3-D constraints on spatially parallel shape perception

GLYN W. HUMPHREYS

University of Birmingham, Birmingham, England

and

NICK DONNELLY

University of Kent, Canterbury, England

We report evidence from three sets of experiments dealing with spatially parallel grouping of parts in single objects. A first set of experiments demonstrates that parts can be encoded in a spatially parallel manner in three-dimensional (3-D) objects, while there is a serial selection of parts across objects. A second set of experiments further shows that grouping in 3-D is less affected by eliminating collinearity between the parts of objects than grouping in two dimensions, suggesting that 3-D constraints operate directly on visual grouping. A final pair of experiments demonstrates that rotating the elements in the plane, to make a physically unstable 3-D object, disrupts the benefit found with 3-D stimuli when collinearity is eliminated. The evidence indicates that there is rapid and spatially parallel encoding of 3-D object descriptions in vision.

Visual Search and Feature Binding

Object and shape recognition in humans is impressively rapid. With masked exposures of only 100 msec, humans can identify objects drawn unpredictably from a wide range of categories (Thorpe & Imbert, 1989), and with rapid serial visual presentations of 16 or 32 objects each for 80 msec or so, detection of a predesigned target can be nearly perfect (Potter, 1976). Yet, in contrast, in many apparently simple search tasks, for instance for a red O target among red X and green O distractors, target detection is relatively slow and shows linear increases with the number of distractors present. The slopes on the search functions in such tasks can also be of the order of 25–40 msec over items (e.g., Quinlan & Humphreys, 1987; Treisman, 1988; Treisman & Gelade, 1980). The influential feature integration theory (FIT) of attention proposes that serial attentional processes are required to bind together primitive visual features computed in parallel across the visual field (Treisman, 1998); hence search is slow and linearly related to the number of distractors when the target is defined by a conjunction of primitive features. This serial binding process is required not only across feature dimensions (e.g., for color and form) but also within dimensions. Consequently, linear search functions are found for targets defined by conjunctions of form fea-

tures (e.g., + vs. † and ‡) as well as for color × form targets (Duncan & Humphreys, 1989; Humphreys, Quinlan, & Riddoch, 1989). According to FIT, spatially serial binding of features may be avoided only when we perceive a known object, when stored knowledge may be used to “glue” features together (Treisman, 1988, 1998; Treisman & Gelade, 1980). Yet the available data suggest that stored knowledge must be recruited extremely rapidly on-line, given our ability to identify within 100 msec an unpredictable set of objects, each with multiple features (Thorpe & Imbert, 1989). Similarly, neurophysiological studies demonstrate rapid activation of cells in high-level areas, such as the infero-temporal lobe, by complex stimuli (e.g., faces; see Tovee & Rolls, 1995); maximal information may be coded within the first 50 msec of stimulus presentation by such cells. It is unclear how high-level representations could be recruited so rapidly were there not some binding of at least the primitive form features present. In the present study we assess whether there is spatially parallel binding of form information in vision, and whether such binding is influenced by high-level constraints concerning 3-D relationships between image features.

In fact more recent research on visual search does suggest that, within the form domain, primitive features are bound using spatially parallel processes. However, the efficiency of search for form conjunctions depends on grouping between distractors and between targets and distractors. Thus, with homogeneous distractors, search functions for targets defined by a conjunction of form features can be relatively flat, consistent with spatially parallel detection. Homogeneous distractors can group on the basis of shape identity, enabling them to be segmented from the

This work was supported by grants from the Biology and Biotechnology Research Council and from the Medical Research Council (U.K.) to G.W.H. We thank Kyle Cave for suggesting Experiment 3. Correspondence should be addressed to G.W. Humphreys, Behavioural Brain Sciences, School of Psychology, University of Birmingham, Edgbaston, Birmingham B15 2TT, England (e-mail: g.w.humphreys@bham.ac.uk).

target. In contrast, with heterogeneous distractors, which do not group, search for form conjunctions is slow and linearly related to the number of distractors present (Duncan & Humphreys, 1989; Humphreys et al., 1989). Such data are consistent with spatially parallel binding of features into groups; search efficiency varies according to whether the groups are consistent or inconsistent with the segmentation of targets from distractors.

Form Grouping and Visual Search

In the search studies of Duncan and Humphreys (1989) and Humphreys et al. (1989), grouping was based on identity between two-dimensional (2-D) shapes. Parallel grouping of form elements does not depend only on shape identity, however. In a study closely related to the present one, Donnelly, Humphreys, and Riddoch (1991) linked visual search procedures (varying the numbers of elements present) with a shape discrimination task (see also Baylis & Driver, 1994; Donnelly, Weekes, Humphreys, & Albon, 1998). In one condition, participants were required to detect a target vertex that pointed toward a central fixation point, while distractors were vertices that pointed away from fixation. On trials when the target was absent in this condition, the distractors could group on the basis of closure and collinearity between their end segments to form simple shapes—a square, a pentagon, and a hexagon—when four, five, or six distractors were present (see Figure 1a). Note that the distractors in this study were locally nonidentical but could group on the basis of relationship present across the whole display. Note also that targets, when present, could have exactly the same local features as distractors, so that their definition was based on a conjunction of shape and position relative to fixation. Despite these constraints, search performance was minimally affected by the number of vertices present, indicating that the vertices were bound together, to form a primitive shape description, on the basis of spatially parallel processes. Interestingly, when the orientations of the target and distractor vertices were reversed, so that targets pointed away from distractors and fixation, search was relatively difficult and linearly affected by the number of vertices present (see Figure 1b for examples of these displays). In this last condition targets and distractors had the same local orientations as in the first condition, so the parallel search functions found with closed shapes cannot be attributed to this factor. Instead, it appears that spatially parallel binding operated using display properties such as collinearity and closure (see also Donnelly et al., 1998, Elder & Zucker, 1993, and Gilchrist, Humphreys, Riddoch, & Neumann, 1997, for further evidence on the spatially parallel computation of closure and collinearity, respectively). Nor is it the case that grouping between form elements only gives rise to 2-D descriptions of displays. Enns and Rensink (1990, 1991) used displays with boxlike line drawings that could be perceived as having an orientation in 3-D space. They had participants search for targets that had the same 2-D line orientations as distractors, but targets and

distractors differed in their orientations in 3-D. Search was relatively efficient, with the slope of the functions being 10 msec/item or less. Such efficient search functions are unlikely to be realized by serial processes in the brain (Crick, 1984).

Humphreys, Keulers, and Donnelly (1994) had participants search for a target cylinder that differed in size from distractor cylinders. Targets and distractors were presented against backgrounds with strong linear perspective cues that could generate size illusions so that either targets or distractors were perceived as larger (or smaller) than was really the case. When the size illusion was consistent with the discrimination between targets and distractors (making the target appear larger and distractors smaller when search was for a large target cylinder), performance was facilitated; when the illusion was inconsistent with the discrimination, performance was slowed (measured relative to baseline conditions in which no size illusion was present). In this case, the perceived size of the stimuli in 3-D space affected search performance (see also Rensink & Enns, 1995, for other evidence that early grouping between form elements can disrupt visual search for apparently more primitive components of the group).

However, in the studies of Enns and Rensink (1990, 1991) and Humphreys et al. (1994), it is not clear whether the 3-D descriptions of the displays emerged only following the completion of grouping between 2-D components (e.g., from grouping between particular 2-D junctions; Enns & Rensink, 1991), or whether 3-D descriptions interacted with the grouping taking place between the elements in 2-D. Two-dimensional and 3-D representations from line drawings may be derived serially in discrete stages; first there is grouping based on 2-D cues and then there is activation of 3-D object descriptions. Alternatively, an interactive effect of grouping between 2-D and 3-D cues could take place if 3-D representations were activated in cascade, prior to completion of 2-D grouping.

The Present Study

In the present study, we investigated in more detail the relations between spatially parallel grouping processes in 2-D and 3-D space, using an extension of the search procedures introduced by Donnelly et al. (1991).¹ Participants were presented with displays composed of vertex elements derived from the 2-D projection of a 3-D cube (Figure 2). In the full cube display, there were 7 vertices present, 4 composed of 3 line elements, and 3 composed of 2 line elements. When the vertices fell in the positions of the corners of the cube, the target was absent. A target, when present, fell (at random) in one of the three locations occupied by the 2-element vertices. The cubes themselves were randomly oriented to face either to the left or the right. With a right-oriented cube (Figure 2), targets fell randomly at the upper back left; the lower front left; or the lower back right corner locations (Figure 2 shows a lower front left target). With a left-oriented cube, targets fell randomly at the upper back right, the lower

Figure 1. Displays used by Donnelly et al. (1991). (a) Closed shapes. (b) Nonclosed shapes. Only 4- and 6-element displays are shown.

front right, or the lower back left corner locations. Targets could also take either of two forms; either they were a reflected version of 1 of the 2-element vertices (so that 1 element pointed away from the surface of the shape), or they corresponded to one of the 2-element corner junctions with an element missing ("reflected" and "missing" targets, respectively). The display size was manipulated by presenting either a single face (the front, top side), or two faces (front+side, top+side), or all three faces of the cube, so that 4, 6, or 7 vertices were present (in the 1-, 2-, and 3- [full cube] face conditions).² In some experiments

the 3-element vertices were replaced by 2-element vertices (in control conditions we show that this change was not important for generating the effect with the cube displays; Experiment 2C).

The reflected targets were composed of the same elements as distractors in the displays, so that formally, these targets were defined by a conjunction of their shape and their location (as in Donnelly et al., 1991).³ The missing targets were defined by the absence of an element in a distractor vertex. Visual search for the absence of a feature is typically rather difficult and may require serial rather than

Figure 2. Full cube displays used in Experiment 1A. Here only one target position and one cube orientation are shown. Also, with 4- and 6-element displays, all faces and face pairings from the cubes were presented equally and randomly.

parallel search (Treisman & Souther, 1985). Accordingly, search for each type of target should be relatively difficult if parallel search is contingent solely on the individual features comprising targets and distractors.⁴ However, if the vertices group to form a single shape description, and if this shape description is sufficient to support the target-distractor discrimination, then performance should be relatively efficient and unaffected by the number of vertices present.

The present experiments are organized into three main groups. In Experiment 1A we report basic data on search for cube displays containing 4, 6, or 7 vertices, and in Experiment 1B we report a control study designed to test performance when the same vertices were present but were not configured into a cube. The contrast between performance in these two studies indicates the effects of grouping between the vertices on search. In Experiment 1C we provide a further control, in which just 2-element vertices were used, and these vertices were arranged sometimes to be consistent with 2-D squares where the front square partially occluded the back square. Here we ask whether

performance suffers when the junctions group into two shapes (the occluded squares) relative to when they group into a single coherent shape description (with the cubes, but not with the occluded squares). In their Experiment 4, Donnelly et al. (1991) presented participants with either single shapes (squares, pentagons, hexagons) or two shapes, one either side of fixation. Targets (vertices pointing into the center of each shape) could appear in either of the two shapes, when both shapes were present. Donnelly et al. found that there was an increment in reaction times (RTs) on 2- relative to 1-shape trials. They proposed that, although the elements within each shape grouped in a spatially parallel manner, participants selected each shape serially, generating the overall RT increment when two shapes were present. This overall effect on RTs did not interact with the display size. If this serial selection of shapes also holds for the occluded squares here, there should be an RT increment for 6- or 7-vertex displays (when two shapes appear) relative to 4-vertex displays (which configure into a single shape; see also Baylis & Driver, 1993; Duncan, 1984).

The second group of experiments examined the effects of within-plane rotation of individual elements on search. Using 2-D shapes, Donnelly et al. (1991) found that rotating the vertices within the plane, so that line segments were no longer collinear, led to reliable increases in the effects of the number of vertices on performance (though performance remained more efficient than when the vertices were reflected so that they could no longer be grouped to form closed shape descriptions, as in Figure 1b). Thus the parallel computation of 2-D shape descriptions is disrupted by eliminating collinearity between edge segments along the boundaries of the shapes. It is possible, though, that activation of a 3-D object description is less finely tuned to 2-D collinearity; for example, activation of a 3-D representation would provide a further source of evidence to help group elements together, when grouping in 2-D is weak. If this holds, then any increased effects of the number of vertices present, with rotated vertices, should be larger when the vertices form 2-D shape descriptions than when they form 3-D descriptions of objects. Note that such a finding would also go beyond previous studies on 3-D coding in search tasks. In the previous studies, it is possible that the 3-D descriptions arose only following the completion of grouping operations in 2-D space (see above). However, if the activation of 3-D descriptions is not disrupted by a variable that disrupts grouping in 2-D space, we can propose that grouping is influenced by 3-D descriptions in an interactive manner (rather than 3-D descriptions being formed at a processing stage after 2-D grouping is complete).⁵ This was tested by examining the effects on search of rotating vertices within the plane, with displays that group into 3-D objects or 2-D shapes (the occluded squares) and also with displays that do not group (random vertices).

In a third set of experiments we tested the effects of whole-object rotation. Whole stimuli were rotated so that, in the case of the cubes, forms were balanced on a bottom apex. This is a physically unstable position in 3-D space, making it unlikely that we have stored 3-D representations for objects at this orientation. However, the same rotation should not affect grouping by local 2-D cues such as collinearity (providing these are unaffected by orientation anisotropies, though we demonstrate that this is not the case). We assessed whether object rotation influenced the ability of participants to group and select the parts of the stimuli together, to produce efficient performance in the search task. In particular, we examined whether the robustness of 3-D forms to local vertex rotations (to eliminate collinearity) remained when the 3-D forms themselves fell at an unstable orientation.

GENERAL METHOD

Apparatus and Stimuli

Stimulus presentation, timing, and data collection were controlled using an Apple Macintosh IIci and VScope software. There were two types of full cube displays, with the object either facing left or right, and the other displays (with 4 or 6 rather than 7 vertices, or with random vertices or occluded squares) were derived

from this model. In all cases, the vertices were centered on the same locations where vertices were centered in the full cube displays. In the 4-vertex displays, the stimuli covered only one face of the original cube, and it was randomly the front, top, or side face. In the 6-vertex displays the stimuli covered two faces, and comprised the front + top, the front + side, and the top + side faces randomly. The target, when present, could appear at one of three locations in each full cube display (one of six locations in general, across the two cube orientations), and at the equivalent locations in the 4- and 6-vertex displays. Thus, within each 4- or 6-vertex display the target's location was more predictable than in a 7-vertex display (in each 4-vertex display, the target could appear only in one location; in each 6-vertex display, the target could be in one of two locations), although across displays, the target positions were matched across the 4-, 6-, and 7-vertex displays.

Individual line elements were 2 cm long and 1 mm across. The spacing between the elements in each corner of the full cube display was 2 cm, so that the full cube, when present, covered an area of about 6 cm vertically $\times$ 6 cm horizontally, and it was viewed from 50 cm. Fixation fell at the midpoint between the two 3-element vertices on the front surface of each cube (the front right vertices in the left-facing cube, and the front left vertices in the right-facing cube). The equivalent initial fixation position was maintained in all conditions. (Note that this meant that the initial fixation fell away from the shapes in some of the 4-vertex displays [e.g., for the top surface of the cube], but this had to be the case if the target and fixation positions were to be maintained across all display types.) A fixation cross was presented for 300 msec, followed by the displays. Display times were unlimited and participants were free to move their eyes after the start of the display.

Participants responded "present" or "absent" by pressing the "M" or "Z" key on the keyboard. The stimuli appeared as black lines against a white background. The fixation stimulus was a black cross.

Design and Procedure

Participants typically received all the conditions within an experiment. The display size was blocked (4-, 6-, or 7-vertex displays) to maximize the likelihood that participants would use stored templates for the grouped displays (if they could). There were equal numbers of target present and target absent trials. When different targets could occur across trials (reflected or missing), each appeared at random on half the target present trials. All the trials were administered during a single experimental session, with short rests between each display size block. Participants also received 20 practice trials before each block with the appropriate stimuli for the block. The order of the blocks was rotated across participants. Half the participants responded "present" with their preferred hand and half responded "absent." There were separate sets of participants in each experiment.

EXPERIMENT 1

Effects With 3-D and 2-D Stimuli

Experiment 1 provides the basic test of search for form conjunction (reflected) targets or for targets defined by the absence of a distractor element (missing targets). In Experiment 1A, full cube stimuli were presented in the 7-vertex condition, and the two and single faces of the cube were presented in the 6- and 4-vertex conditions. In Experiment 1B, the same vertices were presented as in Experiment 1A, but the individual stimuli were randomly assigned to the locations of the vertices in the cube displays (Figure 3). Only missing targets were used. For any single participant, the assignment of distractor vertices to their locations was held constant across the

Figure 3. Random vertex displays used in Experiment 1B.

trials, to match the constant assignment of vertices in Experiment 1A. In Experiment 1C, 2-element vertices were used to create either "occluded squares" or "partial cubes," the difference between these stimuli being the presence of L junctions at the lower back right or left position (depending on the orientation of the original cube) and at the upper back right or left position (Figures 4 and 5). In occluded squares, two L junctions projected at 90° onto two of the virtual edges of the front square in the display (Figure 4). These orthogonal projections create T junctions with the virtual edge. T junctions provide a strong cue for the segmentation of objects into separate parts (Enns & Rensink, 1991; Perkins, 1968), and the T junction created here generated the perception of two 2-D squares, one occluding the other, rather than the perception of a single, coherent shape (see Experiment 1C for converging evidence). In the partial cube displays, the vertices were matched to those used in the occluded square displays (e.g., the 3-element vertices in the standard cube displays were replaced with 2-element vertices), except that (1) the lower back right or left vertex and (2) the upper back left or right vertex were no longer Ls but in-

stead each created an internal angle of 135° and projected respectively to (1) the front lower right or left vertex or (2) the front upper left or right vertex (Figure 5). In such displays, continuation from the lower back right (or left) vertex to the lower front right (or left) vertex and from the upper back left (or right) vertex to the upper front left (or right) vertex should enable the vertices to be encoded as a coherent (3-D) shape (again, see Experiment 1C for converging evidence). By comparing performance on occluded square and partial cube displays, we provide an additional test of whether the present cube displays are grouped to form a description of single perceptual objects. This can be measured by comparing RT increments when two objects are present (occluded squares) relative to one object (partial cubes).

Experiment 1A Basic Data With Cube Displays

Method

There were 10 participants, 6 female and 4 male, all of whom took part as members of a research participant scheme. All had either normal or corrected-to-normal vision.

Figure 4. Occluded square displays used in Experiment 1C.

There were three blocks, each of 120 trials, for the 4- (single face), 6- (two faces) and 7- (full cube) vertex displays, respectively. In each block there were 60 absent trials, 30 trials with a reflected target present, and 30 trials with a missing target present. For each display type, half the displays were derived from a right-facing cube and half from a left-facing cube. For 4-vertex displays, each single face from the cube was presented equally often (20 absent trials each, 10 for each type of present response); for 6-vertex displays, each pairing of two surfaces from the cubes was presented equally often (20 absent trials, 10 for each type of present trial). The two types of present response occurred randomly within each block of present trials. Stimuli were derived from the left- and right-facing cubes on half the trials each. Participants responded as quickly as they could without making errors.

A trial began with a central fixation cross for 500 msec, followed by the display. The display remained on until participants made a response.

Results

Mean correct RTs are given in Figure 6, and error rates in Table 1. RTs were analyzed in a 3 (condition: target present missing, target present reflected, target absent) $\times$ 3 (display size: 4, 6, and 7 vertices) repeated measures analysis of variance (ANOVA), averaging the data across the three target locations, the three types of

4- and 6-vertex displays, and the two display orientations (left and right facing). There was a reliable main effect of condition [$F(2,18) = 8.35, p < .01$]; the effect of display size was not reliable nor was the condition $\times$ display size interaction (both $F < 1.0$). Present reflected responses were faster than present missing and absent responses, which did not differ ($p < .01$, Newman-Keuls analyses).

The error rates were low and there was no sign of a speed-accuracy tradeoff. Errors were not analyzed further.

Discussion

The findings were straightforward. RTs were flat as a function of the increase in the display size from 4 to 7 vertices. Such flat search functions are consistent with performance being dependent on spatially parallel computations to discriminate targets (Crick, 1984; Treisman & Gelade, 1980). Reflected targets were somewhat easier to detect than missing targets, perhaps because the missing targets tended to be grouped into the shape representations, slowing their detection. Nevertheless, there remained a minimal effect of display size, indicating that serial processes were still not required for target detection.

Figure 5. Partial cube displays used in Experiment 1C.

The flat search functions we observed in Experiment 1A could have arisen for various reasons. It may be that no display size effects occurred because display elements grouped into a single object description, allowing them to be encoded together. Alternatively, however, either (1) the differences between targets and distractors were so salient that targets could be detected efficiently irrespective of the context provided by distractors, or (2) the predictability of the target's location led to efficient search. Experiment 1B sought to test these alternative accounts. In Experiment 1B, for each participant, distractors were randomly (but constantly across trials) assigned to the positions held by the vertices in the cube displays in Experiment 1A (for each display size). If the configural context provided by the distractors was unimportant, performance again ought to be efficient here and minimally affected by the display size.

Experiment 1B Random Displays

Method

There were 8 participants, 4 male and 4 female, all with normal or corrected-to-normal vision. All were paid £2 for participating.

There were 120 trials per participant, 60 absent and 60 target missing, 20 at each display size. Only target missing trials were used because it was not possible with random vertex displays to define targets separately from distractors. Display size was blocked, with the order of the blocks randomly determined for each participant.

Results

The mean correct RTs are given in Figure 7, and the error rates in Table 1.

There were reliable main effects of condition (present vs. absent) and display size (4, 6, or 7 vertices) [$F(1,7) = 9.09$, $p < .025$, and $F(2,14) = 14.18$, $p < .001$]. There was also a reliable interaction between condition and display size [$F(2,14) = 4.21$, $p < .05$].

For both present and absent responses, RTs increased linearly as a function of display size. For present responses, the slope of the RT display size was 15 msec/item, and the linear component accounted for 98% of the variance. For absent responses, the slope of the search function was 21 msec/item, and the linear component accounted for 99.9% of the variance. The slope for absent trials was steeper than for present.

There were again minimal errors and no evidence of a tradeoff with accuracy.

Figure 6. Mean correct reaction times (RTs, in milliseconds) in Experiment 1A, with full cubes.

Experiment 1A Versus Experiment 1B

To compare performance in the two subexperiments, a further ANOVA was conducted on the RT data with experiment as a between-subjects factor (including data from only the target missing and absent trials in Experiment 1A). There was an interaction between experiment and display size [$F(2,16) = 9.21, p < .01$], which was not qualified by an interaction with condition (present vs. absent). The effects of display size were reliably greater in Experiment 1B than in Experiment 1A.

Discussion

Unlike in Experiment 1A, there were now reliable effects of display size on performance and strong trends for linear search functions on both present and absent trials. The data here indicate that differences in target-distractor features, or the predictability of the target's location, were not sufficient to produce the flat search functions in Experiment 1A (as the experiments were matched

in this respect). Instead it appears that the configured context provided by distractors facilitates target encoding. Nevertheless, we should also point out that, given that participants here searched for the absence of a feature (for the missing target), search was relatively fast even with random displays. In previous studies examining search for the absence of a feature (e.g., Treisman & Souther, 1985), slopes of 25–30 msec/item have been observed. Also, the slope ratio for present:absent responses here was 1:1.4, rather than the 1:2 ratio expected from serial, self-terminating search (Treisman & Gelade, 1980). It is possible that this relative efficiency was due to the predictability of the target's location. Unlike in most visual search procedures, targets could appear in only one of three locations here (within a given display, or one of six locations when the different orientations of the stimuli across trials are taken into account). Reducing locational uncertainty should benefit search, though not so much as having distractors that provide a configural context (Experiment 1A).

Table 1
Error Rates in Experiment 1

Stimuli	Target Present Reflected: Display Size			Target Present Missing: Display Size			Target Absent: Display Size		
	4	6	7	4	6	7	4	6	7
Experiment 1A									
Full cubes	.02	.01	.02	.01	.03	.02	.00	.01	.01
Experiment 1B									
Random displays	.01	.03	.04	.01	.02	.02	.01	.01	.02
Experiment 1C									
Occluded squares	.01	.03	.03	.00	.02	.03	.01	.01	.01
Partial cubes	.02	.02	.02	.01	.00	.01	.01	.00	.01

Figure 7. Mean correct reaction times (RTs, in milliseconds) in Experiment 1B, with random cubes.

Experiment 1C was conducted to assess whether stimuli had to be assimilated into a single object description for efficient search to occur, or whether search would be equally efficient when items grouped into two object descriptions. Observers took part in two blocks of trials. In one they received occluded square stimuli (Figure 4); in the other they received partial cube stimuli (Figure 5). With the occluded square stimuli, it is possible that each square is encoded as a separate stimulus and selected independently. In that case, there should be an RT increment in the 6- and 7-vertex displays (when separate "front" and "back" shapes may be coded) relative to the 4-vertex displays (when only a single shape was likely to be encoded; see Donnelly et al., 1991, Experiment 4). The partial cube stimuli provide a control for the occluded squares, since they contain similar vertices but group (on the basis of continuation between the vertices) into a single object description.

Experiment 1C Occluded Squares Versus Partial Cubes

Method

There were 12 participants, 6 female and 6 male, who received course credit for taking part. All had either normal or corrected-to-normal vision.

There were two blocked conditions (with occluded squares and partial cubes), each with 108 trials (54 present, 54 absent; 18 trials at each display size). Unlike in the previous experiments, on 6-vertex displays, the front surface of the original cube (from Experiment 1A) was always present to increase the likelihood that stimuli were encoded into separate front and back areas, with the square being encoded as the front shape. The 4-vertex displays were based on the vertices present for each surface of the original cube. Display size

was again blocked (within a condition block), and the different displays occurred randomly within a display size block. The order of the condition blocks was balanced over participants, and the order of the display sizes was balanced across the condition block groups. Only missing targets were used, on present trials.

An independent group of 20 observers was also shown the full cube (7-vertex), partial cube, and occluded squares stimuli and asked to decide whether they thought that there were one or two objects present. All 20 reported one object with the partial cubes and two objects with the occluded squares.

Results

Figure 8a shows the mean correct RTs across participants for occluded squares, and Figure 8b shows the RTs for partial cubes. The mean error rates are given in Table 1.

RTs were analyzed in a three-factor repeated measures ANOVA, with the factors being display type (occluded squares vs. partial cubes), condition (present vs. absent), and display size (4, 6, or 7 vertices). There were reliable effects of display type [$F(1,11) = 8.76, p < .05$], condition [$F(1,11) = 16.83, p < .01$], and display size [$F(2,22) = 13.33, p < .01$]. These effects were qualified by two-way interactions between display type and condition [$F(1,11) = 7.54, p < .05$], display type and display size [$F(2,22) = 6.24, p < .05$], and condition and display size [$F(2,22) = 1.582, p < .001$]. The three-way interaction between display type, condition, and display size was also statistically significant [$F(2,22) = 11.36, p < .01$].

The three-way interaction was broken down by further ANOVAs on each display type. For the occluded squares (Figure 8a), there were reliable effects of condition and display size [$F(1,11) = 15.32, p < .01$, and $F(2,22) = 14.37, p < .01$, respectively]. There was also a condition

Figure 8. (a) Mean correct reaction times (RTs, in milliseconds) in Experiment 1C, with occluded squares. (b) Mean correct RTs in Experiment 1C, with partial cubes.

× display size interaction [$F(2,22) = 9.62, p < .01$]. RTs increased from 4- to 6- and 7-vertex displays ($p < .01$, Newman-Keuls test), though there was no further increase from 6- to 7-vertex displays. Absent responses were slower than present, and this effect increased with display size.

For the partial cubes (Figure 8b), the effects of condition and of display size were not significant; nor was the condition × display size interaction (all F s < 1.0).

Overall, the error rates were low and did not tradeoff with accuracy (Table 1).

Discussion

Performance differed for occluded squares and partial cube displays. With occluded squares, there was an RT increment when 6- and 7-vertex displays were presented relative to when 4-vertex displays occurred. This RT increment suggests that with 6- and 7-vertex displays, participants either processed the vertices less efficiently in parallel or they engaged a serial selection operation, selecting first one part and then another part of the display. Interestingly, there was no additional increment in RTs from 6- to 7-vertex displays, which is consistent with both 6- and 7-vertex displays being parsed into two separate surfaces (a front and a back shape) and with each surface being selected independently. The RT increment from 4- to 6- and 7-vertex displays matches that observed by Donnelly et al. (1991) with spatially separated 2-D shapes, who also found an RT increment when two rather than one shape was present. The present data suggest that similar selection processes operate on (1) spatially overlapping (occluding) 2-D shapes separated in depth and (2) 2-D shapes separated in the picture plane (see also Duncan, 1984, for a similar argument using a different paradigm). The serial selection of each shape in the occluded squares display will also lead to absent responses being slower than present responses (and to a larger RT increment on absent responses). On present trials, the target will be in the first shape selected half of the time, and in the second shape on the other half. On absent trials, responses can be made only after both shapes have been selected.

This account of serial selection of each square surface is also supported by an analysis of the standard deviations in RTs across participants. For the standard deviations there was also a three-way interaction between display type, condition, and display size [$F(2,22) = 9.48, p < .01$]. Standard deviations were increased for the occluded squares relative to the partial cubes, and this increased for Display Sizes 6 and 7 relative to Display Size 4. Interestingly, for occluded square displays, standard deviations were higher on target present than on target absent trials, and this increased for Display Sizes 6 and 7 relative to Display Size 4 [$F(2,22) = 5.73, p < .01$ for the display size and presence-absence interaction]. This increase in the standard deviations on present trials would be expected if targets are detected in the square selected first in half the trials and in the square selected second on the remaining trials. On absent trials, both squares should always be selected.

In contrast to the results with occluded squares, with partial cube displays, performance was minimally affected by the number of vertices present, and present and absent responses did not differ. These data resemble those in Experiment 1A with full cube stimuli. The findings suggest that the vertices in the partial cubes displays grouped to form a single shape description that enabled targets to be detected in a spatially parallel manner (with no increments from 4- to 6- and 7-vertex displays).

Summary of Experiment 1

Experiment 1 has shown that targets defined as parts of a coherent shape description can be detected in a spa-

tially parallel manner (Experiment 1A and with partial cubes in Experiment 1C). Control studies demonstrated that spatially parallel search was not simply due to feature differences between targets and distractors or to the predictability of the target locations in the present study (Experiment 1B). Also, there were RT increments (and increases in standard deviations of RTs) with larger display sizes when the vertices segmented into separate object descriptions (with occluded squares displays) relative to when they grouped into a single object description (Experiment 1C). Taken together, the results are consistent with there being spatially parallel binding of form information in vision. Also, this parallel binding can generate representations of 3-D objects (cubes, in the present instance), so that clear performance differences emerge when vertices configure into a 3-D object description relative to when they configure into independent 2-D shape representations separated in depth (Experiment 1C).

EXPERIMENT 2

3-D Constraints on Parallel Grouping

Experiment 1 showed that vertices can be grouped in parallel into 3-D object representations. This supports other recent results in the visual search literature (e.g., Enns & Rensink, 1991; Humphreys et al., 1994). However, in neither the present Experiment 1 nor in the previous studies is it clear whether the formation of 3-D object descriptions was contingent on the completion of grouping based on 2-D image cues, or whether 3-D shape descriptions can interact with and supplement 2-D grouping. To explain the existing data, an account in terms of 2-D grouping is sufficient: Performance is unaffected by the number of vertices present when vertices group in 2-D. Experiment 2 set out to evaluate whether 3-D shape descriptions directly influence grouping. To do this, we manipulated a variable (rotation of vertices within the plane so that they were no longer collinear) to disrupt grouping in 2-D. Previous data have shown that search performance with 2-D displays is affected by the number of vertices when vertices are noncollinear. If 3-D object representations are constructed after elements have grouped using 2-D information, the effects of rotating any vertices in the plane should be at least as large with 3-D objects (e.g., full and partial cubes) as with 2-D shapes (e.g., with occluded squares). On the basis of previous results, we expect that making the vertices noncollinear, by rotation, will lead to stronger effects of display size with 2-D shapes relative to when the vertices are collinear (Donnelly et al., 1991). Such increased effects of display size will provide a marker for grouping operating on 2-D representations.

We examined the effects of within-plane rotation of vertices using the four displays from Experiment 1: full cubes, partial cubes, occluded squares, and random vertices. In each case, vertices were rotated 15° within the plane. Performance with occluded squares should provide a measure of the effects of within-plane rotation of vertices on grouping in 2-D, which can be used to

Figure 9. Ocluded square displays with rotated vertices used in Experiment 2A.

compare the effects of the same rotations with the other displays.

Experiment 2A Occluded Squares

Method

There were 12 participants, 6 female and 6 male, all with normal or corrected-to-normal vision, and all of whom participated for course credits.

The displays were the same as for the occluded squares condition in Experiment 1C except that each vertex was rotated 15° within the plane around its corner position. This resulted in shapes that were partially closed but noncollinear, no longer having aligned end terminators. Example displays are shown in Figure 9. There were 108 trials, 54 present and 54 absent, at each display size. In the 6- and 7-vertex displays, half the displays faced right and half left.

Results

Figure 10 gives the mean correct RTs as a function of the display size; the error rates are shown in Table 2. There were reliable main effects of both condition (present vs. absent) and display size [$F(1,11) = 16.83, p < .01$, and $F(2,22) = 17.24, p < .001$]. There was also a reliable

condition $\times$ display size interaction [$F(2,22) = 9.42, p < .001$]. On present trials, RTs increased as a function of the display size, with this increase being reliable both from 4- to 6-vertex displays ($p < .01$, Newman-Keuls test) and from 6- to 7-vertex displays ($p < .05$, Newman-Keuls test). The RT increases from 4- to 6- and from 6- to 7-vertex displays were also reliable on absent trials (both $ps < .01$, Newman-Keuls test). The interaction arose because the effects of display size were larger on absent than on present trials. The slopes on the search functions were 18 and 14 msec/item on present and absent trials, and the linear components of the search functions accounted for 98% and 99% of the variance, respectively. Error rates were low and did not tradeoff with accuracy (Table 2).

Experiment 2A Versus Experiment 1C

To assess the effects of within-plane rotation on occluded square displays, performance in Experiment 2A was compared with that on the occluded square displays in Experiment 1C, with experiment as a between-subjects factor. Of primary importance is the effect of experiment and any interactions of experiment with the effects of

Figure 10. Mean correct reaction times (RTs, in milliseconds) in Experiment 2A, with occluded squares.

display size. The main effect of experiment was reliable [$F(1,22) = 8.24, p < .01$], as was the two-way experiment $\times$ display size interaction [$F(2,22) = 6.41, p < .05$] and the three-way interaction between experiment, condition (present/absent), and display size [$F(2,22) = 5.38, p < .05$]. The effects of display size were larger on the occluded squares in Experiment 2A than in Experiment 1C, and this was particularly true on absent trials. Nevertheless, just taking present trials alone, there remained an interaction between experiment and display size [$F(2,22) = 5.20, p < .05$].

Discussion

The results show that there were larger effects of display size (the number of vertices) on both present and absent trials with rotated vertices (in Experiment 2A) than when the vertices were collinear and presented at a canonical orientation (with horizontal and vertical components, in Experiment 1C). In addition, RTs were overall slower with rotated vertices. These effects of rotation match those reported by Donnelly et al. (1991, Experiment 3), using simple 2-D polygons. Donnelly et al. suggested that with weakened grouping by collinearity, either the elements were encoded less efficiently in parallel or spatially serial selection processes were used on some occasions to detect targets. In either case, effects of display size emerge. The present results are consistent with this argument. When collinearity between the vertices is eliminated, the data indicate that it was more difficult to group the vertices constituting each square, leading to increased effects of display size on performance. However, there is at least one other account of the increased display size effects in the study. This is that performance efficiency is determined by the junctions being presented

at a canonical orientation (with horizontal and vertical line components, as in Experiment 1C), in addition to any other effects of grouping between the junctions. Performance might have been disrupted here simply by our own use of vertices rotated to oblique orientations. We return to assess this last possibility in the control study, Experiment 2C, where we use randomly positioned vertices. In Experiment 2B, we assess the effects of rotating the vertices with stimuli from Experiment 1 that seem to be assigned a 3-D description: full and partial cubes.

Experiment 2B Full and Partial Cubes

Method

There were 12 participants, 7 female and 5 male, all with either normal or corrected-to-normal vision. Participation was for course credit.

The displays were equivalent to those used in Experiment 1A (full cubes) and the partial cube displays in Experiment 1C, except that the vertices were each rotated 15° within the plane around each respective corner position. The resulting shapes were again par-

Table 2
Error Rates in Experiment 2

	Target Present: Display Size			Target Absent: Display Size		
	4	6	7	4	6	7
Experiment 2A						
Occluded squares	.02	.04	.05	.02	.02	.03
Experiment 2B						
Full cubes	.01	.02	.02	.02	.01	.01
Partial cubes	.03	.02	.03	.01	.01	.01
Experiment 2C						
Y-Junction displays	.02	.03	.04	.02	.02	.04
L-Junction displays	.01	.03	.04	.01	.02	.02

Figure 11. Full cube displays with rotated vertices used in Experiment 2B.

tially closed, but lacked collinearity between the free ends of the vertices. Example displays are shown in Figures 11 and 12.

The full and partial cube displays were presented in separate trial blocks, as were the three display sizes within each stimulus block. Within each stimulus block there were 108 trials, 36 at each display size (4, 6, and 7 vertices), 18 of which were present and 18 of which were absent. Block order was counterbalanced across participants. Only missing targets were used, on present trials.

Results

The mean correct RTs are shown in Figure 13a for the full cubes, and in Figure 13b for the partial cubes. Error rates are given in Table 2.

The data were analyzed using a three-factor repeated measures ANOVA, with stimulus (full vs. partial cube), condition (present vs. absent), and display size (4, 6, 7) as the factors. There was a reliable main effect of stimulus [$F(1,11) = 10.13, p < .01$]; RTs were faster with the full than with the partial cubes. There were no effects of condition or display size (both F s < 1.0), and no reliable interactions.

Error rates were low and there were no signs of any speed-accuracy tradeoffs.

Full cubes: Experiment 2B Versus Experiment 1A

To assess the effects of rotating the vertices on the full cube displays, the relevant RT data from Experiment 2B (rotated vertices) were compared with those from Experiment 1A (canonical vertices), with experiment as a between-subjects factor. There was a reliable effect of experiment [$F(1,22) = 12.93, p < .01$]; RTs were overall slower with rotated than with canonical vertices. However, there were no other main effects and no interactions with the experiment factor.

Partial Cubes: Experiment 2B Versus Experiment 1C

A similar analysis was conducted on the data with partial cubes, with similar results. There was an overall main effect of experiment [$F(1,22) = 14.24, p < .01$], but no interactions between this factor and other factors in the study.

Discussion

Unlike with occluded square displays, where display size effects increased when vertices were rotated to be non-

Figure 12. Partial cube displays with rotated vertices used in Experiment 2B.

collinear, the effects of vertex rotation on cube displays were confined to an overall increase in RTs; there were no concomitant increases in the effects of display size on performance. The failure to find any effects of display size with rotated cube displays is important and suggests that the 3-D descriptions derived for these displays enabled them to be grouped in a manner that was relatively robust to the reduced 2-D collinearity between the vertices. For example, corner junctions may be grouped in a spatially parallel manner using assumptions about the structure of objects in 3-D that continue to operate even when there are “noisy” representations of 2-D grouping cues. Apparently, tolerance of 2-D noise enables the junctions in cube displays to still be grouped in a spatially parallel manner, even when they are rotated so as to be no longer collinear in 2-D.

Converging evidence for the rotated full and partial cube stimuli still being seen as 3-D objects comes from rating data from an independent set of 20 participants. These participants were asked to rate how “objectlike” the rotated versions of the occluded squares, partial cubes, and full cubes were on a 1–5 scale (5 = *objectlike*). The mean ratings were 1.6, 3.9, and 4.0 for the occluded

squares, partial cubes, and full cubes, respectively. The ratings differed across the stimuli [$F(2,19) = 9.24$, $p < .01$], with the occluded squares being rated as less objectlike than the other stimuli ($p < .01$, Newman–Keuls test, for both comparisons); the partial and full cubes did not differ. The partial and full cubes appeared more objectlike than the occluded squares even when their vertices were rotated to be noncollinear.

The present results suggest that 3-D object descriptions cannot be formulated only after the completion of 2-D grouping, since any costs in 2-D grouping (from vertex rotation) should then be passed on to performance with 3-D shapes. Instead of this, we found that 3-D shape representations continued to show no effects of display size on performance, even though effects were apparent when the elements grouped into 2-D squares (in the occluded squares condition). This is consistent with 3-D representations being activated in a cascade manner, prior to grouping based on cues being completed.

In Experiment 2C, we conducted a further control study in which we examined the effects on search performance of rotating the elements in random vertex displays.

Figure 13. (a) Mean correct reaction times (RTs, in milliseconds) in Experiment 2B, with full cubes. (b) Mean correct RTs in Experiment 2C, with partial cubes.

There were two forms of random vertex display. In Y-junction displays, the vertices were the same as in Experiment 1B, but each was rotated 15° in the plane (Figure 14). In L-junction displays, the vertices were the same as for the occluded squares displays in Experiment 2A, except that they were placed in random locations (different random locations across participants, but constant for a given participant; Figure 15). By means of this control study, we ask whether the shift from having

vertices at canonical to having them at noncanonical (rotated) orientations was the cause of the increased display size effects in Experiment 2A.

Experiment 2C Random Vertices

Method

Twelve participants took part, 5 female and 7 male, for course credit. All had either normal or corrected-to-normal vision.

Figure 14. Y-junction displays with rotated vertices used in Experiment 2C.

In the Y-junction displays, the vertices were the same as in Experiment 1B except that each was rotated 15° within the plane about its corner point. In the L-junction displays, the vertices were the same as in Experiment 2A except that each was assigned to a random location (centered at the location of each element in the occluded square displays). Targets were also presented at the same locations as those with the occluded square displays. Example displays are shown in Figures 14 and 15.

The two display types were presented in separate trial blocks, and the three display sizes were blocked within each display-type block. In each display block there were 108 trials, 54 present and 54 absent, with 18 at each display size. Block order was counterbalanced across participants. Only missing targets were used.

Results

The mean correct RTs are presented in Figure 16, and the error rates in Table 2.

A three-factor repeated measures ANOVA was conducted, with display type (Y vs. L junction), condition (present vs. absent), and display size (4, 6, 7) as the factors. There were reliable effects of condition [$F(1,1) = 18.25, p < .001$] and display size [$F(2,22) = 16.40, p < .001$], but no main effect of display type ($F < 1.0$). The

only interaction was between condition and display type [$F(2,22) = 13.89, p < .01$].

RTs increased with the display size for both present and absent responses, with the effects being larger on absent than on present trials. Averaging across display types, slopes for the present and absent responses were 20 and 31 msec/item, respectively, and the linear components accounted for 98% of the variance for each response type.

The errors showed the same pattern as the RT data and were not analyzed further.

Y Junctions: Experiment 2C Versus Experiment 1B

To assess the effects of within-plane rotation on the random vertex displays, performance with the Y-junction stimuli here was compared with performance on the equivalent canonical-vertex stimuli in Experiment 1B. Experiment was a between-subjects factor. There was a main effect of experiment [$F(1,22) = 8.51, p < .01$] and an interaction between experiment and display size [$F(2,22) = 6.44, p < .01$]. There were no other interactions.

Figure 15. L-junction displays with rotated vertices used in Experiment 2C.

Discussion

There were small but reliable increases in the effects of display size when randomly positioned vertices were rotated within the plane relative to when they were presented at canonical orientations (in Experiment 2C vs. Experiment 1B). It would appear that having horizontal and vertical line elements in displays facilitates search. This may be linked to the “oblique effect” in human vision, whereby the discrimination of lines at horizontal and vertical orientations is superior to that of oblique lines (see Appelle, 1972, and Essock, 1980, for reviews), perhaps reflecting the tuning of early oriented edge filters (Foster & Ward, 1991a, 1991b). In previous studies of search with multielement displays, search asymmetries have also been demonstrated, in which (for example) the detection of an oblique line among vertical lines is easier than the converse task (e.g., Treisman, 1985). At least in part, this may reflect facilitated grouping among vertical relative to oblique elements. With the present displays, facilitated grouping between canonical vertical and hor-

izontal lines would benefit search even with randomly positioned vertices.

The increased effects of display size with rotated junctions occurred here both when Y junctions were present and when they were absent (with L-junction displays; indeed, there was little difference between performance with the two types of display). This rules out one argument concerning the contrasting effects of rotation on cube (Experiment 2B) and occluded square displays—namely that the presence of Y junctions in the cubes led to small effects of display size even with rotated stimuli (e.g., perhaps because Y junctions are less disrupted by within-plane rotations). Here display size effects were increased by rotation irrespective of the presence of Y junctions (note also that, with the partial cube displays in Experiment 2B, there was no interaction between display size and rotation, though there were no Y junctions present in such displays).

Although there were increased effects of display size with rotated random vertices, the magnitude of the increase was small (e.g., an increase of around 5 msec/item

Figure 16. Mean correct (RTs, in milliseconds) with Y-junction and L-junction displays in Experiment 2C.

on present trials in Experiment 2C relative to Experiment 1B). Interestingly, the effects of rotation on the occluded square displays were somewhat larger. In particular, although there were no increases in RTs from 6- to 7-vertex displays in Experiment 1C (in fact there was a slight decrease), with rotated vertices (Experiment 2A), there was a reliable increase of 20 msec. This suggests

that rotation disrupted a particular process with occluded squares over and above any effects with random vertices (e.g., effects due to the introduction of oblique orientations in the displays). One likely possibility is that performance with occluded squares was particularly disrupted by the loss of collinearity between the end segments of the vertices.

Figure 17. Whole-figure rotated versions of the partial cubes used in Experiment 3A.

EXPERIMENT 3 Effects of Whole-Object Rotation

Experiment 2 demonstrated that 3-D shapes are more resistant to effects of vertex rotation than are 2-D shapes, where performance is measured in terms of search efficiency. It follows that reducing the saliency of 3-D information should reintroduce effects of 2-D noise on performance (when vertices are noncollinear). We attempted to do this in Experiment 3 by rotating the complete shapes used in Experiments 1 and 2. In Experiment 3 cube stimuli no longer had a flat base but were rotated to “perch” on the bottom vertex (Figure 17). This manipulation reduces the physical stability of the figure in 3-D, though the degree of collinearity in 2-D remains the same. It is unlikely that we have stored representations of 3-D shapes at orientations that are physically unstable. Under these conditions, the involvement of stored 3-D descriptions in search may decrease.

Experiment 3A tested performance with whole-object rotation when the vertices remained collinear. Grouping by 2-D cues here should lead to efficient search performance. Experiment 3B used stimuli that had both whole

and local vertex rotations, so that individual vertices were no longer collinear. Now neither 2-D nor 3-D constraints should support efficient grouping between elements; do effects of display size emerge even with cube displays under this circumstance?

Experiment 3A Collinear Vertices

Method

Twelve participants were used, 8 female and 4 male. All had either normal or corrected-to-normal vision. The stimuli were the same as the partial cubes used in Experiment 1C except that all the elements in each display were rotated so that in the vertex (cube) condition, the bottom-most element was a vertex pointing down (Figure 17). As a 3-D shape, the partial cube is no longer in a stable position. Effects should mean that 3-D information exerts a lesser impact here than when the bottom vertices were horizontal. Indeed, ratings collected from the same participants as those used for rating the displays in Experiments 2A and 2B suggested that the stimuli were seen less easily as single objects here than formerly. The mean rating of how object-like the 7-vertex displays were was 2.3, which was reliably lower than the ratings given to the partial cube displays used in Experiment 2B [$F(1,19) = 12.62, p < .01$], even though the vertices here were collinear and the vertices in Experiment 2B were not. There were 108 trials, 54 present and 54 absent, with 18 at each display size.

Figure 18. Mean correct reaction times (RTs, in milliseconds) with the partial cubes in Experiment 3A.

Results

The mean correct RTs are given in Figure 18 and the error rates are shown in Table 3.

There were no reliable effects on RTs of either condition (present vs. absent) or display size; nor was the condition $\times$ display size interaction significant (all F s < 1.0). More errors occurred on present relative to absent trials, but there were few effects of display size on errors.

Experiment 3A Versus Experiment 1C (Partial Squares)

To assess whether whole-object rotations disrupted performance, the data from Experiment 3A were compared with those from the partial cubes condition in Experiment 1C (where the stimuli were the same but the whole display had a stable base). The experiment was a between-subjects factor.

There was a main effect of experiment [$F(1,22) = 7.96, p < .01$], but no interactions involving this factor. RTs were overall slower in Experiment 3A, but effects of display size (in particular) were no larger when the whole figures were rotated. More errors also tended to be made in Experiment 3A, but again there was no tradeoff with display size.

Discussion

The effects of rotating the whole display in the plane were to slow overall RTs and to increase the error rate. Clearly participants found it more difficult to detect the target when both it and individual distractors were oblique than when they were horizontal and vertical (as in Experiment 1C, Figure 5). This may be due to the factors discussed after Experiment 2C, including the oblique effect. Interestingly, we found no evidence for increased

effects of the number of elements on search performance here relative to earlier. Grouping between the individual vertices seemed to be as efficient here as when the objects were in a physically stable position and were in a canonical orientation. Grouping occurred efficiently even though participants rated the stimuli as less objectlike. This suggests that grouping by collinearity is sufficient to minimize effects of display size in search. It also indicates that grouping by collinearity operated well even with line elements at oblique orientations. In Experiment 3B we examined performance when the influence of grouping by collinearity was reduced by rotating individual vertices in the plane. The stimuli were the same as in Experiment 2B with partial cubes except that the whole displays were transformed so that the bottom vertex faced down (Figure 19).

Experiment 3B Noncollinear Vertices

Method

There were 12 participants, 6 male and 6 female, all with either normal or corrected-to-normal vision. The procedure was exactly the same as for Experiment 2B with partial vertices, except that all

Table 3
Error rates in Experiment 3

Stimuli	Target Present: Display Size			Target Absent: Display Size		
	4	6	7	4	6	7
Experiment 3A						
"Partial cubes"	.03	.03	.04	.01	.01	.01
Experiment 3B						
"Partial cubes and vertex rotation"	.03	.04	.04	.01	.02	.02

Figure 19. Stimuli with whole-figure rotations and individual vertex rotations used in Experiment 3B.

displays were rotated in the plane. The mean rating of how object-like the 7-vertex stimuli were was 1.8. This was reliably lower than the ratings given to the same stimuli before whole-object rotation in Experiment 2B [$F(1,19) = 15.36, p < .01$].

Results

Mean RTs are shown in Figure 20 and the error rates are given in Table 3. There were reliable main effects of both condition and display size [$F(1,11) = 14.92$ and $F(2,22) = 16.36$, both $ps < .01$]. The condition $\times$ displays size interaction was also reliable [$F(2,22) = 13.67, p < .01$]. RTs increased with the display size from 4- to 6- and from 6- to 7-vertex displays for both present and absent responses ($p < .01$, Newman-Keuls for all comparisons). The slopes on the search functions were 20 and 26 msec/item on present and absent trials, and the linear components accounted for 99% of the variance in each case. There was no indication of a speed-accuracy tradeoff (Table 3).

Experiment 3B Versus Experiment 1C (Partial Squares)

Performance in Experiment 3B was contrasted with performance in Experiment 1C (nonrotated partial

squares). There was a main effect of experiment [$F(1,22) = 15.43, p < .01$], a two-way interaction between experiment and display size [$F(2,22) = 13.25, p < .01$], and a three-way interaction between experiment, condition (present/absent), and display size [$F(2,22) = 8.46, p < .01$]. There were stronger effects of display size in Experiment 3B than in Experiment 1C, and this was most pronounced on absent responses. For present trials alone, however, the interaction between display size and experiment remained significant [$F(2,22) = 6.39, p < .01$]. There were no signs of a speed-accuracy tradeoff across experiments.

Discussion

When whole figures were rotated into an unstable position, and when individual vertices were rotated to be no longer collinear, search efficiency decreased. Now there were reliable effects of the number of elements on search performance. This result contrasts with all other cases where we have examined search with partial cube stimuli, where minimal effects of display size have been apparent (Experiments 1C, 2B, and 3A). Rotating the individual vertices to make them noncollinear was not suf-

Figure 20. Mean correct reaction times (RTs, in milliseconds) in Experiment 3B (using stimuli with whole-figure rotations and individual vertex rotations).

ficient to disrupt performance with these displays when the partial cube was in a physically stable 3-D position (in Experiment 2B). Rotating the whole object but keeping the vertices collinear (Experiment 3A) was not sufficient either. There had to be changes in both the stability of the 3-D representation (due to whole-figure rotation), and in 2-D grouping by collinearity to disrupt performance. Under these circumstances it appears that grouping became less efficient and some serial components were introduced into search. The lack of interaction between 2-D rotation and display size in Experiment 2B can be attributed to the 3-D shape descriptions facilitating grouping. By rotating the whole object here, we reduced the likelihood that a stored 3-D description was activated, and there were then effects of the number of distractors in displays.

GENERAL DISCUSSION

We have reported three experiments that have examined whether 3-D object descriptions directly affect spatially parallel grouping of form elements. In Experiment 1 we found that as long as form junctions group into a single object description, search performance was little affected by the number of junctions present (Experiment 1A). This was not due to the ease of the particular search task (Experiment 1B). Also, when the vertices covered a matching spatial area but grouped into two separate object descriptions (the occluded squares displays), there was an RT increment relative to when the elements grouped into a single object description (Experiment 1C). This last result replicates a finding reported by Donnelly et al. (1991),

but extends it from performance with two spatially separated shapes to a situation with spatially overlapping shapes. The result is consistent with there being serial selection of visual objects even if the parts within each object are selected together (see also Duncan, 1984).

The objects used in Experiment 1A were cubes constructed from disconnected vertices. It may be that the 3-D representation of the cubes is derived only after grouping using 2-D constraints have been completed (e.g., based on 2-D collinearity or closure between the vertices). The same may also hold true of other studies where it has been argued that 3-D properties of displays determine search (e.g., Enns & Rensink, 1990, 1991; Humphreys et al., 1994). In Experiment 2 we tested whether 3-D descriptions directly influence grouping, in addition to any grouping effects occurring in 2-D. With occluded square displays, we found that performance was disrupted when the vertices were rotated so that they were no longer collinear: Effects emerged of the number of distractors present (Experiment 2A). This suggests that 2-D grouping is disrupted when collinearity is eliminated (see also Donnelly et al., 1991, Experiment 3). In contrast, with stimuli that yielded a 3-D description (full and partial cubes), rotating the vertices slowed performance overall but did not lead to the emergence of any effects of display size (Experiment 2B). This last result indicates that even with rotated and noncollinear vertices, the parts of cubes are encoded in parallel. Thus, a variable that is sufficient to disrupt grouping in 2-D (rotation of the individual vertices) was not sufficient to disrupt grouping with 3-D objects. Such a finding is not consistent with the view that 3-D object descriptions are de-

rived only following the grouping of elements in 2-D. However, it is consistent with the proposal that 3-D constraints are used directly to influence grouping.

In Experiment 3, we found that reducing the stability of the 3-D representation by whole-figure rotation did disrupt grouping and search. When whole figures were rotated into an unstable position in 3-D space and when component vertices were noncollinear, effects of the number of elements did influence performance (Experiment 3B). Introducing collinearity in 2-D helped in this case (Experiment 3A). These data are consistent with both 2-D and 3-D cues being used together to form shape descriptions in a spatially parallel manner. Partial grouping based on 2-D cues may be necessary to activate a 3-D shape description, but once this occurs, the 3-D description may help local elements so that they can be selected together.

Perkins (1968) first argued that if corner junctions in shapes are assumed to be orthogonal, the 3-D orientation of stimuli can be computed from the angles in the junctions (see also Enns & Rensink, 1991, and Muller & Dawson, 1990, for an extension of this proposal to a large class of polyhedral objects). If the 3-D orientation of the stimuli does not need to be recovered, but only the three-dimensionality of the object, then accurate computation of the angles between edges may not even be necessary; three-dimensionality may be computed with quite broad tolerance bands. Indeed, within the space of possible objects that may be consistent with the rotated vertices, there may be fewer 3-D than 2-D possibilities. For instance, the rotated vertices are consistent with a squashed cube in 3-D, whereas in 2-D they are not consistent with any known shape. Thus, inferences concerning the three-dimensionality of objects may be computationally useful in visual encoding. Once a 3-D interpretation is given to the display, the effects of junction rotation will be relatively slight, as we observed.

The present results also indicate that full junction information does not need to be present for 3-D object descriptions to be derived; the vertices alone are sufficient. Enns and Rensink (1991), however, failed to find spatially parallel search with multiple cube stimuli constructed only from vertices. Does this mean that vertices can only be grouped into a single object description, and that there are limits on the computation of multiple object representations from the same information? Clearly this is an empirical issue. One factor that may mitigate against parallel search with multiple objects constructed from vertices, though, is grouping across rather than within objects. With multiple objects, neighboring vertices can be similar even when they are parts of different objects. Grouping between similar neighbors from different objects will disrupt search, particularly when the neighbors are parts of target and distractor objects (see Duncan & Humphreys, 1989). Interestingly, using displays with multiple Kanisza figures, Davis and Driver (1998) found evidence for parallel grouping within multiple objects. They found that it was difficult to detect targets that had a missing section when this section was occluded by the

illusory shapes formed by grouping within the figures (see also Rensink & Enns, 1998). Grouping across objects may have been limited in this study by strong within-object grouping based on the new illusory shapes. The result suggests that grouping can operate in parallel within multiple figures, although in a competitive way with grouping across figures.

To date, evidence for parallel encoding of parts into a shape description has come from studies using relatively simple and familiar shapes—squares, pentagons, hexagons, cubes, and so on (e.g., Donnelly et al., 1991). Using a similar procedure to that employed here (manipulating the number of parts in discrimination tasks), Baylis and Driver (1994) found that the number of parts making up random shapes had little effect on symmetry judgments, but this factor had strong effects on judgments of whether parts on each side of the shape were repeated. This suggests that parts can be assembled in parallel even with unknown shapes. However, the evidence with such shapes is confined at present to symmetry detection, which may not require full articulation of the parts of objects or explicit encoding of the shape description, as may be required for the “part” detection tasks used here. Humphreys and Watson (2000) used a visual matching task with face stimuli in which participants had to decide whether two simultaneously presented faces were identical. On *different* trials, the identity of one facial feature was changed across the faces. Across trials, the number of features present in the faces was varied. Humphreys and Watson found little effect of increasing the number of features when the identity of one of the two faces was maintained across trials. However, there were large linear effects of the number of features when the identities of the faces varied across trials. Humphreys and Watson suggested that the features in the faces were bound together in a spatially parallel manner by matching to a memory template of the face that was constant across trials. When face identities varied, so that a memory template could not be used, serial matching of faces took place. In this case, feature binding seemed to require access to stored memory representations. Whether this holds for common objects other than faces is an issue facing future research.

More generally, the present data indicate that descriptions of 3-D objects are used directly in visual grouping. The primitive 3-D shape descriptions that result can be used to guide immediate action and object recognition, even if more detailed (spatially serial) processing is required for finer recognition to occur (see Kosslyn, Alpert, & Thompson, 1995, for PET evidence on this).

REFERENCES

- APPELLE, S. (1972). Perception and discrimination as a function of stimulus orientation: The “oblique effect” in man and animals. *Psychological Bulletin*, *78*, 266-278.
- BAYLIS, G. C., & DRIVER, J. (1993). Visual attention and objects: Evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception & Performance*, *19*, 451-470.

- BAYLIS, G. C., & DRIVER, J. (1994). Parallel computation of symmetry but not repetition within single visual shapes. *Visual Cognition*, **1**, 377-400.
- CRICK, F. (1984). Function of the thalamic reticular complex: The search light hypothesis. *Proceedings of the National Academy of Sciences*, **81**, 4586-4590.
- DAVIS, G., & DRIVER, J. (1998). Kanisza subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 169-184.
- DONNELLY, N., HUMPHREYS, G. W., & RIDDOCH, M. J. (1991). Parallel computation of primitive shape descriptions. *Journal of Experimental Psychology: Human Perception & Performance*, **17**, 561-570.
- DONNELLY, N., WEEKES, B., HUMPHREYS, G. W., & ALBON, A. (1998). Processes involved in the computation of a shape description. *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 1119-1130.
- DUNCAN, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, **113**, 501-517.
- DUNCAN, J., & HUMPHREYS, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, **96**, 433-458.
- ELDER, J., & ZUCKER, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, **33**, 981-991.
- ENNS, J., & RENSINK, R. A. (1990). Influence of scene-based properties on visual search. *Science*, **247**, 721-723.
- ENNS, J., & RENSINK, R. A. (1991). Preattentive recovery of three-dimensional orientation from line drawings. *Psychological Review*, **98**, 335-351.
- ESSOCK, E. A. (1980). The oblique effect of stimulus identification with respect to two classes of oblique effects. *Perception*, **9**, 37-46.
- FOSTER, D. H., & WARD, P. A. (1991a). Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society of London: Series B*, **243**, 75-81.
- FOSTER, D. H., & WARD, P. A. (1991b). Horizontal-vertical filters in early vision predict anomalous line-orientation identification frequencies. *Proceedings of the Royal Society of London: Series B*, **243**, 83-86.
- GILCHRIST, I., HUMPHREYS, G. W., RIDDOCH, M. J., & NEUMANN, H. (1997). Luminance and edge information in grouping: A study using visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 464-480.
- HUMPHREYS, G. W., KEULERS, N., & DONNELLY, N. (1994). Parallel visual coding in three dimensions. *Perception*, **23**, 453-470.
- HUMPHREYS, G. W., QUINLAN, P. T., & RIDDOCH, M. J. (1989). Grouping processes in visual search: Effects with single- and combined-feature targets. *Journal of Experimental Psychology: General*, **118**, 258-279.
- HUMPHREYS, G. W., & WATSON, D. G. (2000). *Parallel template-based matching of facial features*. Manuscript submitted for publication.
- KOSSLYN, S. M., ALPERT, N. M., & THOMPSON, W. L. (1995). Identifying objects at different levels of hierarchy: A position emission tomography study. *Human Brain Mapping*, **3**, 107-132.
- MULLER, J. A., & DAWSON, R. J. M. (1990). Reconstructing polyhedral scenes from single two-dimensional images: The orthogonality hypothesis. In P. K. Patel-Schneider (Ed.), *Proceedings of the Eighth Biennial Conference of the CSCI* (pp. 238-244). Palo Alto, CA: Morgan Kaufman.
- PERKINS, D. N. (1968). Cubic corners. *MIT Research Laboratory of Electronics Quarterly Progress Report*, **98**, 207-214.
- POTTER, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning & Memory*, **2**, 509-522.
- QUINLAN, P. T., & HUMPHREYS, G. W. (1987). Visual search for targets defined by combinations of color, shape, and size: An examination of the task constraints on feature and conjunction searches. *Perception & Psychophysics*, **41**, 455-472.
- RENSINK, R. A., & ENNS, J. (1995). Preemption effects in visual search: Evidence for low-level grouping. *Psychological Review*, **102**, 101-130.
- RENSINK, R. A., & ENNS, J. (1998). Early completion of occluded objects. *Vision Research*, **38**, 2489-2505.
- THORPE, S., & IMBERT, M. (1989). Biological constraints on connectionist modelling. In R. Pfeifer, Z. Schreier, F. Fogelman-Soulie, & L. Steels (Eds.), *Connectionism in Perspective* (pp. 63-92). Amsterdam: Elsevier.
- TOVEE, M. J., & ROLLS, E. T. (1995). Information encoding in short firing rate epochs by single neurons in the primate temporal cortex. *Visual Cognition*, **2**, 35-58.
- TREISMAN, A. (1985). Preattentive processing in vision. *Computer Vision, Graphics & Image Processing*, **31**, 156-177.
- TREISMAN, A. (1988). Features and objects: The Fourteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, **40A**, 201-237.
- TREISMAN, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London: Series B*, **353**, 1295-1306.
- TREISMAN, A., & GELADE, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, **12**, 97-136.
- TREISMAN, A., & SOUTHER, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, **114**, 285-310.

NOTES

1. A preliminary outline of the present experiments was presented in Humphreys et al. (1994). The full data, details of the method, and discussion of the results were not reported.
2. We describe the results in terms of the number of vertices present, in keeping with visual search paradigms where participants may search all the individual stimuli present. This also enables us to compare directly performance on the grouped displays with performance where grouping does not take place (with randomly arranged vertices). With the grouped displays it is possible to analyze performance in terms of the number of "faces" visible in the shapes, but essentially the same findings follow.
3. Though this is the formal definition of a target, the results demonstrate that target detection was contingent on the context provided by the distractor vertices.
4. Two types of target were used for two reasons. First, reflected targets mirrored the items used in Donnelly et al. (1991). Second, the missing targets were necessary for the "random vertex" control condition (see Experiment 1B), which provided a baseline measure of discrimination when the cube shape was absent. Reflected targets could not be used in the control since, in that condition, there was no way to define a difference between a reflected target and a distractor vertex.
5. With the rotated full and partial cube displays, participants remarked that they tended to see "squashed" cubes, consistent with the stimuli still being encoded as 3-D objects. Interestingly, with the rotated versions of the occluded square displays, participants sometimes reported just seeing the displays as 2-D patterns, rather than seeing two squares in depth. Such reports suggest that the grouping of these displays into descriptions of occluded shapes broke down on at least some occasions.

(Manuscript received September 15, 1996;
revision accepted for publication May 1, 1999.)