in the upper and lower fields, respectively. Our data therefore indicate that close approaches along the radial direction in the upper field (critical spacing 5.9 arcmin) would be the most likely to cause errors in the tracking task and, consequently, that the critical spacing derived from the second experiment is reasonable close to that found in the first experiment (6.3 arcmin). Both of these critical spacings are close to Landolt's (1891) original estimate that 5 arcmin was the finest spacing of dots or stripes, right at the center of gaze, which allowed him to count the number present.

What is the maximum number of items that can be individuated from a single display? Ideally, we could combine the constraints found in the present tasks to create a display that would allow us to intuitively appreciate these data. Anstis (1974) has developed a display which allows us to "see" our cortical magnification factor. In his display, many letters were presented around a small central dot. All the letters were scaled to be equally readable when maintaining fixation on the central dot.

We have created a similar display based on our individuation data (Fig. 14). In Anstis' display the items were scaled to allow equal "readability" throughout the visual field. In our display the items have been scaled and spaced to allow equal "individuation" throughout the visual field. In the display, we used critical densities that should provide high accuracy (90%) at counting all the items in any given ring or ray. Thus, this display has items spaced so that they should all be equally and relatively easily accessible to attention. When the central cross is fixated, most observers find that they can attend and individuate any disk at will. Also, most observers can accu-

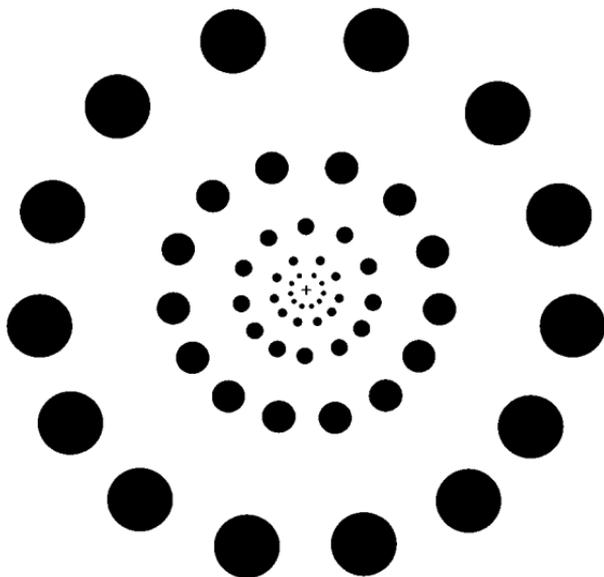


FIG. 14. Seeing one's attentional field. While fixating the center cross, most observers can individuate every item on each concentric circle in this display. To get the best approximation of the conditions of our experiments, hold this figure at a viewing distance of about 14 cm. Individuation should be worse if the figure is either moved farther away or turned upside down.

rately count all the items present in the display (60 in total). To ensure the accuracy of this display, the figure should be held at a distance of about 14 cm so that it occupies approximately 30° of visual angle. When viewed from farther away the task should be more difficult. Note that, as dictated by our experimental results, there are more items present in the lower visual field and, for any item, its immediately adjacent neighbors are more closely spaced tangentially than radially.

Finally, recall that our experiments examined individuation of identical items. Whenever adjacent items differ in some salient feature (e.g., different shapes, different sizes, or colors), they will be easier to isolate even at closer spacings (Kooi et al., 1994). Furthermore, because Experiment 2 only examined eccentricities up to 15° , we can easily add further rings of items at greater eccentricities to Fig. 14 while still preserving individuation performance on the inner rings—in other words, 60 is the maximum number of identical, individuable items in the central 30° of the visual field.

GENERAL DISCUSSION

Three important aspects of attention have been uncovered in the studies reported here. First, both experiments showed that attention has a resolution

that is coarser than that of spatial vision. The first experiment showed that the resolution limit of attentional selection was orthogonal to the capacity limit of attention. The second experiment showed that the grain of attention was inhomogeneous across the visual field: it scaled with eccentricity, it was coarser in the upper visual field, and it was coarser along radial lines from fixation. These are the first detailed measures of the resolution of attention that assess purely the selection of location by attention. Our tasks track location in dense arrays of identical items and so are unaffected by preattentive feature interactions that might degrade the recognition measures used in crowding and flanker tasks.

Attentional Resolution Is Coarse and Inhomogeneous

Although a selection region might scale up to very large sizes to encompass a single large target, or shrink down to individuate a small target (e.g., Eriksen & St. James, 1986), our two experiments showed that the very smallest region that can be selected is much larger than the smallest detail that can be seen. When two items fall within that smallest selection region, they can no longer be individuated. The data of Experiment 1 suggested that the critical center-to-center spacing between items at which individuation and tracking became impossible (chance performance) corresponded to about 6 arc min at 0.25° of eccentricity. That is 5 or 6 times coarser than two-dot resolution at that eccentricity even though the two-dot task has a much higher accuracy criterion (75%). In Experiment 2, spacing for 75% accuracy on each single selection in the stepping task was estimated to be about 3 arcmin right at the fovea. That is 3 times coarser than two-dot resolution at the fovea.

Even more striking than the coarseness of visual attention is its rapid decline with increasing eccentricity. At 15° of eccentricity, for example, the critical spacing for a single selection by attention with 75% accuracy was about 2° of visual angle (Fig. 11). At this same eccentricity, two-dot resolution is about 6 arcmin (extrapolating from the data of Yap et al., 1992; Wertheim, 1887), about 20 times finer. This is one reason why the demonstration of Fig. 1 is so dramatic: At normal reading distance, the middle of the patch of vertical bars is at about 7° of eccentricity. Consequently, the spacing of the bars is well above visual resolution threshold and well below individuation threshold. The difference between the two thresholds is far greater for this peripheral demonstration than it would be for a foveal one.

Although the dissociation between seeing and selecting may seem somewhat counterintuitive, there are other examples of stimuli that can be resolved by one part of the visual system but not another. The “fine-grain motion illusion” (FGMI) is one such example. In the FGMI, two closely spaced items are presented in the periphery of vision. They are set at a spacing which cannot be resolved when the two are presented simultaneously; nevertheless, they produce a sensation of motion when they are presented in succession (Biederman-Thorson et al., 1971). One explanation for the

FGMI is that, at least in the periphery, the motion system has a finer grain than that of the position system. In another example, work by He and his colleagues (He et al., 1996; He, Smallman, & MacLeod, 1995) has demonstrated orientation adaptation following prolonged exposure to grids whose bars are so finely spaced that their orientation cannot be reported. Despite the inability to resolve the lines of the grid, exposure to them affects the visibility of like-oriented grids of slightly wider spacing that can be resolved. Thus, the orientation of the finest grids must have been registered by some part of the visual system that has higher resolution than that available to conscious (reportable) perception.

The coarseness of selection also imposes, in the limit, a space-based character on “object-based” theories of attention (e.g., Baylis & Driver, 1992; Duncan, 1984, 1993; Kahneman, et al., 1992). Specifically, for objects smaller than the smallest available selection region, attention will be unable to select only a single object without also spreading to the adjacent space.

Why Is Visual Selection Coarser Than Visual Resolution?

Why would the resolution of spatial vision be finer than the smallest detail that can be individuated? Or, to put it even more simply, why is the early visual system encoding detail that is unavailable to attention? First, recall that the limit of attentional resolution does not refer to the size of the details being encoded, but to the spacing *between* details. Even an item at the limit of visual resolution is easily individuated if it is alone in the field. Attentional resolution will limit performance only when items are spaced more closely than the minimum required spacing for their particular spatial location and configuration. From the present experiments we can estimate that, near the fovea, a visible but unselectable item must be larger than 1 arcmin in size surrounded by other similar items spaced about 3 arcmin away (about $\frac{1}{20}$ th the width of the index finger held at arm's length). The individuation and tracking of such very small targets surrounded by similar objects may simply be irrelevant to everyday, visually guided performance. But why then should visual details that are more finely spaced than the selection limit be picked up at all by the visual system? One answer is that even though it may not be important to be able to scrutinize individual elements that small, the visible patterns they form can serve as textures that help to classify objects.

The resolution of selection in the fovea may be sufficient for most naturally occurring conditions requiring tracking. However, our data show that the limitations on selection in the periphery are much more severe. In peripheral viewing, the selection of an item is degraded in the presence of even only a few, relatively distant distractors. For example, at only 5° from fixation it becomes difficult to select the central item from a set of three items spaced 1° apart (Fig. 11). Bahcall and Kowler (1999) point out that attention is an exceedingly poor mechanism for enhancing target perceptibility in the periphery—even when the targets are well above acuity limits. They suggest

that, given the availability of saccades, perhaps the only reason one would attend to items outside the fovea is to select the next target for a saccade.

Although moving the eyes to foveate items will bring substantial advantages, it is not an option when tracking multiple, widely spaced targets, nor is foveation desirable in those social contexts where eye contact must be avoided. In these instances, our visually based performance is ultimately limited by the surprisingly coarse resolution of attention in the periphery. It is surprising not only for its coarseness, but also for the extent to which we are *unaware* of this coarseness. Our ability to select relatively fine details from uncrowded fields gives us the impression that we have quite competent visual skills in the periphery. Our data show that this confidence is unfounded.

Of course, bringing a target to the fovea will usually aid individuation and tracking. However, certain strategies of foveating a target may actually be detrimental to tracking. Specifically, moving farther from a display will move items toward the fovea. But, in this case, as the items move closer to the fovea overall item density also increases by an equal proportion. When items and densities are scaled in this manner they actually get harder to individuate as they move toward the fovea (Fig. 11).

Relation to Capacity Limits

Could attentional load and attentional resolution be linked rather than independent? We addressed this question in the tracking task, varying the number of targets to be tracked from one to four. Given that tracking only a single item imposes less overhead than tracking four, we considered the possibility that the surplus resources might be applied to further shrink the attentional field around the target. In this case, there might be no fixed attentional resolution, only a region of selection whose size depended inversely on the resources applied to its maintenance (cf. Lavie, 1995). In the limit, according to this view, attentional selection for single targets might be limited only by visual resolution.

However, our first experiment showed that this did not happen. Tracking was indeed more accurate at moderately small display sizes when tracking only a single target, but the display size at which tracking became impossible was ultimately the same, no matter how many targets were being tracked. This evidence suggests that the spatial limit to selection is an absolute one: It cannot change elastically as more resources become available; it is the same size whether one or four of the regions of selection are being deployed.

Space-Based and Object-Based "Spotlights" of Attention

Research over the past 100 years has suggested that attention may be like a spotlight of extra processing that moves around the visual field (see Cave & Bichot, 1999, for a review). The results from Experiments 1 and 2 seriously challenge this "spotlight of attention" metaphor. Unlike a "spotlight," at-

tention to a region does not simply illuminate the selection region, allowing access to every detail registered within it. To the contrary, selection appears to require that only a single item be present within the selection region in order for it to be individuated and scrutinized. This same point was made by Bahcall and Kowler (1999). In their experiment, the accuracy in reporting 2 items among 24 decreased as the 2 items were placed closer to each other. This was true even though the interitem spacing of the 24 locations remained fixed; only the relative positions taken by the 2 targets changed. As they note, an attentional spotlight or window which improves processing for items within the spotlight predicts better accuracy when both targets fall within a single spotlight. This did not happen.

In our experiments, a spotlight which encompassed a few of the adjacent targets ought to provide access to the details within the "beam." The limit to performance should therefore be the limit of visual resolution. This did not happen. Performance was limited at spacings much coarser than the limits of visual resolution.

Interestingly, the end items of a dense array seem to escape the crowding of the middle items (try attending to the nearest or farthest bar in Fig. 1 while fixating on the plus sign). A simple explanation comes from the geometry of placing selection operators on the array at different locations. If the array is spaced so that even the smallest available selection field covers more than one item, the interior items cannot be individuated. However, the same region of selection can pick up a single item at the edge of the array by offsetting the selection region to cover only that end item and a lot of blank space adjacent to it.

Relation to Cueing Tasks

We noted earlier that cueing tasks in general are not appropriate for measuring the resolution of attention as we have defined it: the smallest spacing at which attention operates. Cueing experiments examined the spread of spatial attention when there is little if anything to constrain it. This was the reason, in our opinion, that estimates of the size of the area of selection vary over 3 orders of magnitude in these tasks. However, one experiment did use displays that were more appropriate for measuring the resolution of attention. In this particular study (Nakayama & Mackeben, 1989), observers performed difficult visual search tasks on extremely small, closely spaced stimuli and the location of the target was cued in advance on each trial. We can make two comparisons to the data of their study. First, in Experiment 4 of their article (their Fig. 10), with ten items around a circular display of 0.5° radius, the accuracy in the identification of the cued target was 87%. The results from our Experiment 2 also show high accuracy (close to 100%) for a stimulus at 0.5° of eccentricity when items had density equivalent to theirs (5 items in our 140° array; see Fig. 10, left panel). Second, in Nakayama and Mackeben's smallest array, with a radius of 8 arcmin, they found 64% perfor-

mance (chance was 33%) with an interitem spacing of 5 arcmin. Because the level of chance performance in our experiment was higher (50%), their result of 64% is equivalent to performance in our experiment of about 75%. We did not have any tests at displays as small as 8 arcmin but we can derive an estimate from Experiment 2 for 75% accuracy in a single selection at this eccentricity [extrapolated using Eq. (4), shown as the middle traces; see Fig. 12]. The estimate is 4 arc min, remarkably close to the result of Nakayama and Mackabben.

Relation to Flanker and Crowding Tasks

Flanker and crowding tasks are well designed to examine the critical spacing for selection and identification of targets. The drawback, as we pointed out above, is that we cannot know whether adjacent distractors act by limiting an item's selection or by degrading the item's representation. This was the motivation we presented for the use of tracking and stepping tasks. Our observers merely had to keep track of position and we attributed the loss of position to the failure of selection mechanisms at close spacings. We feel we have measured, as much as possible, the limits on pure attentional access to location.

Although we have proposed our two tasks as a purer measure of attentional resolution, our data show that the effects of spacing in our tasks are remarkably similar to the effects of spacing measured in the crowding paradigm (e.g., Andriessen & Bouma, 1976; Banks et al., 1979; Kooi et al., 1994; Levi, Klein, & Aitsebaomo, 1985; Walley & Weiden, 1973; Westheimer et al., 1976). In particular, Levi and colleagues (Levi et al., 1985) derived estimates of the size of "perceptual hypercolumns" within which, they claim, adjacent features may interact. Their size estimates agree very closely with our attentional resolution results—both in terms of minimal spacing and in terms of the radial/tangential asymmetry. Because their measurements were only made within the lower visual field, a comparison to our upper/lower visual field asymmetry is not possible.

The fact that similar limits hold for crowding and our tracking/stepping tasks leads us to suggest that it is not the quality of the details available at a location that is limited by adjacent distractors, rather it is the fundamental access to the location. Obviously, if there is no access to the location, there can be no readout of the details. One important objection to this argument is that the nature of the details themselves can determine if there is any crowding (Kooi et al., 1994). Closely spaced items will produce more interference if they are more similar. Such similarity effects could not take place if details were only available following access. However, a closer look at the effect of similarity on crowding reveals that the details of the items may act by triggering a different mechanism of selection.

First, it is not the case that items have to be identical for crowding to occur. The entire crowding literature relies on reporting the identities of adja-

cent stimuli which differ from each other—for example, different letters (e.g., Kooi et al., 1994). In fact, for crowding to be reduced, the adjacent items must differ substantially, say by flanking black items with white ones. We suggest that only differences which support “pop-out” (Treisman & Gelade, 1980) will defeat crowding. For example, a black test surrounded by black distractors must be selected by location—the target is only defined by its position as the center item in the display. However, if as in Kooi et al. (1994), the target is white, but the surrounding distractors are black, it need not be selected by location; it can be selected by feature. The flankers may then be less disruptive.

In displays where the target does not differ from the distractors sufficiently to support selection by feature, there are still two possible sources for the interitem interactions. Given the importance of this argument, let us describe these sources again. First, prior to selection, distortion and mixing of adjacent features (Wolford, 1975; Treisman & Schmidt, 1982) may degrade the representation even if attention succeeds in selecting the intended item. This source may contribute to the critical spacing measured in crowding and flanker studies, but it cannot contribute to the critical spacing in our tracking and stepping studies where target features and identity are irrelevant. If this source is a significant contributor to errors in the crowding and flanker tasks, the critical spacings measured in those tasks should be larger than those estimated by our tracking and stepping tasks. However, as mentioned above, the estimates of critical spacing for the two types of tasks are more similar than they are different. Experiments designed to test exactly this point would be important before drawing a final conclusion. What of the second source? If items are spaced too closely, selection may occasionally pick up an adjacent item in addition to the target. In this case, features may be mixed and degraded *following* selection. If this is the only source of error in the flanker and crowding tasks, critical spacings from these tasks ought to match the critical spacings from our tracking and stepping tasks. To a first approximation, this is what we find.

Cortical Structures Mediating Selection

Our results revealed two properties that help identify the cortical structures mediating attentional resolution: the coarseness of selection and the upper/lower and radial/tangential asymmetries.

Our first indication of the cortical locus of attentional resolution is the coarseness of the selection region. We argue that the resolution is too coarse to be based in cortical area V1. A hypercolumn is the minimal set of cells in area V1 capable of representing all aspects of a small region of space (Hubel & Wiesel, 1974), implying that a hypercolumn, and not an individual cell, would be the minimal useful unit of selection in V1. If attentional resolution were limited by access to single hypercolumns in area V1, what properties would we expect in the tasks explored here?

To begin with, selection based on hypercolumns could explain the encoding of unselectably fine detail by the visual system: the resolution of details smaller than an "attentional receptive field" would be mediated by units in the hypercolumn selective for high spatial frequencies. In this case, some cells of a hypercolumn would respond to the finely spaced bars of a high spatial frequency grating, but no set of cells would respond exclusively to one of the bars in that grating. The finest visible details would, therefore, be resolved as a texture, not as a set of individually addressable elements. In addition, if the minimum size of the selection region were always one hypercolumn, then its size and shape would be fixed in cortical space (each hypercolumn is 1 mm square independent of retinal location). The eccentricity effects found for the resolution of attention could then be attributed to the known mapping of retinal coordinates onto the cortex.

However, several of the present results cannot be addressed by this V1-based, hypercolumn model. For example, we know that there are at least 2400 hypercolumns (1 mm² per hypercolumn and cortical area of 24 cm²; Wandell, 1995) in macaque area V1 (half in each hemifield, the number is probably smaller in humans). But, Experiment 2 indicated only about 60 independently individuatable regions of the central 30° of visual space. Moreover, in a recent study of adaptation and crowding, we (He et al., 1996) concluded that there are situations in which information is registered in area V1 (as shown by orientation-specific aftereffects) but, due to close spacing, selection and report of the information is impossible. Selection, then would appear to occur later than area V1.

As our second piece of evidence for a cortical locus of attentional resolution we turn to the asymmetries found in the current experiments. Since we assumed a constant region of selection at all eccentricities, the inhomogeneities in attentional resolution then become a consequence of the mapping of visual space onto the cortex. We can then search for a particular visual cortex with the appropriate topology. Area V1, for example, does not show the advantages for lower field over upper field or for tangential over radial directions that we found in Experiment 2. On the contrary, many studies have shown that approximately equal areas of early visual cortex (V1 and V2) are devoted to processing the upper and lower visual fields (DeYoe et al., 1996; Horton & Hoyt, 1991; Sereno et al., 1995), suggesting that the site of attentional analysis lies beyond these cortices.

In fact, the first cortical region where we find an overrepresentation of the lower visual field is the occipital-parietal regions (Maunsell & Newsome, 1987). Furthermore, these regions are often linked to spatial attentional control (Driver & Mattingly, 1998; Gazzaniga, 1987; Gazzaniga & Ladavas, 1987; Posner, Inhoff, Friedrich, & Cohen, 1987; Posner, Walker, Friedrich, & Rafal, 1987; Posner, Walker, Friedrich, & Rafal, 1984; Steinmetz & Constantinidis, 1995). Van Essen, Newsome, and Maunsell (1984) have argued that the overrepresentation of the lower field is required

for the control of hand and arm movements. Attention-demanding, skilled actions involving the hands are almost always constrained to the lower field and parietal cortex is the logical site to coordinate vision and motor control. The lower field advantage in our two experiments suggests that the parietal cortex is the likely site of control of selection.

Converging evidence for this parietal locus of attentional selection comes from two sources. First, recent neuroimaging research (Culham et al., 1998) using an attentive tracking task quite similar to our task in Experiment 1 shows that, while early visual areas are not preferentially active during performance of this task, parietal areas do show increases in activation. Additionally, patients who have lesions to the parietal areas are completely unable to perform attentive tracking or to enumerate more than two or three items (François Michel and Marie-Anne Hénaff, personal communication, 1997).

Parietal Pointers

How would selection operate in the parietal cortex? Perhaps the selected information is simply transferred ahead from area V1 to a representation in the parietal cortex where it is analyzed in greater detail. Although others have suggested that portions of the visual input are copied to higher levels (e.g., Kosslyn & Koenig, 1995; Olshausen, Anderson, & Van Essen, 1995), these proposals have typically mentioned the inferotemporal cortex as the likely site. However, as Pylyshyn (e.g., Pylyshyn, 1989; Pylyshyn & Storm, 1988) noted, there is no need to duplicate the selected information to another region if it is possible simply to point to the area of interest. A pointer or index could then indicate the location of the selected items, say, in area V1. When required for further analysis, the properties of the item at that location could be read out via the pointer. In this case, there is no need for a higher cortical area to represent the full range of orientations, sizes, and colors that we might experience from the attended object.

The idea of using a set of pointers to index objects or regions of interest in early visual areas is similar to Pylyshyn's (1989, 2000) "fingers of instantiation," which pointed to the objects of interest. However, whereas Pylyshyn assumed only a list of pointers, we are proposing a two-dimensional map for the set of pointers. We assume that this spatial layout comes from the crude retinotopy available in parts of the parietal cortex (Blatt, Andersen, & Stoner, 1990). A pattern of activity in this "selection" area of the parietal cortex is then a salience map of the current visual input. Similar salience maps have been proposed by Koch and Ullman (1985) in the LGN; by Crick (1984) in the thalamic reticular nucleus; by Robinson and Petersen (1992) in the pulvinar; by Goldberg and colleagues (e.g., Duhamel, Colby, & Goldberg, 1992; Goldberg, Colby, & Duhamel, 1990; Gottlieb, Kusunoki, & Goldberg, 1998; Wurtz, Goldberg, & Robinson, 1980) in the LIP region of the parietal cortex; by Lu and Sperling (Lu & Sperling, 1995) in their model of high-level motion perception; and by Triesman in her feature integration

theory (Treisman & Gelade, 1980), which proposed a master map of spatial locations.

The spatial layout seems to be a necessary aspect of the selection mechanism. Recall Yantis's (1992) findings that spatial grouping factors improved the tracking of multiple targets. A spatial layout offers the simplest mechanism to support grouping of multiple selection locations based on spatial regularities. "Grouping," like "chunking" (Miller, 1956), is an umbrella term for the reduction in the complexity of the stimulus description that becomes possible when a set of items has some redundancy, similarity, or familiarity. The suggestion is that this reduction in perceptual load would support the improved tracking performance in the conditions of Yantis's study (1992) where spatial grouping was available in the display. For example, it would be easier to maintain tracking of multiple locations when the locations themselves could be derived from a simpler description such as the corners of a rigid square or a nonrigid but convex polygon.

Because of the large receptive fields in the parietal area (e.g., Blatt et al., 1990), any single unit in this parietal salience map could only indicate a selected location with crude precision. If two items were to fall within the indicated area further individuation would not be possible. Because of this lack of individuation, visual areas further along the stream of processing would not be able to determine which properties went with which item within the region. In fact, higher levels would treat all properties from the selected region as belonging to one entity.

Conclusions

In conclusion, we have shown that selection has a coarse grain, much coarser than visual resolution and we claim that this grain is a basic and unchangeable aspect of the access of attention to spatial locations, independent of the capacity limit of attention. The inhomogeneities of the spatial limits to selection—the drop off with eccentricity, the lower field advantage, and the advantage of tangential over radial spacing—are important and idiosyncratic guideposts in the effort to find the neural structures responsible for selection by visual attention. We suggest, as have many others, that the parietal area is the seat of this selection mechanism and that it acts by pointing to the spatial coordinates (or cortical coordinates) of items of interest rather than by holding a representation of the items themselves.

REFERENCES

- Andriessen, J. J., & Bouma, H. (1976). Eccentric vision: Adverse interactions between line segments. *Vision Research*, **16**, 71–78.
- Anstis, S. (1974). A chart demonstrating variations in acuity with retinal position. *Vision Research*, **14**, 589–592.

- Aubert, H., & Förster, R. (1857). Beiträge zur Kenntniss des indirekten Sehens. *Graefe's Archives of Ophthalmology*, **3**, 1–37.
- Bahcall, D. O., & Kowler, E. (1999). Attentional interference at small spatial separations. *Vision Research*, **39**, 71–86.
- Banks, W. P., & Prinzmetal, W. (1976). Configurational effects in visual information processing. *Perception & Psychophysics*, **19**, 361–367.
- Banks, W. P., Larson, D. W., & Prinzmetal, W. (1979). Asymmetry of visual interference. *Perception & Psychophysics*, **23**, 447–456.
- Baylis, G., & Driver, J. (1992). Visual parsing and response competition: The effect of grouping factors. *Perception and Psychophysics*, **51**, 145–162.
- Behrmann, M., & Moscovitch, M. (1994). Object-centered neglect in patients with unilateral neglect: Effects of left-right coordinates of objects. *Journal of Cognitive Neuroscience*, **6**, 1–16.
- Biederman-Thorson, M., Thorson, J., & Lange, G. D. (1971). Apparent movement due to closely spaced sequentially flashed dots in the human peripheral field of vision. *Vision Research*, **11**, 889–903.
- Blatt, G. J., Andersen, R. A., & Stoner, G. R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *Journal of Comparative Neurology*, **299**, 421–445.
- Bouma, H. (1970). Interaction effects in parafoveal letter recognition. *Nature*, **226**, 177–178.
- Bouma, H. (1973). Visual interference in the parafoveal recognition of initial and final letters of words. *Vision Research*, **13**, 767–782.
- Broadbent, D. (1982). Task combination and selective intake of information. *Acta Psychologica*, **50**, 253–290.
- Burr, D. C., & Ross, J. (1982). Contrast sensitivity at high velocities. *Vision Research*, **22**, 479–484.
- Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, **197**, 551–566.
- Cave, K. R., & Bichot, N. P. (1999). Visuospatial attention: Beyond a spotlight model. *Psychonomic Bulletin & Review*, **6**, 204–223.
- Chambers, L., & Wolford, G. (1983). Lateral masking vertically and horizontally. *Bulletin of the Psychonomic Society*, **21**, 459–461.
- Comtois, R. (1999). VisionShell: A Macintosh Programming Environment. In distribution by MicroML. St-Hyacinthe, Canada. (E-mail: ml@mblink.net)
- Crick, F. (1984). The function of the thalamic reticular complex: The searchlight hypothesis. *Proceedings of the National Academy of Sciences of the USA*, **81**, 4586–4590.
- Culham, J. C., Brandt, S. A., Cavanagh, P., Kanwisher, N. G., Dale, A. M., & Tootell, R. B. (1998). Cortical fMRI activation produced by attentive tracking of moving targets. *Journal of Neurophysiology*, **80**, 2657–70.
- DeYoe, E. A., Carman, G. J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., & Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. *Proceedings of the National Academy of Sciences of the USA*, **93**, 2382–2386.
- Downing, C. J. (1988). Expectancy and visual-spatial attention: Effects on perceptual quality. *Journal of Experimental Psychology: Human Perception & Performance*, **14**, 188–202.
- Downing, C. J., & Pinker, S. (1985). The spatial structure of visual attention. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 171–187). London: Erlbaum.
- Driver, J., & Baylis, G. C. (1989). Movement and visual attention: The spotlight metaphor

- breaks down. *Journal of Experimental Psychology: Human Perception and Performance*, **15**, 448–456.
- Driver, J., & Halligan, P. (1991). Can visual neglect operate in object-centered coordinates? *Cognitive Neuropsychology*, **8**, 475–496.
- Driver, J., & Mattingley, J. B. (1998). Parietal neglect and visual awareness. *Nature Neuroscience*, **1**, 17–22.
- Duhamel, J. R., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, **255**, 90–92.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, **113**, 501–517.
- Duncan, J. (1993). Similarity between concurrent visual discriminations: Dimensions and objects. *Perception & Psychophysics*, **54**, 425–430.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, **16**, 143–149.
- Eriksen, C. W., & Hoffman, J. E. (1973). The extent of processing of noise elements during selective encoding from visual displays. *Perception & Psychophysics*, **14**, 155–160.
- Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, **40**, 225–240.
- Gatti, S. W., & Egeth, H. (1978). A failure of spatial selectivity in vision. *Bulletin of the Psychonomic Society*, **11**, 181–184.
- Gazzaniga, M. S. (1987). Perceptual and attentional processes following callosal section in humans. *Neuropsychologia*, **25**, 119–133.
- Gazzaniga, M. S., & Ladavas, E. (1987). Disturbances in spatial attention following lesion or disconnection of the right parietal lobe. In M. Jeannerod (Eds.), *Neurophysiological and neuropsychological aspects of spatial neglect* (pp. 203–213). Amsterdam, The Netherlands: Elsevier.
- Goldberg, M., Colby, C., & Duhamel, J. (1990). Representation of visuomotor space in the parietal lobe of the monkey. *Cold Spring Harbor Symposia on Quantitative Biology*, **55**, 729–739.
- Goolkasian, P. (1999). Retinal location and its effect on the spatial distribution of visual attention. *American Journal of Psychology*, **112**, 187–214.
- Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, **391**, 481.
- He, S., Cavanagh, P., & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, **383**, 334–337.
- He, S., Smallman, H. S., & MacLeod, D. I. A. (1995). Neural and cortical limits on visual resolution. *Investigative Ophthalmology and Visual Sciences*, **36**, (Suppl.), 2010.
- He, Z. J., & Nakayama, K. (1994). Perceived surface shape not features determines correspondence strength in apparent motion. *Vision Research*, **34**, 2125–2135.
- Horton, J. C., & Hoyt, W. F. (1991). The representation of the visual field in human striate cortex. A revision of the classic Holmes map. *Archives of Ophthalmology*, **109**, 816–824.
- Hubel, D., & Wiesel, T. (1974). Uniformity of monkey striate cortex: a parallel relationship between field size, scatter and magnification factor. *Journal of Comparative Neurology*, **158**, 295–306.
- Hughes, H. C., & Zimba, L. D. (1987). Natural boundaries for the spatial spread of directed visual attention. *Neuropsychologia*, **25**, 5–18.
- Humphreys, G., & Riddoch, M. (1993). Interactions between object and space vision revealed

- through neuropsychology. In D. Meyer & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 139–168). Hillsdale, NJ: Erlbaum.
- James, W. (1890). *The principles of psychology*. New York: Holt.
- Johnston, A. (1986). A spatial property of the retino-cortical mapping. *Spatial Vision*, **1**, 319–332.
- Johnston, A. (1987). Spatial scaling of central and peripheral contrast-sensitivity functions. *Journal of the Optical Society of America*, **4**, 1583–1593.
- Johnston, A., & Wright, M. J. (1986). Matching velocity in central and peripheral vision. *Vision Research*, **26**, 1099–1109.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, **24**, 175–219.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, **4**, 219–227.
- Kooi, F. L., Toet, A., Tripathy, S. P., & Levi, D. M. (1994). The effect of similarity and duration on spatial interaction in peripheral vision. *Spatial Vision*, **8**, 255–279.
- Kosslyn, S. M., & Koenig, O. (1995). *Wet mind: The new cognitive neuroscience*. New York: Free Press.
- Kowler, E., & Steinman, R. M. (1977). The role of small saccades in counting. *Vision Research*, **17**, 141–146.
- Kowler, E., & Steinman, R. M. (1979). Miniature saccades: Eye movements that do not count. *Vision Research*, **19**, 105–108.
- Kramer, A. F., & Jacobson, A. (1991). Perceptual organization and focused attention: The role of objects and proximity in visual processing. *Perception & Psychophysics*, **50**, 267–284.
- Landolt, E. (1891). Nouvelles recherches sur la physiologie des mouvements des yeux. *Archives d'ophtalmologie*, **11**, 385–395.
- Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology Human Perception and Performance*, **21**, 451–468.
- Lavie, N., & Driver, J. (1996). On the spatial extent of attention in object-based visual selection. *Perception & Psychophysics*, **58**, 1238–1251.
- Levi, D., Klein, S. A., & Aitsebaomo, A. P. (1985). Vernier acuity, crowding, and cortical magnification. *Vision Research*, **25**, 963–977.
- Levi, D. M., Klein, S. A., & Yap, Y. L. (1987). Positional uncertainty in peripheral and amblyopic vision. *Vision Research*, **27**, 581–597.
- Lu, Z., & Sperling, G. (1995). Attention-generated apparent motion. *Nature*, **377**, 237–239.
- Maunsell, J. H., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, **10**, 363–401.
- Miller, G. A. (1956). The magical number seven, plus or minus two. *Psychological Review*, **63**, 81–97.
- Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision Research*, **29**, 1631–1647.
- Nakayama, K., He, Z. J., & Shimojo, S. (1995). Visual surface representation: a critical link between lower-level and higher level vision. In S. M. Kosslyn (Ed.), *An invitation to cognitive science* (pp. 1–70). Cambridge, MA: MIT Press.
- Olshausen, B. A., Anderson, C. H., & Van Essen, D. C. (1995). A multiscale dynamic routing circuit for forming size- and position-invariant object representations. *Computational Neuroscience*, **2**, 45–62.
- Posner, M. I. (1978). *Chronometric explorations of the mind*. Hillsdale, NJ: Erlbaum.

- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, **32**, 3–25.
- Posner, M. I., Inhoff, A. W., Friedrich, F. J., & Cohen, A. (1987). Isolating attentional systems: A cognitive-anatomical analysis. *Psychobiology*, **15**, 107–121.
- Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, **13**, 25–42.
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, **109**, 160–174.
- Posner, M. I., Walker, J. A., Friedrich, F. J., & Rafal, R. D. (1984). Effects of parietal injury on covert orienting of attention. *Journal of Neuroscience*, **4**, 1863–1874.
- Posner, M. I., Walker, J. A., Friedrich, F. A., & Rafal, R. D. (1987). How do the parietal lobes direct covert attention? *Neuropsychologia*, **25**, 135–145.
- Polyshyn, Z. W. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial indexing model. *Cognition*, **32**, 65–97.
- Polyshyn, Z. W. (2000). Situating vision in the world. *Trends in Cognitive Science*, **4**, 197–207.
- Polyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, **3**, 179–197.
- Robinson, R. G., & Peterson, S. E. (1992). The pulvinar and visual salience. *Trends in Neurosciences*, **15**, 127–132.
- Rovamo, J., Virsu, V., Laurinen, P., & Hyvarinen, L. (1982). Resolution of gratings oriented along and across meridians in peripheral vision. *Investigative Ophthalmology and Visual Science*, **23**, 666–670.
- Sagi, D., & Julesz, B. (1986). Enhanced detection in the aperture of focal attention during simple discrimination tasks. *Nature*, **321**, 693–5.
- Scholl, B. J., & Polyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, **38**, 259–290.
- Schwartz, E. L. (1980). Computational anatomy and functional architecture of striate cortex: A spatial mapping approach to perceptual coding. *Vision Research*, **20**, 645–669.
- Sears, C. R., & Polyshyn, Z. W. (2000). Multiple object tracking and attentional processing. *Canadian Journal of Experimental Psychology*, **54**, 1–14.
- Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., & Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, **268**, 889–893.
- Shulman, G. L., Wilson, J., & Sheehy, J. B. (1985). Spatial determinants of the distribution of attention. *Perception & Psychophysics*, **37**, 59–65.
- Steinman, B. A., Steinman, S. B., & Lehmkuhle, S. (1995). Visual attention mechanisms show a center-surround organization. *Vision Research*, **35**, 1859–1869.
- Steinmetz, M. A., & Constantinidis, C. (1995). Neurophysiological evidence for a role of the posterior parietal cortex in redirecting visual attention. *Cerebral Cortex*, **5**, 448–456.
- Toet, A., & Levi, D. M. (1992). The two-dimensional shape of spatial interaction zones in the parafovea. *Vision Research*, **32**, 1349–1357.
- Townsend, J. T., Taylor, S. G., & Brown, D. R. (1971). Lateral masking for letters with unlimited viewing time. *Perception & Psychophysics*, **10**, 375–378.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, **12**, 97–136.
- Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, **14**, 107–141.

- Tsal, Y. (1983). Movement of attention across the visual field. *Journal of Experimental Psychology: Human Perception & Performance*, **9**, 523–530.
- Tsal, Y., & Lavie, N. (1988). Attending to color and shape: The special role of location in selective visual processing. *Perception & Psychophysics*, **44**, 15–21.
- Van Essen, D. C., Newsome, W. T., & Maunsell, H. R. (1984). The visual field representation in striate cortex of the macaque monkey: Asymmetries, anisotropies, and individual variability. *Vision Research*, **24**, 429–448.
- Virsu, V., Naesaenen, R., & Osmoviita, K. (1987). Cortical magnification and peripheral vision. *Journal of the Optical Society of America*, **4**, 1568–1578.
- Virsu, V., & Rovamo, J. (1979). Visual resolution, contrast sensitivity, and the cortical magnification factor. *Experimental Brain Research*, **37**, 475–494.
- Walley, R. E., & Weiden, T. D. (1973). Lateral inhibition and cognitive masking: A neuropsychological theory of attention. *Psychological Review*, **80**, 284–302.
- Wandell, B. A. (1995). *Foundations of vision*. Sunderland, MA: Sinauer.
- Warren, H. (1897). The reaction time of counting. *Psychological Review*, **4**, 569–591.
- Wertheim, T. (1887). Über die Zahl der Seheinheiten im mittleren Teile der Netzhaut. *Graefes Archives of Ophthalmology*, **33**, 137–150.
- Westheimer, G., Shimamura, K., & McKee, S. P. (1976). Interference with line-orientation sensitivity. *Journal of the Optical Society of America*, **66**, 332–338.
- Wolford, G. (1975). Perturbation model for letter identification. *Psychological Review*, **82**, 184–199.
- Wolford, G., & Chambers, L. (1983). Lateral masking as a function of spacing. *Perception & Psychophysics*, **33**, 129–138.
- Wolford, G., & Chambers, L. (1984). Contour interaction as a function of retinal eccentricity. *Perception & Psychophysics*, **36**, 457–60.
- Wurtz, R. H., Goldberg, M. E., & Robinson, D. L. (1980). Behavioral modulation of visual responses by monkeys: Stimulus selection for attention and movement. In J. M. Sprague & A. N. Epstein (Eds.), *Progress in psychobiology and physiological psychology* (pp. 48–83). New York: Academic Press.
- Yantis, S. (1992). Multielement visual tracking: Attention and perceptual organization. *Cognitive Psychology*, **24**, 295–340.
- Yap, Y. L., Levi, D. M., & Klein, S. A. (1989). Peripheral positional acuity: Retinal and cortical constraints on 2-dot separation discrimination under photopic and scotopic conditions. *Vision Research*, **29**, 789–802.

(Accepted January 20, 2001; published online May 15, 2001)