

Curve tracing: A possible basic operation in the perception of spatial relations

PIERRE JOLICOEUR

University of Waterloo, Waterloo, Ontario, Canada

SHIMON ULLMAN

Massachusetts Institute of Technology, Cambridge, Massachusetts

and

MARILYNN MACKAY

University of Saskatchewan, Saskatoon, Saskatchewan, Canada

The two experiments in this study suggest that fast internal tracing of curves is employed by the visual system in the perception of certain shape properties and spatial relations. The experimental task in the first experiment was to determine, as rapidly as possible, whether two Xs lay on the same curve or on different curves in a visual display. Mean response time for "same" responses increased monotonically with increasing distance along the curve between the Xs. The task in the second experiment was to decide either that a curve joining two Xs was unbroken or that the curve had a gap. Decision times again increased as the length of the curve joining the Xs was increased. The results of both experiments suggest that people can trace curves in a visual display internally at high speed (the average rate of tracing was about 40° of visual angle per second). Curve tracing may be an important visual process used to integrate information from different parts of a visual display.

The perception of shape properties and of spatial relations plays an important role in the visual recognition of objects, in visually guided manipulation, and in navigation in the environment. It is not surprising, therefore, that humans can quickly and effortlessly perceive a large variety of spatial relations among items in a visual scene. Even a simple display such as Figure 1 contains a large number of spatial relations that can be perceived immediately when required. For example, (1) Is there a closed curve in the display? (2) How many dots are there? (3) Is there a dot inside the closed curve? (4) Is there an X lying on a curve? (5) Are there two Xs lying on a common curve? Human observers can answer such questions accurately and immediately by merely looking at the figure.

For an observer to decide that a particular spatial relation exists between two elements in a display, he or she must somehow consider the elements in relation to each other and compute the relation between them. The human visual system can establish a variety of shape properties

and spatial relations with intriguing proficiency. The perception of spatial properties and relations that are complex from a computational standpoint nevertheless often appear to us as immediate and effortless. When we perceive such relations, we usually do not know how we arrive at our decision—we simply see the relation. Therefore, systematic empirical study is required to unravel the processes underlying the perception of spatial relations. These processes are still largely unknown even in the case of seemingly simple relations.

The general approach we propose to this difficult problem is to assume that the visual system incorporates a set of basic operations (elementary processes) that can be concatenated to yield answers to more complex questions regarding spatial relations (Ullman, 1984). To illustrate this approach, we shall consider a simple example.

An Example: Two Xs on a Curve

Consider the problem of perceiving one of the relations in Figure 1, namely, that two Xs lie on a common curve. What processes might underlie the perception of such a relation? Clearly, it does not seem plausible to postulate the existence of a dense array of "two-Xs-on-a-curve" detectors specialized for this task. The problem can be solved, however, by applying an appropriate sequence of operations to the image. For example, one could use the following procedure: First, locate an X in the visual field. Next, check whether this X intersects a curve. If the X

We thank Stephen Kosslyn for extensive and useful comments on an earlier draft of the manuscript, and Merrill Hiscock and Thomas Wishart for providing equipment used in the experiments reported in this paper. We also gratefully acknowledge the support of the Center for Cognitive Science at the Massachusetts Institute of Technology. The research was supported by University of Saskatchewan internal research grants 3-773-936 and 1-400-012-790-R, and NSERC Grant A2579 awarded to Pierre Jolicoeur. Requests for reprints should be sent to P. Jolicoeur, Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.




Figure 1. A simple visual display illustrating the variety of possible spatial relations between relatively few elements.

does intersect a curve, the location is marked for future reference. Then the curve can be traced (in one or two directions) until another X is found—in which case the display contains a curve with two Xs. If another X is not found, the tracing is aborted upon reaching a curve termination (for an open curve) or a previously marked location (for a closed curve).

Several elementary processes are necessary to achieve the above computation: processes that locate Xs in a display, determine that an X lies on a curve, mark locations in a display, track curves, detect curve terminations, and recognize previously marked locations. In addition, miscellaneous decision processes are required to evaluate the outcome of the application of the elementary processes and to initiate appropriate action depending on the outcome. Taken together, the ordered application of elementary processes and the associated decision processes constitute a *visual routine*. In the above example, we described a possible visual routine (not necessarily unique) that would yield an answer to the two-Xs-on-a-curve problem for a relatively general set of visual displays.

Summary: Two Hypotheses

The above discussion illustrates two hypotheses we wish to put forward, one general, the other more specific. The general hypothesis is that the processes involved in the perception of a variety of abstract shape properties and spatial relations can be described in terms of a finite set of basic operations. Using different combinations of these operations, the visual system can assemble different processes, or visual routines, to extract an unbounded va-

riety of properties and relations. This is a general hypothesis that cannot be tested directly in a single experiment, but that may be supported by a collection of experiments if the same basic operations can be identified in the processes subserving the perception of various abstract spatial properties and relations.

The specific hypothesis is that a curve-tracing visual routine may be used to determine whether two symbols lie on a common curve or on different curves. The tracing routine in the example acts to bring together information from separate locations belonging to the same curve. The individual Xs could, in principle, be detected by local processes operating independently on different parts of the image. A more global operation is required to integrate this local information and establish that the two Xs lie on a common curve. We propose the possible use of curve tracing for this task. Such a tracing process is probably not the only possible method for integrating information from different parts of the same figure (see Minsky & Papert, 1969); however, curve tracing is a plausible mechanism for integrating information that pertains to the same curve or boundary.

In the following sections we describe a set of experiments that examined the possible use of curve tracing in two specific tasks. In the General Discussion we discuss the applicability of this process more generally to the perception of shape properties and spatial relations.

EXPERIMENT 1

The task in this experiment was to decide, as rapidly as possible, whether two Xs in a visual display were on the same curve or on different curves. The main manipulation consisted of systematically varying the distance along the curve between the Xs. We expected that the task would require that one of the curves in the display be scanned or traced and that longer distances to be traced would result in longer response times. Furthermore, if subjects in fact took longer to answer questions associated with greater curve distances, all else being equal, we would have evidence for the existence of a curve-tracing operation.

Each display in the experiment contained two separate curves and two Xs, and both Xs intersected a curve. Thus, subjects could not respond simply by noting that an X intersected a curve, or that there was only one X—all displays contained two Xs and each X was located on a curve. In half of the trials the two Xs were on the same curve; in the other half, the two Xs were on different curves. When the Xs were on the same curve, the distance along the curve separating the Xs was varied systematically. If subjects can trace arbitrary curves or contours, then they will take longer to solve problems that require tracing longer distances along the curve.

To ensure that longer response times reflect curve tracing, however, longer distances should not be associated with other factors that could cause longer response times. For example, distance to be scanned should not be cor-




Figure 2. One of the basic patterns used in this study. The dots along the curves indicate the locations of the second X used to create a set of displays from this basic pattern. There were no dots in the actual displays.

related with retinal eccentricity (distance from the fovea). In all displays used in this experiment, we held constant all other relevant factors. One X was always at the center of the display and the other was at a constant retinal eccentricity in all displays and for all curve distances. Thus, the two Xs in all displays were separated by the same physical distance. Furthermore, curve distance was varied within each pattern of curves, which ensured that curve distance was not confounded with different patterns of curves. Therefore, if greater distances along the curves are associated with longer response times, we will have good evidence that subjects used a tracing process to perform the task.

In addition to the above manipulations, the exposure duration of the stimuli was also varied. The two dura-

tions used were 250 and 2,500 msec. The short duration was included to discover whether subjects could trace rapidly over an arbitrary curve in the absence of eye movements, which would suggest an internal (cognitive) tracing mechanism. The longer duration was included as a control condition in which performance could be examined without possible complications from the use of a brief stimulus exposure (e.g., tracing over a fading representation).

Method

Materials. The stimuli consisted of figures each composed of two nonintersecting curves and two Xs. Figure 2 illustrates the construction of the displays. In all displays, one of the Xs was at the center of the display and intersected one of the curves. In half of the displays, the two Xs were on the same curve; in the other half, the two Xs were on different curves. Eight basic displays containing two curves with one X at the center were created; these displays differed in the shape of the two curves they contained. Figure 5 shows each of the eight basic patterns of curves (adapted for Experiment 2). Each of these basic displays was constructed such that eight locations for the second X could be identified; four of these locations were on the same curve as the central X and four were on the other curve. The dots in Figure 2 illustrate the eight possible locations for the noncentral X used for that particular display. The actual displays used did not contain any dots, only a second X, which was located at one of the marked locations. As can be seen in Figure 2, there were four possible distances along the curve between the two Xs when they were both on the same curve. The curve distance between Xs always increased by the same increment or unit; thus, the difference in curve distance between any two adjacent X locations was constant. One complete set of stimuli used in the experiment is displayed in Figure 3. The top row in Figure 3 contains stimuli in which the Xs lie on the same curve. The leftmost pattern in the top row has the smallest curve distance (1 unit), whereas the rightmost pattern has the largest curve distance (4 units). The second and third patterns have intermediate curve distances (2 and 3 units, respectively). In the stimuli displayed in the bottom row, the Xs lie on different curves. In this case, the distance from the noncentral X to one of the ends of the noncentral curve increases from the leftmost to the rightmost pattern in the figure.

In all displays, the physical distance between the central X and the other X was identical for all curve distances (see Figures 2, 3, and 5). Thus, any observed differences in response time for dis-




Figure 3. One complete set of displays used in Experiment 1. Top row shows patterns used in "same" trials with increasing curve distance from left to right. Bottom row shows patterns used in "different" trials and illustrates how the noncentral X was systematically moved relative to the ends of the noncentral curve.

plays with different curve distances between the Xs could not be due to differences in factors associated with retinal eccentricity or with the physical distance between the Xs.

The average size of the displays was 65.5×73 mm (displays ranged from 55×70 mm to 70×75 mm; see Figure 5 to appreciate relative sizes of the basic displays). Eight copies were made of each figure and each copy was marked with two 6×6 mm Xs, each positioned on a curve (see Figure 3). Stimuli were presented 800 mm from the subject and subtended an average visual angle of 5.9° . The four distances between the Xs along the same curve corresponded to 2.2° , 4.4° , 6.6° , and 8.8° of visual angle.

In the other four stimuli, the second X was located on the curve that did not contain the central X (see Figures 2 and 3). The physical distance between the two Xs was 25 mm (1.8° of visual angle) in all displays. The figures were mounted on tachistoscope cards with the central X positioned at the center of the card. Half of the stimuli for each basic figure were mounted at an orientation of 0° and half were rotated 90° before being mounted on the card, to reduce the likelihood that subjects would recognize the basic displays with repeated presentations.

One additional basic pattern of curves was created for use in practice trials.

Procedure. The task was to press one of two buttons depending on whether the two Xs were on the same curve or on different curves. The subjects were instructed to respond as rapidly as possible while keeping errors to a minimum. The instructions stressed accuracy because we were primarily interested in response times.

Stimuli were presented using a Gerbrands three-field tachistoscope (Model Number T-3B-1). Between trials, a dark fixation point on a white background was present in the center of the visual field. On each trial, the experimenter asked the subject if he or she was ready, and the trial was initiated when the subject responded affirmatively. Each trial consisted of the continuation of the fixation-point for 400 msec, followed by a blank field of 500 msec, followed by the display for either 250 or 2,500 msec, depending on the condition. Half of the subjects were in the 250-msec display condition and half were in the 2,500-msec condition.¹ The display was followed by a 2,000-msec blank postexposure field. The onset of the target display triggered a millisecond timer that was stopped when the subject pressed one of the response buttons. After each trial, the response and reaction time were recorded by the experimenter.

The subjects were informed that one X would always be at the fixation point in the display, and they were asked to maintain their positions at the viewer during and between trials to minimize variations in performance due to changes in dark adaptation level.

The test session began with practice trials (8 in the 250-msec condition and 16 in the 2,500-msec condition; see Note 1). The subjects then had an opportunity to ask questions and to have the instructions and procedure clarified. The subjects then completed two blocks of 64 trials, with a short rest period between the two blocks. Each block of 64 trials was composed of eight subblocks of 8 trials. Every subblock contained four displays in which the two Xs were on the same curve (one display with each of the four distances between the two Xs) and four in which the Xs were on different curves. Each subblock contained one figure created from one of the eight basic displays. The assignment of figures to subblocks was random within the above constraints. The order of trials within a subblock was randomized anew for each subject and the order of the subblocks was changed from subject to subject according to a Latin square design.

Half the subjects in each exposure condition saw the stimuli in the first block presented in the upright orientation and saw the stimuli in the second block rotated 180° . This procedure was used to minimize the chance that particular displays would be recognized during testing in the second block. Another result of this procedure is that direction of curve tracing (if required) and stimulation of different visual half fields were equated for all curve distances. For the other half of the subjects this order of presentation was reversed (i.e., Block 1 was inverted and Block 2 was upright).

Half of the subjects in each exposure condition (250 and 2,500 msec) responded "same" with the dominant hand and responded "different" with the nondominant hand. The other half had the opposite hand-response pairing. Subjects were informed when they made an error. Trials with incorrect responses were repeated at the end of the set in which the error was committed. Subjects were not informed that trials were repeated in this manner.

Subjects. The subjects were 16 undergraduates at the University of Saskatchewan. Eight subjects were assigned to the 250-msec exposure condition and 8 were assigned to the 2,500-msec exposure condition. All of the subjects had normal or corrected-to-normal vision, all were naive with respect to the hypotheses in the study, and no subject participated in the second experiment reported in this paper.

Results

Response times. The response time data were corrected for obvious outliers with the following rule: Any observation greater than the median plus two times the interquartile range (in a given cell—for each subject, response category, and curve distance) was replaced by the median of all the observations in that cell (see Tukey, 1977). This procedure resulted in the replacement of 3.9% of the observations.

The mean response time for each curve distance, response ("same"/"different"), and exposure duration is shown in Figure 4. Preliminary analyses revealed no important effects of testing blocks and this factor was not included in the analyses reported here. The response times for each subject, curve distance, and response were aver-

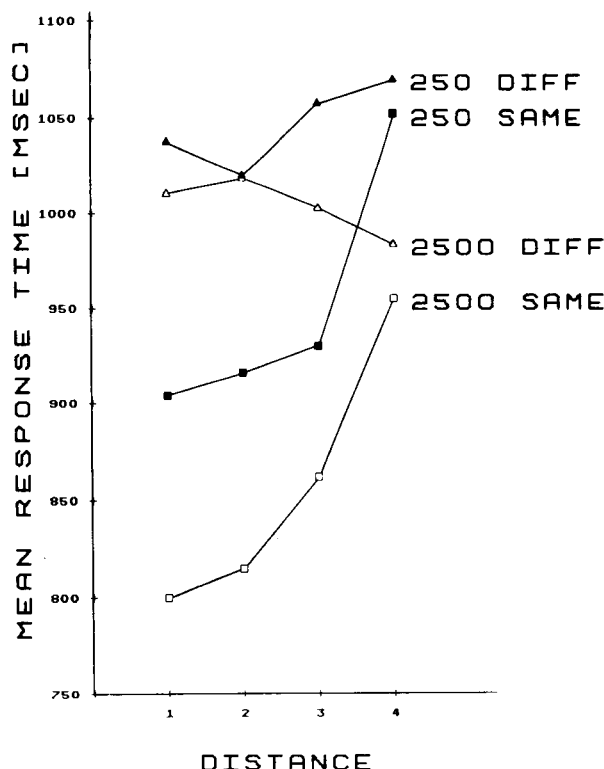


Figure 4. Results for Experiment 1. Filled symbols display the results for the 250-msec exposure condition; unfilled symbols display the results for the 2,500-msec condition. Squares show results for "same" trials; triangles show results for "different" trials.

aged and the means were submitted to a repeated measures analysis of variance. There were eight means for each subject, four for "same" responses and four for "different" responses. The seven degrees of freedom for these eight conditions were partitioned in the following manner: three for the distance effect within "same" responses, three for the distance effect within "different" responses, and one for the difference between "same" and "different" responses. The three degrees of freedom assigned to the distance factor were decomposed into the usual polynomial trends (linear, quadratic, and cubic).

As is evident in Figure 4, the mean response times for "same" responses increased monotonically with increasing distance between the two Xs in both exposure conditions, which resulted in a significant linear trend over distance for "same" responses [$F(1,14) = 19.05$, $p < .0007$, $MSe = 9.877 \times 10^3$]. This result suggests that subjects did trace the curves separating the two Xs in order to perform the task and that scanning longer distances took more time. The increase in response time for "same" trials over distance also had a significant nonlinear component, which was reflected in the quadratic trend over distance [$F(1,14) = 8.61$, $p < .011$, $MSe = 4.113 \times 10^3$].

"Same" responses were faster than "different" responses, on average [$F(1,14) = 13.88$, $p < .0025$, $MSe = 3.394 \times 10^4$]. Furthermore, there was no overall effect of distance on the speed of "different" responses ($F < 1$). As can be seen in Figure 4, however, there was a slight tendency for the trends in response times over distance in "different" trials to differ depending on exposure duration. This tendency was reflected by a marginal interaction between exposure duration and the linear trend over distance for "different" responses [$F(1,14) = 4.56$, $p < .051$, $MSe = 2.323 \times 10^3$]. This was the only significant effect attributable to exposure duration. The apparent difference between response times across the two display durations on "same" trials was not significant ($F < 1$).

Separate analyses were also performed on the data from each exposure duration. Both analyses revealed a strong linear component in the response times for "same" trials [$F(1,7) = 6.98$, $p < .034$, $MSe = 4.998 \times 10^3$, and $F(1,7) = 13.5$, $p < .008$, $MSe = 7.800 \times 10^3$ for the 250- and 2,500-msec conditions, respectively]. However, there were no trends over distance in "different" trials in either exposure duration ($p > .079$ in all cases).

Errors. The percentages of trials resulting in errors at each distance and in each exposure condition are presented in Table 1. The pattern of error rates does not suggest any speed-accuracy trade-offs. The error rate for each distance in "same" and "different" trials for each subject was submitted to an analysis of variance in which exposure duration was a between-subjects factor and "same"/"different" and distance were within-subject factors. There were more errors on "same" trials (7.0%) than on "different" trials (4.7%) [$F(1,14) = 6.32$, $p < .025$]. Errors tended to increase with increasing distance [$F(1,14) = 6.37$, $p < .025$], as tested by the linear trend across

Table 1
Mean Error Rate (Percent) for Each Exposure Duration, Response, and Distance Along the Curve in Experiment 1

	Distance			
	1	2	3	4
250 msec				
"same"	4.7	6.2	10.9	6.2
"different"	2.3	2.3	2.3	6.2
2,500 msec				
"same"	3.1	6.2	7.8	10.9
"different"	5.5	5.5	7.8	5.5

Note—For "same" trials, "Distance" represents the distance along the curve between the two Xs, and each unit distance corresponded with 2.2° of visual angle. For "different" trials, "Distance" represents the distance of the noncentral X to one of the ends of the curve not passing through the center of the display.

distance. The only other significant effect in the analysis was the exposure duration \times trial type ("same"/"different") \times distance interaction [$F(3,42) = 2.84$, $p < .05$]. The interaction between exposure duration and trial type with the linear trend across distance was only marginally significant, however [$F(1,14) = 3.18$, $p < .097$], and we will not consider these marginal effects further.

Discussion

The main results were as we would expect if the visual system has the capacity to trace curves in a visual display rapidly and internally: More time was required to solve the two-Xs-on-a-curve problem when longer distances along the curve separated the two Xs. Given that the two Xs were always separated by the same physical distance and were always at the same retinal eccentricity, these factors could not contribute to the pattern of means we observed. Because mean response time increased with curve distance in the 250-msec exposure duration, we have good evidence that curves can be traced without eye movements, which suggests the involvement of a relatively rapid internal perceptual process.

Response times were slower, on average, for "different" trials than for "same" trials (see Figure 4). This result is what would be expected if subjects traced the entire curve that was drawn through the fixation point in search of a second X and responded "different" only when they failed to find another X on the first curve.

The results suggest that subjects traced along the curve joining the two Xs in order to solve the problem and that tracing longer distances takes more time. Factors other than distance per se, however, may affect tracing time. For example, curves with greater curvature may require more time to trace per unit distance than less curved contours. In our displays, however, the average curvature was about the same for different curve segments beyond the first unit distance from the central X (see Figure 5). The effect of curvature on tracing rate is currently under empirical investigation.

The pattern of error rates is also consistent with the notion that subjects traced the central curve from the central X. Suppose that subjects traced the central curve from

the central X in search of the other X, and that they tended to respond "different" with increasing probability as the tracing process progressed along the curve without reaching another X. The error rate on "same" trials would increase as the curve distance between the Xs increased, which is what we found (see Table 1).

During debriefing after the experiment, several subjects spontaneously described how they felt they performed the task. No consistent story emerged from these commentaries other than that most subjects, after viewing the displays, either "saw" that the two Xs lay on the same curve or "saw" that they lay on different curves. Many subjects expressed genuine surprise and skepticism upon hearing our hypotheses and no subject reported consciously scanning along the curve before making a decision. This aspect of the task is noteworthy because the response time data clearly show a dependence on curve distance. Given that response times were about 1,000 msec or less, and that subjects did not have the subjective impression of tracing, the results suggest that if curve tracing is in fact used, it is a fast and internal process and therefore is potentially a basic operation used by internal visual routines.

One aspect of the data in Experiment 1 was unexpected. As is evident in Figure 4, the increase in response time with distance was somewhat nonlinear. However, the most important aspect of the results is that response time did increase with distance. The details of the relationship between curve distance and tracing time are left as topics for future research.

Performance was similar in the two exposure durations, as can be seen in Figure 4. The 2,500-msec exposure duration allowed ample time for the subjects to move their eyes while displays were still in view. If the subjects waited for eye movements to take place before responding in the 2,500-msec condition, then we would expect longer response times than in the 250-msec condition, in which eye movements would be useless. However, response times were not longer in the 2,500-msec condition. Given the similarity of results in the two conditions, it appears that curve tracing may be performed internally even when subjects could track a curve with the eyes (see Yarbus, 1967).

EXPERIMENT 2

The results of Experiment 1 provide good evidence for the existence of a high-speed process capable of tracing curves internally. The primary goal in this experiment was to gather additional evidence for a curve-tracing process in the context of a different task. If the results from this experiment also provided support for the existence of a curve-tracing process, we would be confident that the results from the first experiment were not due to a superficial aspect of the task used.

In Experiment 2, all displays contained two nonintersecting curves and one of the curves always had two Xs on it. The distance between these Xs along the curve was varied as in the first experiment. In half of the displays

a small gap interrupted the curve between the two Xs, whereas in the other half of the displays the curve joining the Xs was continuous. The task was to decide, as rapidly as possible, whether or not there was a gap along the curve joining the Xs.

The data from no-gap trials (trials with a continuous contour joining the Xs) should provide additional evidence for the existence of a curve tracing process. A reasonable strategy for deciding that the curve joining the two Xs does not contain a gap is to trace the entire curve between the two Xs and to note that there is no gap. In this case, response times should increase systematically as longer distances separate the two Xs along the curve. For gap trials (trials with a gap in the curve joining the two Xs), there are several possible processing strategies with corresponding predicted patterns of results. One of the more straightforward possibilities is that subjects trace the curve from the central X and make a response as soon as they locate a gap on that curve. In this case, response times should reflect the curve distance between the central X and the gap. Another possibility is that subjects always check that the gap is in fact located between two Xs when they find a gap on the curve intersecting the central X. In this case, response times should reflect the curve distance between the Xs. Yet another possibility is that subjects first locate a gap and then trace the curve to either side, looking for Xs. Response times in this case should reflect the curve distance between the target gap and the nearest X. However, this latter strategy would not be very efficient in the present task because all displays contained several gaps on the curve that did not contain the Xs (the noncentral curve). Thus, subjects should prefer one of the first two strategies.

Method

Subjects. Sixteen University of Saskatchewan undergraduates with normal or corrected-to-normal vision participated as unpaid subjects. No subject in this experiment had participated in the first experiment reported in this paper.

Materials. The stimuli were constructed using the same basic patterns of curves used in Experiment 1. However, in all displays the two Xs were on the same contour. Figure 5 illustrates the different basic patterns of curves used throughout this study and the placement of the Xs and the gaps in the present experiment. For each basic pattern of curves, four distances between the two Xs were used. Displays in which the curve between the two Xs did not contain a gap were thus almost identical to the stimuli used in the "same" trials of Experiment 1 (the only difference was the presence of gaps along the noncentral curve; see Display A in Figure 5). Displays with a gap between the two Xs were constructed as follows: For displays with distances of 2, 3, and 4 units between the Xs, the location of gaps corresponded with the possible position of the noncentral X between the two Xs (e.g., the location of the dots on the curve passing through the central X in Figure 2). Therefore, when there was a gap between the two Xs in displays with curve distances of 2, 3, or 4 units, the gap was at a constant physical distance to the central X, and all gaps were at the same distance from the fovea. Thus, physical distance and retinal eccentricity of gap locations could not influence the results. For displays with a distance of 2 units between the two Xs, there was only one location at which a gap could occur. An example display is shown in Figure 5 (Display C). For displays with a distance of 3 units there




Figure 5. Example stimuli in Experiment 2. Display A shows a no-gap trial with a distance of 1 unit between the Xs. B–H show gap trials. In B and C, the distance between the Xs is 1 and 2 units, respectively. In D and E, the Xs are 3 units apart; the gap is at the 1-unit mark in D and at the 2-unit mark in E. F, G, and H show gap trials with 4 units between the Xs; the gap is at the 1-unit mark in F, at the 2-unit mark in G, and at the 3-unit mark in H. This figure also shows the eight basic patterns of curves used in Experiments 1 and 2.

were two possible locations (see Figure 5, Displays D and E), and for displays with a distance of 4 units there were three possible locations (see Figure 5, Displays F, G, and H). For displays with a curve distance of 1 unit between the Xs, the gap was located midway along the curve joining the two Xs and was thus somewhat closer to the fovea than were the other gaps (see Display B in Figure 5).

For each curve distance between the Xs, there were 24 displays in which there was a gap between the two Xs, with an equal number of displays using each possible gap location. Also, for each curve distance between the Xs, there were 24 displays in which there was no gap (three replications of each of the eight basic patterns of two curves used in Experiment 1). In every display, between two and four gaps were placed on the curve that did not intersect the Xs (the number was determined at random for each display). These gaps were located at the possible locations that were used for the noncentral X in "different" trials of Experiment 1 (see Figure 2). Thus, these gaps were at the same physical distance from the central X and at the same distance from the fovea as "target" gaps in gap trials. This procedure ensured that subjects could not respond simply by detecting the presence of any gap—all displays contained gaps, even displays in which the appropriate response was "no gap."

An additional set of displays was used in practice trials. These displays had an equal number of gap and no-gap trials, with each possible curve distance between the two Xs used once in a gap trial and once in a no-gap trial.

Procedure. The task consisted of deciding whether or not there was a gap in the curve joining the two Xs, and was to be accomplished as rapidly as possible while keeping errors to a minimum. The stimuli were presented using the same equipment and trial structure used in the 2,500-msec condition of Experiment 1. The subjects were instructed to press the response key labeled "gap" if the display contained a gap in the curve joining the two Xs, and to press the key labeled "no gap" if there was no gap.

The session began with 16 practice trials followed by a short question and answer period to ensure that subjects understood the instructions. There were 192 experimental trials in all, with an equal number of gap and no-gap trials. The gap and no-gap trials were divided equally among the four possible curve distances separating the Xs. The gap was in the same location in gap trials with dis-

tances of 1 and 2 units (24 trials for each distance). In displays with 3 units separating the two Xs, each of the two possible gap locations was used in 12 trials. In displays with 4 units separating the two Xs, each of the three possible gap locations was used in 8 trials.

The trials were divided into eight blocks of 24; each block contained an equal number of gap and no-gap trials for each curve distance. For any given curve distance in the gap trials, an equal number of trials using a particular gap location occurred in each block of trials. Displays were assigned to each block at random within the blocking constraints outlined above. The order of trials within a block was random, with the constraint that no more than five consecutive trials require the same response. All subjects saw all 192 experimental trials, with a short rest period between the third and fourth blocks and between the sixth and seventh blocks. The order of blocks was counterbalanced between subjects, using a Latin square design.

Results

Response times. The data from each subject and curve distance was first corrected for obvious outliers by the rule that any observation greater than the median plus two times the interquartile range be replaced by the median of the observations for that subject, response, and distance (see Tukey, 1977). This procedure resulted in the replacement of 1.1% of the observations. A preliminary analysis revealed that practice (blocks) did not interact with any other factor in the experiment, and this factor was not considered further.

The mean response time for each curve distance between the Xs in no-gap trials and in gap trials is shown in Figure 6. Consider first the results from no-gap trials. These data were submitted to a repeated measures analysis of variance with subjects and distance as factors. As is evident in Figure 6, the mean decision time increased monotonically with increasing distance between the two Xs [$F(3,45) = 18.51, p < .0001, MSe = 3.995 \times 10^3$].




Figure 6. Mean response time for no-gap and gap trials at each distance along the curve between the two Xs in Experiment 2. Each unit distance along the curve corresponded to 2.2° of visual angle.

The monotonic increase in response times with distance was reflected in the linear trend in response times over distance [$F(1,15) = 38.86$, $p < .0001$, $MSe = 5.644 \times 10^3$]. None of the other trends were significant ($F < 1$ in all cases). These results provide good additional evidence for the use of a curve-tracing visual routine.

Now consider the results from the gap trials. The mean response time for each curve distance between the Xs and for each gap location is given in Table 2. First, we analyzed the data from trials in which more than one gap location was used, to determine whether the distance of the gap from the central X had any effect on response times. Consider the results from trials in which the two Xs were separated by 4 units along the curve (see the rightmost column in the top portion of Table 2). These data were submitted to a separate repeated measures analysis of variance in which subjects and gap location were factors. There were no significant differences between the mean response times depending on gap location ($F < 1$).

Now consider the data from trials with a distance of 3 units between the Xs (see the second-to-last column in the top portion of Table 2). In this case, the mean response time for gaps 1 unit from the central X along the curve was longer than that for gaps 2 units from the central X along the curve [$F(1,15) = 4.47$, $p < .052$, $MSe = 8.254 \times 10^3$], although this difference was only marginally significant.

Together, the results across gap locations for 3 and 4 units of distance between the Xs are inconsistent with the notion that the subjects traced the curve from the central X to the noncentral X and responded as soon as they detected a gap along the curve. If the subjects had done this, response times would have increased with an increase in distance between the central X and the gap. As can be seen

in Table 2, however, this pattern of results was not observed. Rather, it seems that subjects always traced along the curve until they encountered the second X before making a response. Figure 6 also displays the mean response time for the gap trials averaged over all gap locations for each curve distance between the Xs. The mean response times for gap trials mirrored the times for no-gap trials. A separate repeated measures analysis of variance on data from gap trials, in which we considered subjects and distance between Xs as factors, confirmed that response times increased with curve distance between the Xs [$F(3,45) = 45.15$, $p < .0001$, $MSe = 3.234 \times 10^3$]. The linear component of the increase in mean response times was highly significant [$F(1,15) = 75.20$, $p < .0001$, $MSe = 5.489 \times 10^3$]. The quadratic component was also significant [$F(1,15) = 13.80$, $p < .003$, $MSe = 1.825 \times 10^3$], reflecting the slight nonlinearity of the results evident in Figure 6.

Finally, we also performed a repeated measures analysis of variance in which we combined the data from no-gap and gap trials. In this analysis we considered subjects, trial type (no-gap/gap), and distance as factors. This analysis confirmed that no-gap trials resulted in shorter response times than gap trials (no-gap, 1,000 msec; 1,067 msec) [$F(1,15) = 12.61$, $p < .003$, $MSe = 1.151 \times 10^4$]. As can be seen in Figure 6, the effect of distance was similar for gap and no-gap trials [$F(3,45) = 2.61$, $p > .06$, $MSe = 2.918 \times 10^3$].

Considering all the results displayed in Figure 6, the average rate of curve tracing was 48° of visual angle per second.

Errors. The percentages of trials on which errors occurred across conditions are presented in Table 2. In the no-gap condition, the number of errors was highest for trials with 1-unit distance between the Xs and was lower for other distances, over which the error rate was relatively constant. For gap trials, errors tended to increase with dis-

Table 2
Mean Response Time (RT, in msec) and Error Rate for Gap and No-gap Trials for Each Curve Distance Between the Xs in Experiment 2

		Curve Distance Between Xs (in Units)			
Gap Distance		1	2	3	4
Gap Trials					
RT	.5	979			
% error		1.3			
RT	1		1,013	1,116	1,196
% error			2.9	1.6	2.3
RT	2			1,048	1,177
% error				1.0	1.6
RT	3				1,214
% error					6.3
No-gap Trials					
RT		928	964	1,026	1,082
% error		2.3	1.8	1.3	1.6

Note—The results for gap trials are displayed for each gap distance to the central X along the curve. Each unit of gap distance and curve distance between Xs corresponded to 2.2° of visual angle. In both cases, distance along the curve was measured from the central X.

tance overall. Thus, the results do not appear to be due to speed-accuracy trade-offs. The error rate for each subject, trial type (gap/no-gap), and distance between the Xs was submitted to a repeated measures analysis of variance with these three factors. There were no significant effects ($p > .17$ in all cases).

Discussion

The main results are straightforward: Longer curve distances between the Xs in both no-gap and gap trials are associated with longer response times. Thus, this experiment provides additional evidence that subjects make use of a curve-tracing routine to solve a problem involving spatial relations.

Results from trials in which there was a gap between the Xs suggest that subjects traced the entire curve joining the two Xs. Perhaps one would have expected subjects to respond as soon as a gap was encountered along the curve, that is, to adopt a self-terminating mode of processing. Whether exhaustive processing on gap trials is a basic limitation or is due to some incidental aspect of the task will need to be addressed in future research. Given that self-termination appeared to take place in Experiment 1 ("same" trials) and in Experiment 2 (gap and no-gap trials) as a function of the location of the noncentral X, however, we suspect that the exhaustive processing on gap trials may have been the result of the specific instructions in the experiment, rather than the result of a fundamental property of curve tracing. The instructions given to subjects asked them to respond with the "gap" key when they saw a gap *between the two Xs*. It is possible that subjects interpreted our instructions literally and checked to see *than the gap* they found along the curve intersecting the central X was in fact between two Xs. To do this would require tracing the entire length of the curve between the Xs. This issue is left for future research.

As is evident in Figure 6, gap trials were uniformly responded to more slowly than no-gap trials (a difference of 67 msec). This difference is consistent with the notion that subjects performed an additional check when they detected the presence of a gap on the curve intersecting the central X (e.g., ensuring that the gap was indeed between two Xs). Alternatively, the additional time may simply reflect a tendency to respond "no-gap" and to change response selection when a gap is detected. Another possibility is that gap trials required rechecking more often than no-gap trials. This possibility is related to Krueger's (1978) noisy-operator theory for perceptual matching. In our experiment, it is possible that internal noise in the visual system was more likely to result in the presence of a spurious gap in a continuous curve than in the presence of a spurious continuous curve between two points or a spurious filling in of an existing gap. If so, the detection of a continuous curve between two points would be unlikely to result from internal noise, and thus would not require additional checking. In contrast, the detection of a curve with a gap could have resulted from

internal noise rather than from the presence of an actual gap, and thus the perception of a gap may have required additional processing (rechecking perhaps) to ensure that the gap was not spurious.

Finally, the fact that the subjects were slower in gap than in no-gap trials argues against a strategy in which gaps are located first in the display, followed by tracing of the curve to either side. In no-gap trials, all the distractor gaps would need to be processed before a "no gap" response could be made, whereas in gap trials, only half of the gaps in the display, on average (assuming random selection), would need to be processed. Thus, if the subjects first looked for gaps and then traced, no-gap trials should have resulted in longer response times than gap trials, which is the opposite of what was observed.

GENERAL DISCUSSION

Curve Tracing in the Experimental Task and in Natural Viewing

In our introduction, we made two suggestions: that a tracing operation may be used in integrating local information along a curve, and that visual routines may play a role in the perception of abstract shape properties and spatial relations. The results reported in this paper suggest that, in certain visual tasks, people use a rapid internal visual process to trace curves in a display. The time needed to trace a curve is monotonically related to the distance to be traced, with longer distances requiring more time. In Experiment 1, the tracing operation provided an identity or "sameness" operator: It served to verify that the two Xs were marked on one and the same curve, rather than on two disconnected curves.

On a more general level, our results suggest that the seemingly immediate and effortless perception of spatial properties and relations may be misleading. It may conceal, in fact, the application of elaborate and specialized processes, processes we referred to in the introduction as visual routines. These routines can be broken down into more elementary constituents, one of which is the internal curve-tracing operation suggested by this study. It is also possible that the curve-tracing operation itself is decomposable into more primitive operations; it may itself be a visual routine, rather than a primitive operation.

In this study we employed schematic drawings, rather than natural scenes, because such simplified stimuli allow more flexibility in adapting the stimuli to the operation under investigation. However, it seems reasonable to expect that curve tracing is also employed in the perception of natural scenes, for three reasons. First, the experimental displays provided a simplified version of a problem that is in fact quite general: the integration of relevant information from different locations. If an image is first analyzed primarily by local operators, then a subsequent stage is required for tying together the relevant local information. A curve-tracing operation can play a useful role in this integration stage. Second, both psychological (e.g., Attneave, 1954) and physiological (e.g.,

Hubel & Wiesel, 1968) evidence has suggested that contours and boundaries play a fundamental role in natural perception. A process that participates in the analysis of contours, therefore, could play an important role in perception in general. Third, insofar as we examine tasks for which the visual system exhibits a high degree of proficiency, we are likely to be exploring fundamental capacities of the system. It seems reasonable to hypothesize that a useful operation that is used spontaneously and proficiently by the visual system in these tasks is also employed in other tasks.

A number of differences between the simplified experimental conditions used in this study and more complex natural scenes are noteworthy. First, the experiments employed a simple type of curve composed of a dark contour against a light background. Such contours are rare in natural scenes. If a curve-tracing operation is indeed employed by the visual system, as we have suggested, it should be applicable to the types of contours and edges found in natural scenes. A direct examination of this notion would require additional studies. Second, under more complex conditions, curve tracing is probably only one of several operations available in the system. We do not wish to imply that curve tracing by itself is sufficient to establish shape properties of, and spatial relations among, contours in a scene. Rather, we suggest that curve tracing may be one of a set of operations out of which more elaborate processes can be constructed.

Alternative Explanations

We have suggested that a plausible interpretation of our experimental results is that sequential curve tracing was employed in both experiments. There are, however, a number of alternative interpretations worth considering before we accept this conclusion. The first possibility is that the pattern of results arises from complexity effects, rather than from sequential tracing. In this view, complex figures are more difficult to analyze visually than simple ones, and the longer response times in our study were associated with more complex figures. This hypothesis is unlikely, however, because identical curves were used in the different distance conditions (the only difference was in the location of the second X which resulted in different distances between the Xs along the curve; see Figure 3).

A second possibility is that the entire curve was processed in parallel, as opposed to the proposed sequential tracing. There are, in fact, known methods that operate in parallel over the entire curve for solving problems similar to the ones in our experiments (i.e., for determining whether two given points lie on the same connected component) (Cook, 1983; Shiloach & Vishkin, 1982). Although it is sometimes difficult to distinguish between sequential and parallel implementations, given only input and output relationships, it seems that parallel schemes (at least of the types that have been found so far) for this particular task are unlikely to be embodied in the visual system, for two reasons. First, they require complicated

numerical calculations and assume total connectivity (i.e., each processor in the network communicates directly with every other one). Second, the totally parallel schemes are inherently insensitive to the size of the input (in our case, curve length). Parallel schemes thus have the advantage that long curves are processed almost as rapidly as shorter ones, but this comes at the expense of complicated computations and a very large number of connections. Although it may be possible to degrade the performance of the parallel schemes and make them somehow more dependent on curve length, this would be rather ad hoc. A more likely hypothesis is that the visual system has compromised on a solution that is more sequential in nature and therefore depends more on curve length.

Finally, tracing could proceed along the curve not in a continuous manner, or in very small steps from one point to the next, but in larger "chunks." In this case, the distance effect could be explained in one of two ways: (1) with increased distance along the curve the number of chunks increases, or (2) the number of chunks remains constant, and the size of the individual chunks is responsible for the distance effect. The second alternative seems unlikely, because it seems to entail knowledge of the length of the curve separating the two Xs in a particular trial in order to adjust the chunk size. It is not clear how the visual system could estimate the length of the curve separating the two Xs, given that this distance varied at random from trial to trial and that no other aspects of the displays were correlated with curve distance. The second alternative also begs the question of why processing a larger chunk should take longer than processing a smaller chunk. The first alternative, on the other hand, appears entirely plausible and worthy of future study.

Possible Relations to Mental Scanning and Orienting Attention

The results obtained in this study bear some resemblance to the results of studies in the areas of mental scanning and orienting of visual attention. Kosslyn and others have used image scanning as a "tape measure" for mental images, and have argued that mental images preserve metric distance information (see Finke & Pinker, 1982, 1983; Jolicoeur & Kosslyn, 1985; Kosslyn, 1980; Kosslyn, Ball, & Reiser, 1978). In general, studies of mental scanning have focused on the nature of mental images rather than on the nature of the scanning mechanism. The scanning process itself has been of only peripheral interest in these studies, and little is known about the exact nature of mental scanning other than that more time is required to scan longer distances (but see Reed, Hock, & Lockhead, 1983).

Nonetheless, on the surface, results from mental scanning experiments appear similar to ours. Experiments on mental scanning have consistently shown a linear relation between response time and distance scanned across on a display or in an image. There are, however, a number of important differences between our work and work on mental scanning. In most mental scanning experiments,

subjects have been instructed to imagine a small black speck moving as rapidly as possible within the image from one point to another. This task requires that the subjects actively imagine a speck moving within a mental image; the process under study is a conscious one. In our experiments, the subjects were not instructed to trace curves, to perform mental scanning, or to imagine a speck moving along the central curve. Rather, they were asked to perform the task as rapidly as possible by any possible means. Furthermore, they were not informed that curve distance was varied systematically, and in postexperiment conversations, they did not report consciously scanning the curves in our displays. It is unlikely, therefore, that our subjects were engaging in conscious mental scanning.

In Experiment 1, displays were sometimes in view for 2,500 msec; thus, the task was performed while the display was continuously present perceptually. Therefore, it is likely that the subjects made use of a perceptual representation rather than an imagined one. Most image scanning experiments require subjects to scan between two locations in a straight line, and the locations are not usually connected by a curve or boundary in the image, as they were in our experiments. However, Reed et al. (1983) asked subjects to perform conscious mental scanning over different types of contours (straight line, curved spiral, square spiral). One finding of interest was that the type of scan path affected the scanning rate: square-spiral paths were scanned at the slowest rate, whereas straight-line paths were scanned at the fastest rate. Of greater interest here, however, is that even in the fastest condition the scanning rate was about $6.7^\circ/\text{sec}$, which is much slower than the average scanning rate of $40^\circ/\text{sec}$ found in our curve-tracing experiments. This difference in scanning rate suggests that there may be important differences between conscious mental scanning and curve tracing. Our experiments, therefore, belong more properly in the domain of perception than in the domain of mental scanning.

Although most image scanning experiments involve mental images of patterns that are well learned, Finke and Pinker (1982, 1983) asked subjects to decide whether an arrow pointed to one of the dots in a previously displayed pattern of dots. In those trials in which the arrow did point to a dot, the distance between the arrow and the nearest dot was varied systematically across trials. The mean time to respond that the arrow did point to a dot (in those trials in which it did) increased linearly as the distance between the arrow and the dot was increased. Finke and Pinker's results have been interpreted as evidence of spontaneous image scanning as a means of establishing a particular spatial relation between two elements in a visual display (the relation *points to*). Finke and Pinker assumed that the processes they studied in their paradigm were identical to those used in other mental scanning paradigms. The assumption is plausible, of course, because of the similarity in the functional relationships observed between response time and distance. However, more evidence is required to establish that the same mechanisms are, in fact, used in the various imagery paradigms or in curve tracing.

Studies on the orienting of attention have elegantly demonstrated a dissociation between the locus of attention and retinal position (for relevant reviews, see Posner, 1978, 1980). That is, the visual system seems capable of processing information from cued locations preferentially compared with information from other locations. Furthermore, there is evidence that shifting attention from one location to another takes time, and that the locus of attention passes through intermediate locations (Shulman, Remington, & McLean, 1979). One may think, therefore, that the time needed to make an attentional shift would depend on the distance to be shifted over, and in fact there is some empirical support for this view (Tsal, 1983). Based on the above information, one could postulate that curve tracing and shifts of attention share a common processing component. Recently, however, Remington and Pierce (1984) argued that the time needed to shift attention does not depend on the distance to be shifted across. The exact nature of attentional shifts and the associated underlying mechanisms are still in dispute. Perhaps the similarity between attentional shifts and curve tracing is only apparent. There is no compelling evidence at this time that the two experimental tasks in fact require the use of common processing mechanisms. We do not wish to express strong skepticism regarding possible relationships between curve tracing and attention shifts; we simply raise the possibility that deliberate attention shifts of the type studied in mental scanning experiments or in orienting of attention experiments and the rapid curve tracing suggested by our results may not be one and the same process.²

In any case, it seems likely that rapid internal tracing of curves could provide the visual system with a valuable means of integrating information from different spatial locations. Furthermore, even if future research reveals that the processes of curve tracing, attention shifting, and image scanning share a common scanning operator, our experiments suggest a new role for such a scanning operator: the integration of information along a contour.

REFERENCES

- ATTNEAVE, F. (1954). Some informational aspects of visual perception. *Psychological Review*, **61**, 183-193.
- COOK, S. (1983). The classification of problems which have fast parallel algorithms. In G. Goos & J. Hartmanis (Eds.), *Lecture Notes in Computer Science*, Vol. 158. Berlin: Springer-Verlag.
- FINKE, R. A., & PINKER, S. (1982). Spontaneous imagery scanning in mental extrapolation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **8**, 142-147.
- FINKE, R. A., & PINKER, S. (1983). Directional scanning of remembered visual patterns. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 398-410.
- GORDON, R. B. (1984). *Boundary tracing, image scanning, attentional shifts, and apparent motion: An exploration of the possibility of a common shift operator*. Unpublished master's thesis, University of Saskatchewan, Saskatoon, Saskatchewan, Canada.
- HUBEL, D. H., & WIESEL, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, **195**, 215-243.
- JOLICOEUR, P., & KOSSLYN, S. M. (1985). Is time to scan visual images due to demand characteristics? *Memory & Cognition*, **13**, 320-332.

- KOSSLYN, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- KOSSLYN, S. M., BALL, T. M., & REISER, B. J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception & Performance*, **4**, 47-60.
- KRUEGER, L. E. (1978). A theory of perceptual matching. *Psychological Review*, **85**, 278-304.
- MINSKY, M., & PAPERT, S. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
- POSNER, M. I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum.
- POSNER, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, **32**, 3-25.
- REED, S. K., HOCK, H. S., & LOCKHEAD, G. R. (1983). Tacit knowledge and the effect of pattern configuration on mental scanning. *Memory & Cognition*, **11**, 137-143.
- REMINGTON, R. W., & PIERCE, L. (1984). Moving attention: Evidence for time-invariant shifts of visual selective attention. *Perception & Psychophysics*, **35**, 393-399.
- SHILOACH, Y., & VISHKIN, U. (1982). An $O(\log n)$ parallel connectivity algorithm. *Journal of Algorithms*, **3**, 57-67.
- SHULMAN, G. L., REMINGTON, R. W., & MCLEAN, J. P. (1979). Moving attention through visual space. *Journal of Experimental Psychology: Human Perception & Performance*, **5**, 522-526.
- TSAL, Y. (1983). Movements of attention across the visual field. *Journal of Experimental Psychology: Human Perception & Performance*, **9**, 523-530.
- TUKEY, J. W. (1977). *Exploratory data analysis*. Reading, MA: Addison-Wesley.
- ULLMAN, S. (1984). Visual routines. *Cognition*, **18**, 97-159.
- YARBUS, A. L. (1967). *Eye movements and vision*. New York: Plenum.

NOTES

1. The two conditions in Experiment 1 were conducted as separate experiments. The 250-msec condition was carried out first. The two experiments were identical in all respects except that 8 practice trials were used in the 250-msec condition and 16 were used in the 2,500-msec condition.

2. The only study we know of that investigated whether attention shifting, curve tracing, and mental scanning share a common scanning component did not provide evidence of such a shared component (Gordon, 1984). Each subject in that study participated in a battery of tasks that included image scanning, curve tracing, and attention shifting. The study employed an individual-differences approach and correlated scanning rates from different subjects across the battery of tasks. None of the correlations across the three categories of tasks (attention shifts, curve tracing, and mental scanning) provided evidence for a common processing component.

(Manuscript received March 21, 1985;
revision accepted for publication August 19, 1985.)