

Early completion of occluded objects

Ronald A. Rensink^{1,a,*}, James T. Enns^b

^a*Cambridge Basic Research, Nissan Research & Development, Inc., 4 Cambridge Center, Cambridge, MA 02142-1494, USA*

^b*Department of Psychology, University of British Columbia, Vancouver, BC Canada*

Received 5 February 1997; received in revised form 17 September 1997; accepted December 1997

Abstract

We show that early vision can use monocular cues to rapidly complete partially-occluded objects. Visual search for easily-detected fragments becomes difficult when the completed shape is similar to others in the display; conversely, search for fragments that are difficult to detect becomes easy when the completed shape is distinctive. Results indicate that completion occurs via the occlusion-triggered removal of occlusion edges and linking of associated regions. We fail to find evidence for a visible filling-in of contours or surfaces, but do find evidence for a "functional" filling-in that prevents the constituent fragments from being rapidly accessed. As such, it is only the completed structures—and not the fragments themselves—that serve as the basis for rapid recognition.

Keywords: Visual completion; Visual search; Object recognition; Preattentive vision; Early vision

1. Introduction

One of the more remarkable qualities of human vision is its ability to compensate for missing visual information. Every student of vision can probably recall their sense of disbelief upon first being told that each eye contains a blindspot at the optic disk. But except under controlled conditions, we are largely unaware of this hole in our visual input (Helmholtz, 1867/1962). The visual system appears to fill in this spot with the colors, textures, and forms appropriate to that part of the visual input (Ramachandran, 1992). Interestingly, physiological studies of monkey brain show that the mechanisms underlying such compensation extend all the way down to Area V1, the first stage of cortical processing (Fiorani, Marcello, Gattass, & Rocha-Miranda, 1992).

Although some information about the world is always lost at the sensory interface, much more is usually lost via occlusions by external objects along the observer's line of sight. Such occlusion is pervasive in the real world, so much so that many (if

not most) objects viewed are only partially visible at any given time. And just as we are generally unaware of the interruptions caused by physiological factors, so are we unaware of the interruptions caused by external occluders. If pressed, we can summon our attention and see that the occluded sections of an object are indeed absent in the visual field. But what is actually happening here? Is the "fragmented" interpretation generally used, with completion taking place later on if required? Or are objects completed at an early stage, and our perception of the visible fragments the result of later processes that undo this completion? Can completion be identified with the grouping that links nonoccluded items? Or does it have its own unique characteristics specialized for visible occluders? And does the completion process itself posit new visual elements, e.g., extending contours and filling in surfaces? Or does it simply impose a nonvisual structure onto the elements already present?

The experiments presented here show that partially occluded objects are indeed completed rapidly at early levels of vision. Such completion is found to have many of the characteristics of other rapid-interpretation processes, namely, a use of simple rules based on local context, with operations carried out rapidly and in parallel across the visual field. However, it is also found to have characteristics that differentiate it from the grouping processes that exist at those levels. The

* Corresponding author. Current address: Department of Psychology, University of British Columbia, Vancouver Canada. Fax: +1 604 822-6923; email: rensink@psych.ubc.ca

¹ Much of this work was done while the first author was with the Dept. of Computer Science, University of British Columbia. Parts of this work were presented at the 1992 meeting of the Association for Research in Vision and Ophthalmology, Sarasota, Florida.

experiments yield little evidence that new visual elements are posited. Instead, rapid completion appears to remove occlusion edges and link fragments so thoroughly that a *functional* form of filling-in occurs: it is only the completed structures—and not their constituent elements—that become the effective units of subsequent recognition processes.

1.1 Basic Issues

If it is to be reliable, a recognition system must be able to compensate for the loss of information at various locations in the visual field. Such localized losses are pervasive, originating both in the world (via occlusion by external objects) and in the sensor array itself. Compensation is generally believed to take the form of completion, i.e. forming a representation in some sense the same as that corresponding to an uninterrupted structure. This process is rather complex and not well understood; indeed, relatively little is known even about the ways in which a completed structure might be similar to a representation of its uninterrupted counterpart (see, e.g. Jacobs, 1992). In what follows, we will show that new light can be cast on our understanding of completion by examining the extent to which it is carried out in early vision.

1.1.1. Early versus later completion of visual structure

At what processing levels is visual completion carried out? Could it be done at early levels, where processes are simple and rapid? Should it be done there? A useful place to begin answering these questions is by considering the ways in which localized loss of information affects visual processing.

(i) *Interference with surface recovery.* Localized loss often produces 'holes' that can—if large enough—split a region into disconnected pieces. In contrast, occlusion can cause disconnected surfaces to project to adjacent regions in the image. Thus, image neighborhoods no longer correspond directly to surface neighborhoods in the world. This can create significant problems for surface-recovery algorithms (Williams & Hanson, 1996).

(ii) *Interference with object grouping.* Localized losses can also split apart items in the image that originate from the same object. An important part of object recognition involves forming groups of such items, since these groups can reduce the complexity of the recognition process while increasing its discriminative power (Jacobs, 1996; Lowe, 1985). The interruptions in the image create uncertainties as to which items belong together, thereby interfering with the formation of such groups (Jacobs, 1992).

Note that these disruptions occur at levels that either precede or are involved with the formation of object representations. As such, it would seem that the earlier completion occurs, the better. However, early-level processing is believed to involve simple localized operations carried out on relatively simple undifferentiated structures. Could such processes be

sophisticated enough to carry out the operations required for completion? Could some aspects of completion be sufficiently simple to be carried out by such processes?

Relatively little is known about this issue. It has been suggested (Barrow & Tenenbaum, 1978) that early processes should only be concerned with assigning relatively simple scene properties to the image. It has also been suggested (Lowe, 1985; Marr, 1982) that grouping could occur at early levels, but that completion would not; this would be done at a later stage, where three-dimensional (3-D) object models would be matched against the grouped fragments. However, if early vision tries to obtain as much structure as possible by "quick and dirty" techniques (Enns & Rensink, 1992; Rensink, 1992; Rensink & Enns, 1995) the possibility exists that it might not be content with simple grouping, but may attempt to go beyond the information given and complete various aspects of object structure as soon as possible.

1.1.2 Completion versus grouping

Whenever studying a completion process, an important issue is its relationship to grouping. This relationship can in principle take on a number of forms, ranging from complete identity to a complete separation of processes. To see how this comes about, consider the different ways in which visual information about the world can be lost:

(i) *Gaps in the sensor array.* These can be due not only to blindspots, but also to large blood vessels in the retina or lesions at any point along the visual pathway (see, e.g., Ramachandran & Gregory, 1991). In all cases, the location of the gap is fixed. This allows the system to learn to compensate for the loss; any interruption of the stimulus at that point can be reasonably attributed to the sensor and not to the world.

(ii) *Imperfect coupling between sensor and world.* This is due to sensor effects such as stimulus interactions causing loss of edge continuity (see, e.g., Marr, 1982) or noise washing out signals too low in contrast. As in the case of sensor gaps, signals can be lost from large regions of the visual field. But here, the areas of loss do not have fixed locations. It is therefore not always possible to determine whether a lack of signal originates in the sensor, or whether it corresponds to an interruption in the world.

(iii) *Occlusion in the world.* This is caused by opaque objects along the observer's line of sight. As in the case of imperfect coupling, the locations of signal loss are not fixed. However, in this case there is a visible cause for the loss of the signal, so that the loss of information can be attributed to the world, and not to the sensor.

Each of these factors is different and so in principle may require a different strategy if compensation is to be maximally effective. For sensor gaps, information loss can be reasonably ascribed to the sensor, and so

compensation is made to the representation of the world. Indeed, the high reliability of this ascription may be the reason why the filling-in here is effectively irreversible.

For interruptions caused by sensor coupling, a more complex situation arises. Instead of evidence of absence (as in the case of sensor gaps), there is now only absence of evidence. Compensation must therefore be more cautious and tentative. This often takes the form of *grouping*, which places a nonvisual link between those fragments thought to belong to the same object in the world. Other compensation mechanisms include the formation of illusory lines or surfaces (see, e.g. Grossberg & Mingolla, 1985). In all these cases, there is no confusion between the illusory and the real, and the fragments in the image remain readily visible.

For interruptions due to occlusion, the situation differs yet again. Unlike the case of sensor gaps, the locations of loss are not fixed. Unlike the case of sensor coupling, there is evidence of absence. And unlike both, two levels of structure—occluded as well as occluding—may need to be represented simultaneously (Williams & Hanson, 1996). Even ignoring issues of structure, however, compensation for occlusion is simply a more difficult task, since it requires identifying which items in the image correspond to the occluded object and which to the occluder (e.g. Jacobs, 1992, ch. 7). Completion of occluded stimuli may involve some of the same mechanisms used for the other two cases (see e.g., Durgin, Tripathy, and Levi, 1995; Kellman & Shipley, 1991). However, it could also involve different processes specialized for the constraints peculiar to this type of information loss.

In what follows, we will consider only this latter type of completion. As such, the term "completion" will be used here as shorthand for "completion of occluded objects", also known as "amodal completion" (e.g., Gerbino & Salmaso, 1987).

1.1.3. Nature of the completed structures

If rapid completion of occluded objects does occur, another important issue is the nature of what might be posited to compensate for localized information loss. In what ways might a completed structure be the same as a representation of the corresponding uninterrupted structure? At least four types of compensation are possible:

(i) *Total restoration* : Positing of all occluded visual elements, mainly via extension of existing boundaries and surface properties. This would be an image of the object as it would appear in the absence of occluders. This kind of restoration has been proposed for sensor gaps (Ramachandran, 1992).

(ii) *Boundary restoration*: Positing of boundaries (via extension), but not surface properties. Gaps in contours are effectively removed, so that object shape is restored. Related fragments are linked by virtue of the completed boundaries that connect them. Note that

boundaries and surfaces may well be handled via different processing streams; if so, completion processes could easily be split along these lines (see e.g., Grossberg & Mingolla, 1985).

(iii) *Surface restoration*: Positing of surface properties (via extension), but not boundaries. Related fragments are linked via the contiguous "stuff" (or material) posited to exist between them. Without completed boundaries, the shape of the completed region is at least partly indeterminate.

(iv) *Functional restoration*: No positing of new visual elements. This process is amodal in the strictest sense possible: related fragments are linked entirely by abstract, nonvisual structures. In some sense, this could be considered a special form of grouping triggered by pictorial cues to occlusion.

A related set of issues concerns the status of the fragments linked together in the completed structures. Are these fragments still accessible, so that they can be used whenever needed? Or are they preempted by the completed structures so that they can no longer be readily accessed (Rensink & Enns, 1995)?

1.2 Previous Work

1.2.1. Computational Studies

The early completion of occluded structure is a problem difficult even to formulate clearly, never mind solve. Nonetheless, computational studies have been able to provide some insight into its nature. Many of these studies were based on the grouping of image segments into contours corresponding to the outlines of individual objects. For example, segments have been collected together on the basis of local properties such as segment co-termination (Lowe, 1985), or more global properties of the contours such as their overall smoothness and length (Sha'ashua & Ullman, 1988), closure (Elder & Zucker, 1996), or convexity (Jacobs, 1992, 1996). Such approaches generally made no distinction between occluding and occluded structures during the formation of the groups, preferring to let this determination be a natural outcome of the grouping process.

Although techniques based on these approaches can adequately compensate for losses due to sensor interactions, they often have difficulty with interruptions due to occlusion (see, e.g., Jacobs, 1992). This suggests that something more is needed, presumably something involving the determination of occluding and occluded structure during the completion process itself. One interesting suggestion in this regard is to decompose completion into two independent subprocesses, with the determination of shape largely decoupled from the determination of which fragments belong together (Williams & Hanson, 1996).

1.2.2. Psychophysical Studies

A great deal of experimental work has been carried out on the nature of visual completion and perceptual organization (see, e.g., Pomerantz & Kubovy,

1986). However, studies concerned with isolating the levels involved have been relatively rare. Evidence for completion effects in low-level vision (i.e., vision not involving stimulus-specific knowledge) was found by Weisstein, Montalvo, and Ozog (1972). Observers viewed a grating for an extended period (the adaptation phase) and were then asked to make contrast judgments for a small target in the center of the adapted field (the test phase). Adaptation effects were stronger when a drawing of a 3-D cube was placed in the center of the field than when a blank hexagon was placed there. This suggested that the low-level mechanisms involved in contrast perception could be affected by a process that filled in regions perceived as being occluded.

Completion at low levels has also been reported in studies on binocular stereopsis. Nakayama, Shimojo, & Silverman (1989) used stereo-depth displays in which mosaic fragments (i.e., the fragments corresponding to the purely visible parts of occluded objects) were positioned in one of two depth planes. Identification was much more accurate for fragments in the far plane than for fragments in the near plane. This is consistent with a process that integrated fragments better when they were behind an occluder than when in front. Similarly, He & Nakayama (1992) showed that visual search for Ls against upside-down Ls in the far plane was slow when they were occluded by squares set in the near plane, but fast when the order of the planes was reversed. This indicated a degree of completion for occluded objects, with the completed objects more alike and the targets therefore harder to find.

Investigators have also examined the time course of completion. Gerbino and Salmaso (1987) showed that same-different judgments for complete and occluded shapes could be made as rapidly as for pairs of complete shapes or pairs of occluded shapes. Mosaic fragments required more time to match, suggesting that the completed shape was accessed more quickly than the mosaics. Sekuler & Palmer (1992) used a speeded priming paradigm to show that the completed representation of an occluded object is available within 200 ms of display onset.

1.3. Does Completion Occur at Early Levels?

Although the adaptation and stereopsis studies point towards the existence of a completion process at low levels, they suffer from the drawback that a relatively large amount of time must elapse before the critical percept is formed. As such, it is difficult to determine how much of their effects are due to rapid, early processes, and how much to slower processes that feed information back to lower levels. Conversely, the time-course studies were not designed to isolate the particular levels involved: for example, since only a few items were ever viewed at a time, the effects of focused attention could not be separated out. Thus, although previous research has shown that completion can occur at low levels and also that it can be rapid, it has not shown that it is *early*, i.e., that it is both low-level and rapid. Furthermore, it did not ascertain the

nature of processes based on monocular cues, or the type of completion that occurred.

The experiments presented here examine whether completion does occur in early vision, and if so, what its main characteristics might be. In particular, they examine whether completion occurs at *preattentive* levels, where operations are believed to be rapid, automatic, and carried out in parallel across the visual field (see, e.g., Beck, 1982; Julesz, 1984; Treisman, 1986). At these levels, a distinctive target can be detected quickly by visual search, the speed of search reflecting the degree to which the target differs from the distractor elements (see, e.g., Treisman & Gormican, 1988). Our question therefore is this: if the target and distractor items in a search task contain monocular cues for completion, is speed governed by the distinctiveness of the fragments or of the completed figures? If the latter is the case, the rapid response times associated with monocular cues will show the existence of an object-completion process that acts both rapidly and in parallel across the visual field.

It is of course possible that completion occurs only at the small region of visual space that is focally attended. Outside of this region the representation of objects may be much simpler, perhaps consisting only of a mosaic of shapes and colors. In fact, this exact situation would be expected from feature integration theory (Treisman, 1986; Treisman & Gormican, 1988). Here, an attentional spotlight is needed to establish spatial relations between components and to "glue" them into coherent objects; outside the spotlight, components are disconnected and free-floating. In this view, the only way an array of completed objects could be formed would be for attention to inspect each cluster of image features in turn.

However, recent studies have shown the existence of rapid-interpretation processes that yield preattentive descriptions concerned with scene-based rather than image-based properties (e.g., Enns & Rensink, 1990, 1991; Epstein, Babler, & Bownds, 1992; Ramachandran, 1988; Rensink & Cavanagh, 1993; Sun & Perona, 1996; Wolfe, Friedman-Hill, & Bilsky, 1994). The existence of such processes suggests that a primary goal of early vision is to recover as much scene structure as possible within a few hundred ms (Enns & Rensink, 1992; Rensink, 1992). If so, it may be that some aspects of object completion are also carried out rapidly and in parallel at these levels.

In what follows, we examine whether early visual processes can carry out the most computationally demanding type of completion task: the completion of occluded objects. We examine general issues such as the existence of rapid completion, its relation to rapid grouping, and the nature of the filling-in process. As a consequence, the particular structures that may be involved—such as illusory contours or surfaces—will not be a central focus of the work here.

Finally, it is important to keep in mind that search tasks are not based directly on subjective estimates of stimulus qualities. Instead, they examine the factors




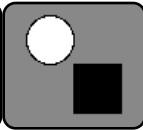




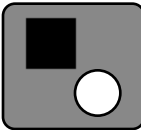
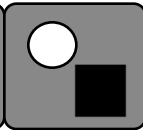






Condition	Targets		Distractors		Rate (ms/item)	
					Present	Absent
1A (Mosaic)					7	8
1B (Occlusion)					36	66
1C					20	45
1D					35	83

Fig. 1 Search items and results of Experiment 1 (Basic Effect). Search is rapid for the free notched squares (Condition 1A), but slow when these contact occluding disks (Condition 1B). Reversing the choice of target and distractor items in Condition 1A causes search to slow down (Condition 1C), showing that the fast search was due to a distinctive feature in the notched squares. A similar reversal of the items of Condition 1B (Condition 1D) shows no such asymmetry, indicating that the occluded squares have no distinctive feature.

that affect search speed, and these factors may or may not correspond to the percepts formed after the action of subsequent processes. But although this may cause some difficulty in relating our results to conscious experience, it has the advantage of a genuinely different perspective, one possibly much closer to the primary processes involved in object recognition.

2. General method

Each of the experimental conditions employed well-known visual search methodology in which observers searched as rapidly as possible for a pre-defined target among a set of distractor items that varied in number (e.g., Enns, 1992; Enns & Rensink, 1991; Treisman & Gormican, 1988). Displays were formed of 2, 8 or 14 items, chosen at random. The target was present on half the trials (chosen randomly) and absent in the other half. Observers were asked to determine the presence or absence of the target as quickly as possible, while maintaining an accuracy level of at least 90%. The primary dependent variable was *search rate*, defined as the slope of the correct response time (RT) over display size. In all cases, the error data were consistent with the RT data, ruling out the possibility of speed-accuracy tradeoffs influencing the results.

Items were positioned randomly in each display, on an imaginary 6 x 4 grid of possible locations. The display area subtended approximately 12° x 8° of visual angle, with each item less than 2° in diameter. In addition, the position of each item was jittered by ±0.5° to minimize the use of item collinearity to aid search.

Each experimental condition tested 10 adult observers with normal or corrected-to-normal visual acuity. Half the observers were naive to RT testing and visual search methodology, while the other half had been tested extensively in other search tasks.

Data analyses focused on mean search rates for present and absent displays in each experiment. Differences between search rates were tested with between-groups analysis of variance when more than two conditions were compared and with independent *t*-tests for simple comparisons. Although some overlap existed in observers in various conditions, none of the comparisons took advantage of this (i.e., none relied on within-subjects analyses). Instead, the more conservative assumption was made that observers were sampled independently. All reported differences were significant at the $p < .05$ level or better. A table of the mean RTs, standard errors of the means and accuracy rates for all conditions is given in Appendix A.

3. Experiments 1-3: Rapid completion of contiguous fragments

The first set of experiments examined whether rapid completion occurs for fragments that remain contiguous when partially occluded. Here, observers searched for a black fragment with a unique shape. In all conditions these fragments were either a "notched" square (i.e., a black square with a notch in one of two possible locations) or an "unnotched" square. Both notched and unnotched squares were paired with white disks, the spatial relations of the disks and squares

Cond.	Targets		Distractors		Rate (ms/item)	
					Present	Absent
2A					6	8
2B					19	44
2C					25	50
2D					5	8

Fig. 2. Search items and results of Experiment 2 (Effect of Depth Ordering). Search for the free notched squares remains faster than for the completed squares, even when different depth planes are involved (Conditions 2A and 2B). The similarity of speeds in Conditions 2B and 2C shows that search for items cannot be based on depth ordering *per se*. Condition 2D shows that search can be fast, even when critical fragments are further away in depth.

being varied across the different conditions (see e.g., Fig. 1). If search can be based on distinctive fragments, the differences in spatial relations should have no effect on search; otherwise, a more complex pattern of search rates will result.

3.1. Experiment 1: Basic Effect

In Condition 1A (Mosaic), observers searched for a notched square against a set of unnotched squares. Squares and disks were kept separate, so that no completion would be expected (Fig. 1A). This condition was designed to be relatively easy, allowing it to be used for comparison against other conditions. Search here was indeed rapid: baseline RTs (extrapolated RTs for 1 object) were between 500 and 600 ms, while search rates had mean RT slopes of 7 ms/item for target-present trials and 8 ms/item for target-absent trials. These values are comparable to those found in other kinds of rapid search (e.g., Treisman & Gormican, 1988).

In Condition 1B (Occlusion), the squares were displaced slightly toward the disks so that they appeared to be completed (Fig. 1B). If the spatial relations of the disks and the squares play no role, or if observers could simply focus on the black fragments, search should be as fast as in the previous condition. However, the displacement caused a dramatic slowdown in search (36 and 66 ms/item), indicating that the targets no longer contained a distinctive visual feature.

Condition 1C switched the targets and distractors of Condition 1A. Search was significantly slower here than in the "unreversed" condition (20 and 45 ms/item), showing that the unnotched squares had no

distinctive features that could speed up search. Condition 1D was a similar switch of the items of Condition 1B. In contrast to the asymmetry found in Conditions 1A-1C, rates here remained the same (35 and 83 ms/item). Search in Condition 1C was significantly faster than in Condition 1D, suggesting that observers might have used the "free" notches to help guide search (e.g., Wolfe, Cave, & Franzel, 1988), or to check each distractor a little more quickly. Thus, the results of these conditions suggest that the free notches give rise to a distinctive feature that is lost when they contact occluders.

3.2 Experiment 2: Effect of Depth Ordering

At least two different factors could explain why notches no longer support rapid search when they contact occluders: (i) an increased similarity between targets and distractors (e.g., Duncan & Humphreys, 1989), or (ii) difficulties in accessing items when fragments are perceived as being at different depths. To determine the extent to which the latter explanation can account for the slowdown, Experiment 2 examined the influence of depth ordering.

In Condition 2A observers searched for a notched square against unnotched squares. All squares were now positioned so that they overlapped the disks, allowing observers to use the notch as they had in the Mosaic condition of Experiment 1, but with the display containing pictorial cues to occlusion. Condition 2B had similar stimuli, but with target disks moved so that the notches were no longer free. If the mere existence of occlusion made search in Experiment 1 more difficult, search should be slow in both Conditions 2A

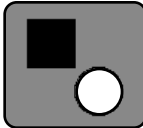
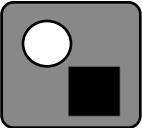

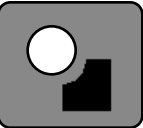






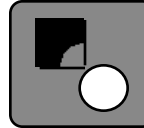
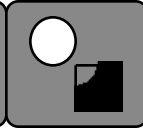




Cond.	Targets		Distractors		Rate (ms/item)	
					Present	Absent
1C					20	45
3A					64	121
3B					30	53
3C					6	15

Fig. 3. Search items and results of Experiment 3 (Type of Completion). For purposes of comparison, an upper row has been added showing the slopes of Condition 1C. Search becomes extremely slow when notched squares are occluded by disks (Condition 3A), indicating that the completed target squares have a form highly similar to the notched squares. Adding square outlines to the distractors (Condition 3B) causes search to speed up, indicating that the distractor squares have now become less (rather than more) similar to the targets. Hiding part of the distractor squares via occlusion (Condition 3C) causes search to speed up even more, something not to be expected if the form of the squares had been maintained by the completion process.

and 2B. Otherwise, Condition 2A should be faster than in Condition 2B, with speeds comparable to those of Condition 1A. Results (Fig. 2) showed that observers could easily detect the mosaic targets of Condition 2A (6 and 8 ms/item), and that this was indeed much faster than for the occlusion targets of Condition 2B (19 and 44 ms/item). Thus, the simple existence of depth ordering did not cause the slowdown found in Experiment 1.

It might be argued that depth order itself acted as a feature, with items in front given preferential status, or items in back somehow suppressed. To test this, Condition 2C switched the items used for targets and distractors in Condition 2B. If some kind of feature assignment were involved, this should speed up search. However, the results (Fig. 2C) showed no such speedup—in fact, a small (though not significant) slowdown occurred. Thus, feature assignment based on depth ordering does not appear to take place.

Another possibility is that it might be more difficult to access items perceived as being further away in depth. In Condition 2D, notched targets were used, but with the critical fragments appearing behind the disks. If depth order were the critical factor, search should be relatively slow; if not, search should be as fast as for the other notched targets. The results (Fig. 2D) were clear: search for occluded targets was as fast as in the other mosaic conditions (5 and 8 ms/item). Taken together, these results show that depth ordering *per se* has little effect on rapid access.

3.3. Experiment 3: Type of Completion

Given that depth ordering cannot explain the slowdown found in Experiment 1, this effect would appear to be due to the increased target-distractor similarity caused by a rapid completion process. But if so, what type of completion takes place at these levels? What is and is not posited?

In Condition 3A, items were the same as in Condition 1C, except that target squares were moved behind the disks so that they were now partially occluded (Fig. 3). If the completed squares were visually similar to the unnotched squares, speeds should be largely unaffected, since only a slight shift in the position of the target parts was introduced. However, the results (Fig. 3A) showed a dramatic slowdown (64 and 121 ms/item). Evidently, the completed square and the notched square were seen as much more similar than the unnotched square and the notched square. This argues against any great degree of boundary or surface extension.

This conclusion is further supported by the results of Condition 3B. Here, the distractor boundaries were augmented by visible segments to yield square outlines (Fig. 3B). If the boundaries in the targets were extended to form squares, the greater target-distractor similarity should cause search to slow down. Instead, search sped up by a factor of 2 (30 and 53 ms/item), indicating that the items were less similar.



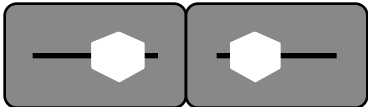
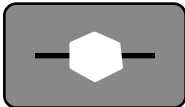
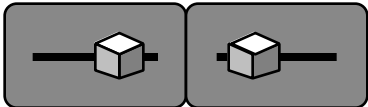
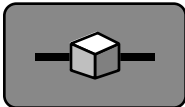
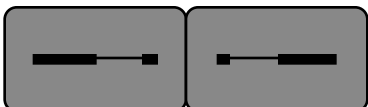
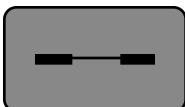
Cond.	Targets	Distractors	Rate (ms/item)	
			Present	Absent
4A			6	4
4B			29	55
4C			46	86
4D			7	15

Fig. 4. Search items and results of Experiment 4 (Preemption of Separated Fragments). Conditions 4A-4C are *distinct-fragment* tasks. Search here becomes *slower* as completion becomes stronger, showing that the fragments cannot be rapidly accessed. When fragments are connected by a visible line (Condition 4D), they do become accessible, showing that the difficulty in access is caused by the presence of the occluders.

In Condition 3C, the distractor fragments of Condition 3A were moved so that they were partially occluded. If visual elements were posited to compensate for this occlusion, search should remain largely unaffected. Instead, it sped up greatly (6 and 15 ms/item), indicating that a distinctive feature had emerged. Such a speedup is difficult to explain if boundaries had been completed and surfaces filled in. But it could be easily explained in terms of the mosaic (i.e., visible) aspects of the squares, such as the width of the targets, or the oblique orientations of the distractors.

If there is no significant extension of boundaries or surfaces, what might account for the slowdown encountered in Experiment 1? Since search in Condition 1A was rapid, the black fragment must have contained some distinctive visual feature not found in Condition 1B. The only differences between the targets and distractors in the two conditions were the short boundary segments and curved segments caused by the notches. Given that no great amount of boundary extension takes place (Condition 3B), the short segments should give rise to similar structures in both conditions. The key factor would therefore appear to be the curved segments, which are no longer accessible when placed against an occluding disk. Evidently, although rapid completion of contiguous fragments does not extend boundaries or fill in surfaces to any great extent, it is able to remove edges caused by the presence of occluders.

4. Experiments 4-6: rapid completion of separated fragments

Having established that rapid completion does exist but that it does not restore occluded regions, the

next step is to determine if it can at least link separated image fragments that correspond to the same object. To do this, we used items generally having the form of a solid bar split into two fragments separated by 1.0° (Fig. 4). The relevant properties here are (1) the shapes of the completed bars, and (2) the shapes of the fragments. As in the previous experiments, these critical fragments (i.e., the bars) were usually accompanied by a secondary item common to both targets and distractors.

4.1. Experiment 4: Preemption of separated fragments

In this experiment, targets differed from distractors in fragment length but not in overall length (see Fig. 4). If the fragments could be rapidly accessed, search in these *distinct-fragment* conditions should be easy: observers could simply respond to the presence of a long or a short bar. But if the fragments were somehow linked and no longer readily available, this would cause slower, more effortful search.

Four conditions were used. Condition 4A was designed to check that rapid search could be based on the isolated fragments. Search here was quite fast (Fig. 4A): rates were 6 ms/item for target-present trials and 4 ms/item for target-absent trials.

Items in Condition 4B were the same as in Condition 4A except that the gaps between corresponding fragments were now occupied by flat 2-D hexagonal patches. This slowed down search considerably (29 and 55 ms/item). Evidently, the presence of an occluder triggered the completion process, causing the bar segments of the targets and distractors to form structures highly similar to each other.

In Condition 4C the 2-D patches were replaced




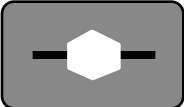

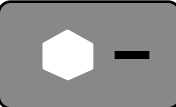
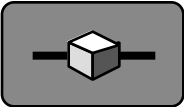
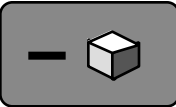
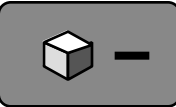





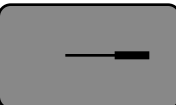
Cond.	Targets	Distractors		Rate (ms/item)	
				Present	Absent
5A				32	60
5B				17	25
5C				10	16
5D				5	12
5E				20	47

Fig.5. Search items and results of Experiment 5 (Speedup vs. Slowdown). Conditions 5A-5C are *distinct-object* tasks corresponding directly to Conditions 4A-4C. Here, search is *faster* when completion is stronger, showing that completion itself does not slow search down. When target fragments are connected by a thin line (Condition 5D) this joining is evident in the rapid search rates. Adding a thin line to the distractor fragments (Condition 5E) slows search somewhat, presumably because of the greater similarity in the overall lengths of target and distractor items.

with drawings of 3-D blocks; search now slowed down even more (46 and 86 ms/item). In contrast, simply connecting the fragments by a solid line (Condition 4D) caused search to speed up again (7 and 15 ms/item). This indicated that the shapes of the connected fragments could be accessed in a way not possible for fragments separated by an occluder.

Thus, these results show that search slows down when occluders are placed between fragments separated in the image. Such a slowdown is not simply due to the removal of edges caused by occlusion (as in Experiments 1-3), for the segments themselves—either isolated or connected by a line—are sufficiently distinct to support rapid search. It is not caused by the simple presence of occlusion, for search can be quite fast under such conditions (cf. Condition 2D). Instead, it appears that preemption takes place in the segments that have been linked together, with only certain aspects of the completed structure being rapidly accessible.

4.2 Experiment 5: speedup vs. slowdown

In previous experiments completion always acted to slow search down. This raises the possibility that the effects found in Experiment 4 might not have been entirely due to completion but might have involved other, unrelated factors. For example, the slowdown in Conditions 4B-4C might have been caused by the introduction of three-dimensionality in the occluders,

which might have created sufficient "noise" to interfere with the search process (e.g., Duncan & Humphreys, 1989; Treisman & Gormican, 1988).

To examine this possibility, Experiment 5 used a set of *distinct-object* conditions similar to those of Experiment 4, but with targets differing in overall length rather than fragment length (Fig. 5). Fragments were the same in both targets and distractors, but distractors contained only a single bar, separated from the occluders by 0.5° . Observers were asked to search for targets based on overall configuration. If fragments could be joined and no slowdown factors operated, search should be easy: observers could simply look for a long bar. Indeed, since Conditions 5A-5C each have a distinct-fragment counterpart in Experiment 4, the pattern of speedup should be an exact reversal of that found in Conditions 4A-4C. But if the fragments could not be joined, or if slowdown factors were dominant, a different pattern would emerge.

The distinct-object conditions are shown in Fig. 5. Condition 5A tested the isolated bar fragments. Search here was relatively slow (32 and 60 ms/item), indicating that observers could not rapidly join the target fragments.

In Condition 5B a flat 2-D occluder was added to targets and distractors. This caused a significant speedup in search (17 and 25 ms/item). Replacing the flat occluders by drawings of 3-D cubes in Condition 5C caused search to speed up even more (10 and 16

ms/item). This reversal of the pattern found in Conditions 4A-4C shows that the slowdown there was not due to any interference caused by the 3-D nature of the occluders, but rather, was due to a strengthening of preemption.

In Condition 5D thin lines were attached to the target fragments to connect them into a single contiguous item (Fig. 5D). As might be expected from the difference in overall length, search was now quite fast (5 and 11 ms/item); adding thin lines to the distractors (Condition 5E) caused it to slow down (20 and 47 ms/item). This pattern of results shows that the thin lines do join with the isolated fragments, and that this joining is at least as effective as the linking triggered by occlusion. However, preemption does not take place in visually joined structures, as Condition 4D shows.

The results from Experiments 4 and 5 indicate that rapid completion does indeed link separated fragments, and that the completed structures — for better or worse—preempt their constituent fragments. Interestingly, completion appears to be a graded phenomenon, its strength depending on the type of occluder: in agreement with the results of Weisstein et al. (1972), completion effects were found to be stronger for images of 3-D occluders than for 2-D occluders.

4.3. Experiment 6: Linkage of separated fragments

Although Experiments 4 and 5 indicate that separated fragments are linked together, they do not shed much light on the nature of this linkage. The results are consistent with the linking of the bar fragments across occluded space to form structures quite separate from the occluders. But it is also possible that the bar fragments and occluders in each item were simply concatenated into an undifferentiated agglomeration. If so, search in Experiments 4 and 5 might have been governed not by the completed bars, but by these agglomerations.

To examine this possibility, Condition 6A used the items of Condition 5C, but with one of the bars in the target removed so that it no longer had a matching fragment. Since the bar fragments were about as wide as the occluders, the (concatenated) targets would be twice as long as the (non-concatenated) distractor fragments, a difference relatively easy to detect (Rensink & Enns, 1995; Treisman & Gormican, 1988). Indeed, when matching fragments exist (Condition 5C), search is relatively fast (10 and 16 ms/item). However, for the unmatched fragments here, search was much slower (29 and 54 ms/item), a rate not significantly different from that for the isolated fragments themselves (Condition 5A).

Another test was carried out in Condition 6B, where items were concatenations of the fragments used in Experiment 4. These items had smaller differences in relative overall length, so that if search was based on concatenated structure, it would be slower than in Condition 6A. However, search was significantly

faster (11 and 16 ms/item); indeed, rates approached those for the fragments alone (Condition 4A), indicating that search was based on the lengths of the bar fragments alone. Evidently, access to individual fragments is possible here in a way that was not possible in Experiment 4. The absence of matching fragments appears to have created a situation similar to Experiments 1-3, where preemption of fragment structure did not occur.

These results indicate that a bar fragment in an item remains effectively isolated unless a matching fragment exists. The separation of the bars from the occluders is likely effected by the presence of T-junctions, which can split apart otherwise connected lines and regions (Enns & Rensink, 1991, 1992). Condition 6A shows that this separation cannot be overcome, even to help search. Rather, the bar fragments remain separate from the occluders at all times. They can only be linked to each other, and this only when an occluder is in the space between them. Note that while unlinked, the visible structure of a fragment is rapidly accessible (Experiments 1-3; Condition 6B); after linking, this is no longer possible (Conditions 4B, 4C).

5. Experiments 7-9: rapid completion vs. rapid grouping

The linking found in Experiments 4-6 raises the issue of whether the completion effects are simply due to grouping. Several types of rapid grouping are known to exist in early vision (e.g., Rensink & Enns, 1995), and some of these have characteristics similar to those of the process found here (e.g., preemption of constituent fragments, reluctance to posit new visual elements). But the findings that occlusion edges are removed and that occluders are needed to trigger linking suggest that rapid completion may involve occlusion-specific mechanisms not found in any general grouping process. A better understanding of the relation between completion and grouping may be attained by further exploration of the ways in which completion depends on occlusion.

5.1. Experiment 7: Effect of gaps

Rapid grouping is known to take place across gaps (e.g., Rensink & Enns, 1995). What about rapid completion? If completion is concerned with linking across occluded spaces only, it should fail whenever gaps are present, since these cannot be accounted for by any visible occluder.

To determine if this is indeed the case, Condition 7A used the distinct-fragment items of Condition 4C, but with 0.4° gaps introduced between the bar fragments and the occluding block (Fig. 7A). Search sped up remarkably (10 and 15 ms/item). In Condition 7B the gaps were reduced to 0.2°; search again was relatively fast (9 and 14 ms/item). In both conditions search was considerably faster than for the no-gap stimuli of Condition 4C, indicating that the

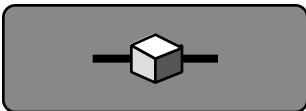
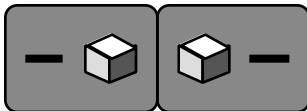
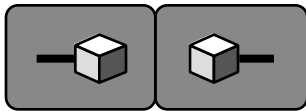
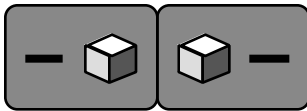
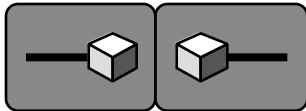
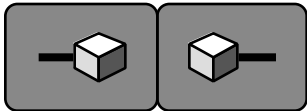
Cond.	Targets	Distractors	Rate (ms/item)	
			Present	Absent
5C			10	16
6A			29	54
6B			11	16

Fig. 6. Search items and results of Experiment 6 (Linkage of Separated Fragments). For purposes of comparison, an upper row has been added showing the slopes of Condition 5C. When one of the matching target fragments is removed (Condition 6A), search slows down considerably, more than would be expected if the overall length of the concatenated fragments could still be used. When the bar fragments in the items contact the occluders (Condition 6B), search is still governed by the length of the fragments rather than by overall concatenated length. This shows that fragments are linked only to matching fragments, and that preemption of fragment length occurs only in linked structures.

distinct fragments could once again be accessed. Evidently, completion had failed to occur under these conditions.

The introduction of gaps in Conditions 7A and 7B caused the bar fragments to be slightly displaced from each other. To check whether this displacement itself might have caused the speedup, Condition 7C used the same items as in Condition 4C, but with gaps created by displacing the fragments vertically (Fig. 7C). However, the task remained quite easy (5 ms/item for target-present trials; 8 ms/item for target-absent trials). In fact, it was as easy as when the bar fragments appeared alone (Condition 4A).

Taken together, these results show that rapid completion is highly sensitive to the existence of gaps, with linking occurring only across a completely occluded space. Such sensitivity is not characteristic of rapid grouping. As such, these results support the position that rapid completion is not simply an expression of general grouping, but involves separate processes specialized to deal with the presence of occluders.

5.2. Experiment 8: Restoration of linked structures

Although Experiments 1-3 showed that new visual elements are not posited in the occluded parts of contiguous fragments, this might still happen for linked structures. Suggestive evidence on this point is the finding that fragments in linked structures are preempted, whereas isolated fragments are not.

Condition 8A used the items of Condition 5C, except that matching bar fragments were added to the distractors and then displaced vertically (Fig. 8A). If a positing of visual elements took place within linked fragments, the rapid search found in Condition 5C should be maintained, since completed targets would have distinctive shapes as well as considerably more "black material" on their surfaces. Search, however, became far slower (45 and 86 ms/item).

Such a slow search indicates that the targets contained no distinctive features. Evidently, the distractor fragments were grouped into items with lengths similar to those of the targets—indeed, connecting the distractor bars with a thin line (Condition 8B) did not cause any significant change in speed (45 and 90 ms/item). It is not entirely clear why a similar grouping did not speed up search in Condition 5A. One possibility is that since rapid grouping has distance limits (Rensink & Enns, 1995), it may have been that without the occluders those fragments were simply too far apart.

In any event, the lack of a distinctive feature in Conditions 8A and 8B argues against any significant positing of visual elements in linked structures. Consequently, the slowdown found for 3-D occluders is probably not related to the filling-in found by Weisstein et al. (1972). Rather, the 3-D occluders appear to simply generate a stronger linkage of the occluded fragments.

5.3. Experiment 9: Effect of occluders

To better understand how the presence of occluders might trigger rapid completion, the final experiment looked at completion effects in line drawings (Fig. 9). This allowed an examination of several aspects of line configuration (including free endings and T-junctions) in order to determine some of the critical properties involved. All conditions were distinct-fragment tests, where the task was to indicate the presence/absence of a long fragment against a background of shorter ones.

Condition 9A contained no explicit occluders. Instead, targets and distractors were composed of isolated fragments (Fig. 9A). As might be expected from Condition 4A, completion here did not occur; observers could easily find the distinctive fragments of the targets amid those of the distractors (3 and 2 ms/item).

Cond.	Targets		Distractors	Rate (ms/item)	
				Present	Absent
4C				46	86

7A				10	15
7B				9	14
7C				5	8

Fig.7. Search items and results of Experiment 7 (Effect of Gaps). For purposes of comparison, an upper row has been added showing the slopes of Condition 4C. The presence of any kind of gap between bar fragments and occluders causes search to greatly speed up, indicating that preemption (and therefore completion) has failed.

In Condition 9B an occluding two-dimensional figure was placed into the gap between the fragments. Similar to what occurred with the solid fragments of Experiment 4, this caused search to slow down (13 and 21 ms/item). Although not quite as slow as for the solid fragments, search did slow down by a factor of more than 4, indicating that some degree of completion had occurred.

Condition 9C examined whether the occluder needed to be a complete figure. Here, the ends of the occluding rectangle were removed, while the remaining segments continued to contact the bar fragments (Fig. 9). Speed remained essentially the same as in Condition 9B (13 and 24 ms/item), indicating that the critical factor was the information at the junctions.

Condition 9D removed the occluding lines of the T-junctions so that only free line endings remained. Speed still remained similar to that of Conditions 9B and 9C (11 and 25 ms/item), showing there is little need for the occluding line in the T-junction; evidently, the stems of T-junctions are treated much the same as free line endings. Note that the stems of L-junctions do not have a similar equivalence (Condition 9A).

In Condition 9E the free line endings were moved so that the gaps were no longer orthogonal to the edges of the occluded figure (Fig. 5). Search now slowed down even further (22 and 47 ms/item), indicating that the orientation of the "virtual gap" was important. Condition 9F removed the line endings by joining them with explicit line segments. Search again became rapid (5 and 7 ms/item)—a rate not significantly different from that of Condition 9A. Thus, the slowdown in Condition 9E was not primarily due to the different edge lengths present, but to a strengthening of completion.

A comparison of Conditions 9E-9F with 9A-9D

indicates that completion is relatively weak (although still operative) for orthogonal gaps. One possibility is that illusory contours might spread out orthogonally from the line endings (e.g., Grossberg & Mingolla, 1985; von der Heydt & Peterhans, 1989) and form virtual lines that weaken linking (cf. Condition 9A).

In any event, these results clearly show that free line endings are linked together when the segments involved are collinear, and that the strength of these linkages can be modulated by interactions with nearby structures. A similar pattern appears to hold for the stems of T-junctions but not the stems of L-junctions. This sensitivity to junction type is consistent with previous research on low-level completion under focal attention (Nakayama et al., 1989) and with the preattentive recovery of 3-D slant from line drawings (Enns & Rensink, 1991; Rensink, 1992).

Thus, it appears that rapid completion is triggered—at least for the line drawings examined here—by the free line endings caused by the presence of occluders. It is tempting to speculate that a similar joining of "free surface endings" generated by the removal of occlusion edges may underlie the linking of solid fragments. A deeper understanding of these matters, however, must await future experiments.

6. General discussion

The experiments presented here provide evidence of a process that uses monocular pictorial cues to rapidly complete partially-occluded objects. Search for an easily-found fragment becomes difficult when it is completed into a shape similar to that of others in the display; conversely, search for a hard-to-find fragment becomes easy when it is completed into a distinctive shape. These findings extend the results of earlier studies based on priming and stereopsis, showing that

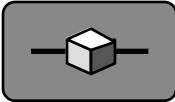
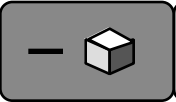
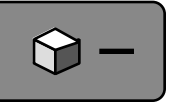
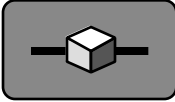
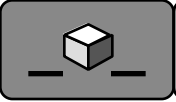
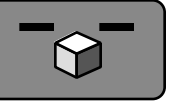
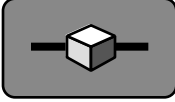
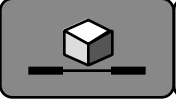
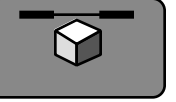
Cond.	Targets	Distractors	Rate (ms/item)	
			Present	Absent
5C		 	10	16
8A		 	45	86
8B		 	45	90

Fig. 8. Search items and results of Experiment 8 (Completion vs. Filling-In). For purposes of comparison, an upper row has been added showing the slopes of Condition 5C. When matching fragments are added to the isolated fragments of Condition 5C (Condition 8A), search slows down greatly, indicating that grouping (although not completion) occurs between fragments separated by gaps. Condition 8B shows that this grouping is relatively strong, since adding a connecting line leaves search largely unaffected. The slow speeds for these conditions would not be expected if target bars had become structures with more surface material than distractor bars.

completion based on monocular cues can take place in the absence of focused attention, and that it can do so within the time spans generally associated with early visual processing. This process does not appear to posit new visual elements to restore the occluded parts of the image. It does, however, remove edges caused by occluders and links together matching fragments so strongly that the completed structures are *functionally* filled in, with subsequent processes unable to rapidly access the constituent parts.

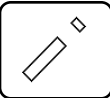
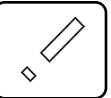
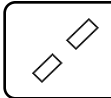

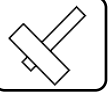

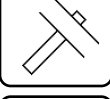





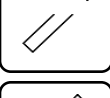



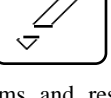
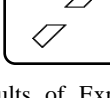
Cond.	Targets	Distractors	Rate (ms/item)	
			Present	Absent
9A		 	3	2
9B		 	13	21
9C		 	13	24
9D		 	11	25
9E		 	22	47
9F		 	5	7

Fig. 9. Search items and results of Experiment 9 (Effect of Occluders). Completion in line drawings depends on free line endings (or stems of T-junctions), and is strongest when these line endings are not orthogonal to each other.

In some ways, the existence of a rapid-completion process is not entirely unexpected. Recent experiments have shown that early vision not only registers image-based features such as orientation, length, and color, but can also recover several scene-based properties, such as shape from shading (Enns & Rensink, 1990; Ramachandran, 1988; Sun & Perona, 1996), slant from line drawings (Enns & Rensink, 1991; Epstein, Babler, & Bownds, 1992), slant from binocular stereopsis (Holliday & Braddick, 1991), and shadows from luminance patterns (Rensink & Cavanagh, 1993). The completion of sensor-gap interruptions is believed to take place in V1, the earliest cortical stage of visual processing (Fiorani et al., 1992). Several types of rapid grouping have been discovered at early visual levels (Elder & Zucker, 1993; Rensink & Enns, 1995), as well as illusory contours (Davis & Driver, 1994), all presumably to help compensate for losses introduced by sensor coupling.

The results obtained here, however, show that early vision can also compensate for losses caused by occlusion. Such a process must be of considerable sophistication: to compensate for occlusion on the basis of monocular cues, not only must items be linked together, but a simultaneous determination must be made about which items correspond to the occluding objects and which to the occluded objects. As such, the existence of rapid completion shows that the "quick and dirty" processes found in early vision can be used for visual problems of considerable structural complexity.

6.1. Nature of the completion process

The experiments here have focused on establishing the existence of rapid completion rather than determining the particular rules of its operation. However, the results do provide at least some indication of the nature of this process and how it operates. To be-

gin with, rapid completion appears to remove any edge caused by occlusion (Experiment 1). This need not be a literal removal—the edge could remain visible, but simply be assigned to the occluding rather than to the occluded structure (Nakayama & Shimojo, 1990). Such a removal does not leave the occluded figure completely indeterminate in shape, since rapid search can still be based on its visible parts (Experiment 2). Evidently, this indeterminacy is local, confined only to the edge segments that have been reassigned.

Next, the linking between separated fragments appears to be sensitive to the nature of the occluders. Completion appears to be a graded phenomenon, with completion stronger for 3-D than for 2-D occluders (Experiments 4 and 5), and for gaps not orthogonal to the edges of the occluded object (Experiment 9). Note that in both cases, the stronger effects are found when there is less likelihood that the occluding structure is part of the occluded one: a 3-D occluder is unlikely to be part of the same object as the 2-D bars, while an orthogonal angle might be considered a "nongeneric" orientation unlikely to result from the arbitrary placement of an occluding object (cf. Nakayama & Shimojo, 1990).

Finally, there does not appear to be much visual filling-in of the completed structures, either in the way of surface "stuff" or of boundaries (Experiments 3 and 8). This is consistent with the results of Treisman & DeSchepper (1996), who found that negative priming is sensitive to the mosaic rather than the completed form of individual shapes. Our findings indicate that attentional access is based on the "exterior" shape of the completed structure: when a visible element is isolated, it effectively is the completed structure, and so can be rapidly accessed (Experiment 2); if linked to a corresponding element, however, rapid access is possible only for the shape of the linked structure, and so rapid access to that element is lost (Experiments 4 and 5). From the point of view of object recognition, this could be considered the result of a filling-in process. But such filling-in is purely functional: although the completed structures are the effective units of attentional access, it is their visible elements alone that affect (preattentive) estimates of visual similarity. In this regard, it is worth noting that structures completed in the absence of focused attention can also serve as conduits of attention (Mattingley, Davis, & Driver, 1997). Evidently, functional filling-in applies not only to attentional access, but to attentional transmission as well.

6.2. *Early vision and visual recognition*

The existence of a rapid-completion process clarifies several issues about the nature and purpose of early vision. It is widely held that early vision should provide a description of the world that facilitates the recognition of objects and events (e.g., Marr, 1982). But what exactly should be described? And how do these descriptions facilitate the recognition process? It has been proposed that early-level descriptions be

restricted to relatively unstructured surface properties, leaving determination of the more abstract properties of individual objects to higher levels (e.g., Barrow & Tenenbaum, 1978). In this view, completion (including the visual filling-in of unstructured properties) might be expected for the relatively invariant interruptions caused by sensor gaps. But it would not be expected for the more complex, ever-changing interruptions caused by occluding objects. Instead, recognition would be carried out via higher-level processes that side-step much of the occlusion problem, e.g., by matching projections of known 3-D models against uncompleted fragments (e.g., Lowe, 1985).

However, the results here indicate that the visual completion of occluded structure is—in spite of its computational complexity—a task important enough to be carried out at early levels. This has strong implications for the role of early vision: not only should it attempt to recover unstructured scene-based properties such as slants and shadows, but it should also provide more sophisticated structures, or "proto-objects" (Rensink & Enns, 1995) that correspond to objects in the world.

What is the nature of these proto-objects? Little is known about them as yet, but a comparison of the results obtained here with those from other studies (e.g., Donnelly, Humphreys, & Riddoch, 1991; Enns & Rensink, 1991; Rensink & Enns, 1995), suggests that they are assemblies of edge fragments with a considerable degree of internal coherence, capable of extending over several degrees of visual angle, and of overlapping each other without confusion of identity. Their formation is triggered by local cues (e.g., co-termination), with their separation into distinct structures beginning at much the same time, also based on local cues (e.g., T-junctions). These local structures may then be grouped into proto-objects according to the more global criteria considered in many computational studies, such as overall smoothness (Sha'ashua & Ullman, 1988) or convexity (Jacobs, 1996). In any event, it is these completed proto-objects—and not the isolated fragments—that become the basic units of the attentional access crucial to high-level recognition. Of course, these fragments can be attended to individually if required, but only at the cost of increased attentional effort.

Visual recognition clearly becomes much easier if based on complete shapes rather than interrupted fragments; if nothing else, it allows a greater amount of bottom-up indexing into higher-level memory, which greatly reduces the number of model-to-image matches that need to be made (Jacobs, 1992, 1996). The difficulty remains that early visual processes cannot be specialized for each particular object in the scene, and so cannot be optimal for all aspects of object recognition. But the world does apparently contain enough local structure that many aspects of occlusion can be rapidly compensated for, with the resulting descriptions reliable enough to serve as the primary bases for subsequent visual processing.

Acknowledgements

This research was supported by grants from NSERC Canada to J.T. Enns and R.J. Woodham. The authors are grateful to Leighton Duerre, Diana Ellis, and David Shore for their help in collecting data, to Bob Woodham for his support of R.R., and to Carol Yin and two anonymous reviewers for their comments on an earlier draft of this paper. Also, thanks to Adriane Seiffert for making us complete this paper.

Appendix A. Detailed description of experimental data

Mean correct response time (RT), standard errors of the mean (SE), and mean accuracy (%) in each of the nine experiments in this study.

Condition	Target Present display Size			Target Absent display Size			RT Slopes (ms/item)	
	2	8	14	2	8	14	P	A
1A	RT 521 SE 16 % 95	578 15 95	617 18 93	564 17 97	617 16 98	663 20 99	7.4	8.3
1B	RT 554 SE 22 % 98	827 51 90	980 50 85	648 23 98	1105 45 97	1428 79 99	35.6	66.0
1C	RT 541 SE 16 % 99	683 24 97	782 31 90	646 28 97	934 48 99	1187 79 99	20.0	45.2
1D	RT 599 SE 15 % 97	749 78 91	995 51 87	613 61 98	1271 71 98	1536 183 99	34.8	83.1
2A	RT 479 SE 14 % 99	526 12 98	546 9 94	511 13 98	559 14 95	603 17 96	5.6	7.6
2B	RT 599 SE 15 % 97	749 78 96	995 51 91	613 61 98	1271 71 99	1536 182 99	18.8	44.4
2C	RT 581 SE 17 % 96	752 33 91	886 32 81	647 23 98	1082 73 99	1244 92 98	25.4	50.0
2D	RT 497 SE 36 % 96	532 26 96	557 30 96	552 33 98	618 33 99	649 39 99	4.8	8.1
3A	RT 733 SE 28 % 98	1207 44 85	1499 65 85	823 33 98	1642 72 99	2272 113 98	63.9	120.9

3B	RT 572 SE 19 % 98	774 19 93	930 27 86	664 22 98	1040 42 98	1297 61 99	29.9	52.8
3C	RT 481 SE 7 % 99	537 13 97	556 10 96	529 11 95	634 12 99	702 23 99	6.3	14.6
4A	RT 448 SE 18 % 99	486 20 97	510 16 93	504 25 98	544 36 98	571 43 98	5.6	4.2
4B	RT 566 SE 27 % 98	783 54 91	909 62 86	639 29 96	1062 99 99	1301 122 98	28.6	55.2
4C	RT 581 SE 27 % 96	886 62 89	1127 106 82	634 27 98	1320 113 99	1668 142 92	45.5	86.2
4D	RT 506 SE 22 % 98	557 21 98	585 26 98	580 29 97	690 28 99	759 44 99	6.7	15.0
5A	RT 584 SE 31 % 98	779 32 95	963 51 90	565 25 99	944 65 95	1282 94 91	31.6	59.8
5B	RT 541 SE 19 % 97	654 27 94	742 38 90	577 23 98	729 43 99	883 59 98	16.7	25.5
5C	RT 532 SE 18 % 98	599 21 93	647 23 90	570 23 99	671 23 99	765 33 98	9.6	16.3
5D	RT 513 SE 23 % 98	555 26 98	570 28 97	572 35 98	657 43 99	712 47 99	4.8	11.7
5E	RT 622 SE 26 % 99	757 34 95	859 29 95	673 34 99	921 55 99	1238 79 99	19.8	47.1
6A	RT 644 SE 23 % 96	850 42 90	991 57 80	711 26 99	1060 96 98	1354 119 96	28.9	53.6
6B	RT 526 SE 16 % 95	600 25 92	653 21 92	570 27 99	676 31 98	763 34 97	10.6	16.1
7A	RT 532 SE 12 % 97	610 18 94	653 29 93	591 18 97	727 24 98	772 33 97	10.0	15.1
7B	RT 500 SE 17 % 97	578 21 94	612 23 90	562 23 97	702 36 99	728 54 99	9.3	13.9

7C	RT	468	506	530	524	570	617	5.2	7.8
	SE	14	16	13	22	19	22		
	%	99	97	93	98	98	98		
8A	RT	628	986	1174	685	1291	1715	45.4	86.1
	SE	14	27	36	13	55	80		
	%	98	87	82	99	99	99		
8B	RT	652	1044	1191	726	1372	1806	44.9	90.2
	SE	15	32	53	18	53	76		
	%	96	85	82	98	98	99		
9A	RT	454	483	494	494	517	516	3.3	1.8
	SE	20	19	19	20	27	23		
	%	98	96	95	95	97	99		
9B	RT	508	585	667	558	712	810	13.2	21.0
	SE	30	40	53	31	49	55		
	%	99	94	84	98	97	96		
9C	RT	549	649	705	601	824	886	13.1	23.8
	SE	29	38	49	33	85	87		
	%	98	93	87	97	98	97		
9D	RT	517	593	649	585	814	888	11.0	25.2
	SE	22	31	39	24	61	65		
	%	96	95	90	97	97	98		
9E	RT	608	799	869	680	1052	1249	21.8	47.4
	SE	39	72	74	41	82	106		
	%	98	87	77	99	98	99		
9F	RT	466	509	521	518	578	596	4.6	6.5
	SE	22	24	20	22	33	38		
	%	97	97	95	98	99	98		

References

- Barrow, H.G., & Tenenbaum, J.M. (1978). Recovering intrinsic scene characteristics from images. In A. Hanson and E. Riseman (Eds.), *Computer Vision Systems* (pp. 3-26). New York: Academic.
- Beck, J. (1982). Textural segmentation. In J. Beck (Ed.), *Organization and Representation in Perception* (pp. 285-317). Hillsdale, NJ: Erlbaum.
- Davis, G., & Driver, J. (1994). Parallel detection of Kanisza subjective figures in the human visual system. *Nature*, 371, 791-793.
- Donnelly, N., & Humphreys, G. W., & Riddoch, M. J. (1991). Parallel computation of primitive shape descriptions. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 561-570.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, 96, 433-458.
- Durgin, F., Tripathy, S.P., & Levi, D.M. (1995). On the filling in of the visual blind spot: some rules of thumb. *Perception*, 24, 827-840.
- Elder, J.H., & Zucker, S.W. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, 33, 981-991.
- Elder, J.H., & Zucker, S.W. (1996). Computing contour closure. *Proceedings of the 4th European Conference on Computer Vision*, vol. 1, 399-412.
- Enns, J.T. (1992). The nature of selectivity in early human vision. In B. Burns (Ed.), *Percepts, concepts, and categories* (pp. 39-74). Amsterdam: Elsevier.
- Enns, J.T., & Kingstone, A. (1995). Access to global and local properties in visual search for compound stimuli. *Psychological Science*, 6, 283-291.
- Enns, J.T., & Rensink, R.A. (1990). Influence of scene-based properties on visual search. *Science*, 247, 721-723.
- Enns, J.T., & Rensink, R.A. (1991). Preattentive recovery of three-dimensional orientation from line drawings. *Psychological Review*, 98, 101-118.
- Enns, J.T., & Rensink, R.A. (1992). A model for the rapid interpretation of line drawings in early vision. In D. Brogan (Ed.), *Visual search II* (pp. 73-89). London: Taylor & Francis.
- Epstein, W., Babler, T., & Bowns, S. (1992). Attentional demands of processing shape in three-dimensional space: Evidence from visual search and precuing paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 503-511.
- Fiorani, M., Jr., Rosa, M.G., Gattass, R., & Rocha-Miranda, C.E. (1992). Dynamic surroundings of receptive fields in primate striate cortex: A physiological basis for perceptual completion? *Proceedings of the National Academy of Sciences*, 89, 8547-8851.
- Gerbino, W., & Salmaso, D. (1987). The effect of amodal completion on visual matching. *Acta Psychologica*, 65, 25-46.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, 92, 173-211.
- He, Z.J., & Nakayama, K. (1992). Surfaces versus features in visual search. *Nature*, 359, 231-233.
- Helmholtz, H. von. (1867/1967). *Treatise on physiological optics* (Vol. 3). J. P. C. Southall (Ed. and Trans.). New York: Dover.
- Holliday, I.E., & Braddick, O.J. (1991). Pre-attentive detection of a target defined by stereoscopic slant. *Perception*, 20, 355-362.
- Jacobs, D.W. (1992). *Recognizing 3-D objects using 2-D images*. Ph.D. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA.
- Jacobs, D.W. (1996). Robust and efficient detection of salient convex groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18, 23-37.
- Julesz, B. (1984). A brief outline of the texton theory of human vision. *Trends in Neuroscience*, 7, 41-45.
- Kellman, P. & Shipley, T. (1991) A theory of visual interpolation in object perception. *Cognitive Psychology* 23: 141-221.
- Lowe, D.G. (1985). *Perceptual organization and visual recognition*. Dordrecht: Kluwer Academic.
- Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman.
- Mattingley, J.B., Davis, G., & Driver, J. (1997). Preattentive filling-in of visual surfaces in parietal extinction. *Nature*, 275, 671-674.
- Nakayama, K., & Shimojo, S. (1990). Toward a Neural Understanding of Visual Surface Representation. *Cold Spring Harbor Symposia on Quantitative Biology*, Vol. LV, 911-924.

- Nakayama, K., Shimojo, S., & Silverman, G.H. (1989). Stereoscopic depth: Its relation to image fragmentation, grouping, and the recognition of occluded objects. *Perception*, *18*, 55-68.
- Pomerantz, J.R., & Kubovy, M. (1986). Theoretical approaches to perceptual organization. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance*, vol. II (ch. 36). New York: John Wiley & Sons.
- Ramachandran, V.S. (1988). Perceiving shape from shading. *Scientific American*, *259*, 76-83.
- Ramachandran, V.S. (1992). Blind spots. *Scientific American*, *266*, 86-91.
- Ramachandran, V.S., & Gregory, R. L. (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature*, *350*, 699-702.
- Rensink, R.A. (1992). The rapid recovery of three-dimensional orientation from line drawings. Ph.D. Thesis (also Technical Report 92-25), Department of Computer Science, University of British Columbia, Vancouver, BC, Canada.
- Rensink, R.A., & Cavanagh, P. (1993). Processing of shadows at preattentive levels. *Investigative Ophthalmology & Visual Science*, *34*, 1288.
- Rensink, R.A., & Enns, J.T. (1995). Preemption effects in visual search: Evidence for low-level grouping. *Psychological Review*, *102*, 101-130.
- Sekuler, A.B., & Palmer, S.E. (1992). Perception of partly occluded objects: A micro-genetic analysis. *Journal of Experimental Psychology: General*, *121*, 95-111.
- Sha'ashua, A., & Ullman, S. (1988). Structural saliency: The detection of globally salient structures using a locally connected network. *Proceedings of the IEEE Conference on Computer Vision*, 321-327.
- Sun, J.Y., & Perona, P. (1996). Preattentive Perception of Elementary Three-Dimensional Shapes. *Vision Research*, *36*: 2515-2529.
- Treisman, A. (1986). Features and objects in visual processing. *Scientific American*, *255*, 106-115.
- Treisman, A., & DeSchepper, B. (1996). Object Tokens, Attention, and Visual Memory. In T. Inui and J.L. McLelland (Eds.), *Attention and Performance XVI* (pp. 15-46). Cambridge, Mass.: MIT Press.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, *95*, 15-48.
- Weisstein, N., Montalvo, F.S., & Ozog, G. (1972) Differential adaptation to gratings blocked by cubes and gratings blocked by hexagons: A test of the neural symbolic activity hypothesis. *Psychonomic Science*, *27*, 89-91.
- Williams, L.R., & Hanson, A.R. (1996). Perceptual completion of occluded surfaces. *Computer Vision and Image Understanding*, *64*, 1-20.
- von der Heydt, R., & Peterhans, E. (1989). Cortical contour mechanisms and geometrical illusions. In D. Lam and C. Gilbert (Eds.), *Neural Mechanisms of Visual Perception* (pp. 157-170). New York: Academic.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1988). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception & Performance*, *15*, 419-433.
- Wolfe, J.M., Friedman-Hill, S.R., & Bilsky, A.B. (1994). Parallel processing of part-whole information in visual search tasks. *Perception & Psychophysics*, *55*, 537-550.