

Global perceptual processing in problem solving: The case of the traveling salesperson

THOMAS C. ORMEROD and EDWARD P. CHRONICLE
Lancaster University, Lancaster, England

The traveling salesperson problem (TSP) consists of finding the shortest tour around a set of locations and is an important task in computer science and operations research. In four experiments, the relationship between processes implicated in the recognition of good figures and the identification of TSP solutions was investigated. In Experiment 1, a linear relationship was found between participants' judgments of good figure and the optimality of solutions to TSPs. In Experiment 2, identification performance was shown to be a function of solution optimality and problem orientation. Experiment 3 replicated these findings with a forced-pace method, suggesting that global processing, rather than a local processing strategy involving point-by-point analysis of TSP solutions, is the primary process involved in the derivation of best figures for the presented TSPs. In Experiment 4, the role of global precedence was confirmed using a priming method, in which it was found that short (100 msec) primes facilitated solution identification, relative to no prime or longer primes. Effects of problem type were found in all the experiments, suggesting that local features of some problems may disrupt global processing. The results are discussed in terms of Sanocki's (1993) global-to-local contingency model. We argue that global perceptual processing may contribute more generally to problem solving and that human performance can complement computational TSP methods.

The traveling salesperson problem (TSP) consists of finding the shortest closed tour around a set of locations in two-dimensional space. In other words, a route is found in which each location is visited only once and the starting point is returned to. The task is of considerable theoretical importance in computer science, since it provides a canonical example of NP-complete problems, a class of decision problem that can be shown to be computationally intractable (Garey & Johnson, 1979). Its computational intractability is nicely illustrated by the fact that the number of possible solutions to a TSP is given by $(n - 1)!/2$, where n is the number of locations (assuming that the direction of the tour is ignored). For a 20-location problem, $(n - 1)!/2$ is equal to 6.08×10^{16} . Assuming a computer program could find solutions at the rate of 1,000 per sec, it would take 6.08×10^{13} sec, or approximately 1.9 million years, to conduct an exhaustive search of the solution space for a 20-point TSP.

Although TSPs are computationally intractable, they have considerable practical significance in domains as

diverse as pipe laying, circuit wiring, and the organization of distribution networks, where optimizing the use of resources (e.g., time, effort, and materials) can determine profitability and success. Hence, considerable efforts have been made in the field of operations research to identify efficient heuristic algorithms that provide solutions of acceptable quality while minimizing computational cost. One approach is to use the relatively inefficient *nearest neighbor* heuristic algorithm, in which the order of point connection is determined by moving from one point to its nearest neighboring point until the tour is complete (Rosenkrantz, Sterans, & Lewis, 1974). A more efficient approach is offered by the 2-Opt heuristic algorithm, in which an initial approximate solution, based on connecting exterior points into a convex hull, is subsequently improved by comparing alternative connections between local point clusters and swapping them where this improves the solution (Sangalli, 1992). However, even the most efficient heuristic algorithms still require in the order of n^3 computations, for an n -point network, in order to compute solutions to within 2%–3% of optimality (Golden, Bodin, Doyle, & Stewart, 1980).

Given the theoretical and practical importance of TSPs, it is surprising that they have received little attention from psychologists. One approach that has been adopted by operations researchers is to use computer-based methods to find partial solutions, which are then completed by human participants (Krolak, Felts, & Marble, 1971). In another study, Polivanova (1974) demonstrated that, with simple problems containing between 4 and 10 points, humans found better solutions to TSP problems that were represented visually as points in a two-dimensional array

A report of Experiments 2 and 3 was presented at the British Psychological Society Cognitive Section Conference, Bristol, September 1995. This research was supported by a grant from the Faculty of Social Sciences, Lancaster University. We gratefully acknowledge the assistance given in data collection by Jacqueline Droogan (Experiment 2) and Adrienne Cutner (Experiments 3 and 4) and the thoughtful comments of Peter Walker and James MacGregor, as well as those of an anonymous referee. Correspondence concerning this paper should be addressed to T. C. Ormerod, Department of Psychology, Lancaster University, Lancaster, LA1 4YF, England (e-mail: t.ormerod@lancaster.ac.uk).

—Accepted by previous editor, Myron L. Braunstein




Figure 1. The traveling salesperson problems used in MacGregor and Ormerod (1996), together with the shortest tour around each set of points (the same problems were used as stimuli in the experiments reported here).

than to those represented as tabular lists of numerical distances between points. These studies indicated that humans apply visually based mechanisms in the solution of TSPs that might complement computer-based heuristic algorithms, but they did not assess the quality of human solutions per se.

MacGregor and Ormerod (1996) conducted the first systematic study comparing the quality of TSP solutions produced by humans with those produced by computer-based heuristic algorithms. In their study, participants drew solutions to a set of 10- and 20-point TSPs whose complexity was varied in terms of the number of interior points that lay within a boundary defined by a simple convex hull around the points. Figure 1 shows the 10-point problems used in their first experiment, together with optimum solutions. They found that the least optimal paths produced by participants were less than 1% longer than the estimated optimal route for both 10- and 20-point problems; best path lengths exceeded the optimum solution for each problem by less than 0.3%. Solutions were also calculated with three heuristic algorithms: the nearest neighbor heuristic algorithm and two variants of the

2-Opt heuristic algorithm (largest interior angle and convex hull cheapest interior; for a description, see MacGregor & Ormerod, 1996). The average percentage above optimum solution achieved by the best path length from each heuristic algorithm was 7.3% for nearest neighbor, 2.2% for largest interior angle, and 2.1% for convex hull cheapest interior. Thus, the path lengths produced by MacGregor and Ormerod's participants were better by at least an order of magnitude than the ones produced by these heuristic algorithms.

Although the experiments of MacGregor and Ormerod (1996) demonstrate that humans can produce TSP solutions of a high quality, the mechanism by which they identify good solutions has yet to be determined. Two candidate explanations can be identified a priori, which differ in whether the account of human performance is in terms of a local or a global analysis of TSP arrays. If human solvers use a local processing strategy, they would select TSP solution components on a point-by-point basis, basing their decisions on the processing of localized problem features. The results of MacGregor and Ormerod are sufficient to rule out the use by human solvers of a near-

est neighbor heuristic. However, there may be others that sample localized features of the problem array more efficiently. Alternatively, human performance may be based on extracting global properties of a TSP stimulus.

MacGregor and Ormerod (1996) hypothesized that complexity of TSPs for humans is a function of the number of points that lie within a convex hull drawn around the perimeter of a problem. They found that response uncertainty (measured by summing, across all possible pairwise connections, the number of participants connecting each pair of points) increased significantly with the number of interior points ($r = .93$ for the 10-point problems, $r = .79$ for the 20-point problems). Their results provide evidence that humans identify a convex hull in solving TSPs and that the ease with which a convex hull can be identified is a function of the number of interior points in the problem. Identification of a convex hull through a point-by-point inspection of localized problem features is, of course, fundamental to many heuristic algorithms.

Given the construction of the problems used by MacGregor and Ormerod (1996), however, it is possible that the obviousness of the convex hull is highly correlated with the figural goodness of the optimal solution. The starting point for the construction of the problems they used was, in effect, a notional circle: Problem exemplars were created simply by taking between one and five dots from the perimeter to the interior of the circle. Hence, "simpler" problems (with fewer internal points) were also more circle-like. One strategy available to participants would, therefore, have been to approach the task by drawing a solution that provided the best, most circular figure. This suggestion finds some empirical support in the work of Pomerantz (1981), who found that when participants were asked to connect points of a random dot array in a manner that illustrated a pattern that they perceived within the array, they frequently connected points in the shortest possible path. As MacGregor and Ormerod themselves point out, "The task of the TSP may happen to parallel what it is natural for the perceptual system to do in any case when presented with an array of dots" (p. 537).

A possible and plausible model for human performance on the TSP is, therefore, that participants initially process the TSP stimulus globally by identification of a best figure fitting the dot array and utilize this figure either to generate a route or to evaluate a presented route, depending on the experimental task in which they are involved. This model makes a number of testable predictions. First, and most obvious, if TSP solutions varying in optimality are rated for goodness of figure, the greater the optimality of the solution, the better should be the rating of good figure. Second, if TSP solutions are presented to participants who are asked to judge whether the solution is in fact optimal, the time taken to respond will reflect the degree of mismatch between perceived best figure and presented route.

The question then arises as to the nature of the perceptual mechanisms by which best figure is itself identified. An accumulating body of evidence (Sanocki, 1993; see also Sanocki, 1991), using the method of part-whole priming, supports an account of perceptual processing of objects wherein some subset of elements of the image of an object is used to generate a global hypothesis about the object: this global hypothesis then constrains further processing at other levels. Sanocki (1993) terms this *global-to-local contingency*. In the context of the TSP, this account implies that a global hypothesis about best figure might be available early in processing and could potentially inform a decision about optimal route, by the process of comparison outlined above. Such a theoretical background allows the derivation of experimental predictions about the effect of the duration of a prime stimulus on decision times for judging TSP optimality.

The four experiments reported here were, therefore, conducted to investigate the processes by which humans identify solutions to TSPs. In the first experiment, the suggestion, above, that optimal solutions to the TSPs used by MacGregor and Ormerod (1996) were perceived as having good figure, was examined. Experiments 2 and 3, using a solution identification task, tested the prediction that TSP solutions are evaluated by judging them against a best figure. In a final experiment, using a part-whole priming method, the question of whether global or local processes are dominant in the derivation of a best figure was examined.

EXPERIMENT 1

Method

Participants. Twenty-four undergraduate students at Lancaster University (17 women and 7 men, mean age 20.8 years) volunteered to participate.

Materials. The 10-point problems shown in Figure 1 were used as materials. The points of each stimulus were placed within a 12-cm² grid (for the decision rules regarding exact point coordinates, see MacGregor & Ormerod, 1996). Five problems had a similar spatial representation, in which the number of internal points within the problem's convex hull boundary was varied between one and five. The sixth problem was the Dantzig, a 10-node TSP with five internal points that is known in operations research to be resistant to many heuristic algorithms (Dantzig, Fulkerson, & Johnson, 1959). For each of the six problems, five solutions were constructed in the ranges 0%, 11%–18%, 21%–28%, 32%–37%, and 41%–47% longer than the optimal solution (the spatial layout of each problem necessitated solutions of slightly different optimality within each range). For convenience, these ranges are subsequently labeled with the midpoints 0%, 15%, 25%, 35%, and 45%.

Design. There were two experimental factors, both repeated measures: optimality, with five levels (0%, 15%, 25%, 35%, and 45%), and problem, with six levels (one through five internal points, plus the Dantzig). The experimental measure was a judgment of good figure of the presented TSP solution on a numerical scale, where 1 represented a *good figure* and 5 a *poor figure*.

Procedure. Judgments of figure were collected individually from participants, using the computerized SuperLab system. Prior to the commencement of the experiment, standardized instructions that explained the nature of the judgment to be made were pre-

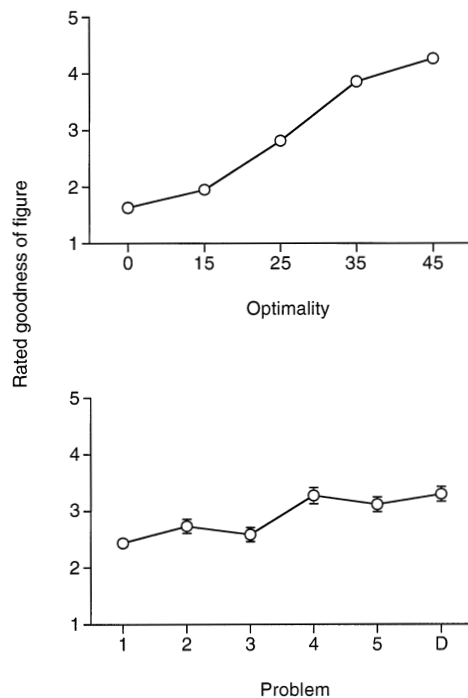


Figure 2. Mean ratings of goodness of figure (1 = best, 5 = worst) for each level of optimality and problem in Experiment 1. Vertical lines depict standard errors of the means (these are omitted for the optimality data, since they fall within the space taken by the plotting point in all cases). Levels of optimality are labeled with the approximate percentage distance over the optimal solution, and levels of problem are labeled with the number of points that lie inside a convex hull drawn around the problem perimeter (D = Dantzig problem).

sented. Each participant received 30 experimental trials presented in a single block. The order of trials was randomized with respect to optimality and problem. For each trial, a single stimulus appeared, consisting of a tour drawn round a set of points. The participant then simply rated the goodness of figure of the tour by pressing the appropriate number on the computer keyboard. No time constraints were placed on each trial; adjacent trials were separated by an interval of 2 sec, during which time a blank screen was shown.

Results

Goodness-of-figure judgments were subjected to a two-way analysis of variance (ANOVA), over the repeated measures factors of optimality and problem. An alpha level of .05 was used for all the statistical analyses in this and subsequent experiments. All reported a posteriori probabilities were corrected for sphericity, using the Geisser–Greenhouse conservative test (Winer, Brown, & Michels, 1991). We also report ω^2 as a measure of effect size for each variable, using the formulas for repeated measures designs given by Dodd and Schultz (1973). There was a significant main effect of optimality [$F(4,92) = 72.843$, $MS_e = 2.628$, $p < .001$, $\omega^2 = .492$], which is illustrated in Figure 2. Post hoc trend analysis revealed that the rated goodness of figure increased linearly with optimality of the solution [$F(1,92) = 283.097$,

$MS_e = 2.628$, $p < .001$]. There was also a significant main effect of problem [$F(5,115) = 28.900$, $MS_e = 0.576$, $p < .001$, $\omega^2 = .052$; see Figure 2]. Post hoc trend analysis revealed that rated goodness of figure declined linearly with the complexity of the problem [$F(1,115) = 113.564$, $MS_e = 0.576$, $p < .001$]. In addition to these main effects, there was a significant interaction of problem with optimality [$F(20,460) = 4.442$, $MS_e = 0.433$, $p < .001$, $\omega^2 = .019$]. Because its effect size was small, in comparison with those of the main effects, and the rank order of mean judgment scores across levels of optimality was identical at each level of problem, this interaction is not considered further.

Discussion

In this experiment, the participants consistently rated optimal solutions to TSPs as having good figures and successively less optimal solutions as having successively less good figures. Furthermore, there was a linear relationship between TSP complexity (as defined by the number of interior points to a convex hull) and good figure judgments: The simpler the TSP, the better the figure was judged to be. These results provide initial support for the suggestion made earlier that human performance on these particular TSPs may be underpinned by a strategy that equates best figure with optimal route. The results of Experiment 1 do not, however, permit an entirely unambiguous conclusion about this issue, since it is still conceivable that TSP solutions are provided through some other set of cognitive operations, unconnected with a judgment of good figure, and that the relationship with good figure demonstrated in the foregoing data is merely coincidental.

If a judgment of the best figure for a particular TSP reliably provides the optimal route, one plausible and reasonable strategy for making judgments about the optimality of a particular TSP solution would be to compute the best figure of the problem (we make no claim, as yet, as to how this might occur) and to compare this best figure with the presented route. Some relevant theoretical background is provided by Sanocki (1991, 1993), who propounded a contingency model in which the processing of late-arriving information about a visual stimulus is contingent on hypotheses derived from earlier arriving information. Early information provides a rough global reference frame; local information must then be cross-checked and localized onto this reference frame, if it matches appropriately. Sanocki (1993) argues that such a model “allows the system to process late-arriving details without the immense computational burdens of an unconstrained set of interpretations” (p. 879). In other words, late-arriving information that is highly consistent with prior hypotheses will be processed with minimal delay. In the context of the TSP solution identification strategy outlined above, the contingency hypothesis makes it likely that a presented optimal route would quickly and accurately be identified as such, since the route would be highly consistent with the best figure. In the same vein, it is likely that gross discrepancies between best figure

and presented route would also be identified quickly, since the prior hypothesis would be quickly found to be inappropriate in the very first attempt to integrate local information with the global reference frame. Presented routes that are slightly less than optimal, however, would cost processing time, since a considerable amount of local/global cross-checking might occur before a mismatch was detected. It might also be expected that more errors of route judgment would occur, since slight mismatches between global and local information might not be detected with great reliability.

In this hypothesis, no effect of TSP complexity is necessarily predicted, since the same processes of good figure identification and matching with presented solution are required in each case. This holds only under the simplifying assumption that the derivation of best figure is not itself contingent on the complexity of the problem.

In order to examine the role of figural goodness in TSP performance further, therefore, Experiments 2 and 3 used a solution identification method to test the foregoing predictions about the effects of solution optimality and problem complexity. This method had the additional advantage that it did not require the drawing of solutions by participants: This is necessarily a serial method of operation, and it was thought desirable to study identification processes in the absence of a production constraint of this nature.

EXPERIMENT 2

Method

Participants. Fourteen undergraduate and postgraduate students at Lancaster University (11 women and 3 men, mean age 24.5 years) volunteered to participate.

Materials. The TSP stimuli used in Experiment 1 were used again here; in addition, a further exemplar for each combination of optimality and problem was created. In the case of the 0% (optimal) stimuli, the second exemplar used a route that was visually similar to the optimal route but exceeded the optimal path length by an average of 2.6%. Four blocks of stimuli were then created by rotating the orientation of each exemplar through 0°, 90°, 180°, and 270°.

Design. There were three experimental factors, all repeated measures: optimality, with five levels (0%, 15%, 25%, 35%, and 45%); problem, with six levels (one through five internal points, plus the Dantzig); and block, with four levels (0°, 90°, 180°, and 270° orientations). The experimental measures were identification times, from presentation of stimulus to response selection, and errors (where *no* responses to 0% stimuli and *yes* responses to any other level of optimality were scored as errors).

Procedure. The data were collected individually, using the computer-based SuperLab presentation system. Each participant received 240 experimental trials presented in four blocks of 60 trials, each block containing stimuli in a single orientation. Within each block, the trials were randomized over optimality and problem. For each trial, the participant was shown a single stimulus consisting of a tour around a set of points. The task was to decide whether the route shown was the shortest (i.e., optimal) tour around the points, and to press a key marked *yes* if the route shown was optimal, otherwise to press a key marked *no*. At the start of each trial, the stimulus was presented, and it remained on the screen until the participant selected a response key. An interval of 2 sec separated each trial, during which a blank screen was shown. The participants rested for

1 min between each block of 60 trials. Each participant took approximately 25 min to complete the experiment.

Results

Mean identification times of all the responses (pooled over exemplars) and proportion of errors for each level of optimality, problem, and block were calculated. Outliers more than two standard deviations from the mean were removed prior to analysis. Three-way ANOVAs were calculated for the identification time and error data, over the repeated measures factors of optimality, problem, and block.

Identification times. Significant main effects were found of optimality [$F(4,52) = 4.877$, $MS_e = 1.092 \times 10^7$, $p = .02$, $\omega^2 = .017$] and block [$F(3,39) = 17.195$, $MS_e = 1.413 \times 10^8$, $p = .001$, $\omega^2 = .070$]. Mean identification times for each level of optimality are shown in Figure 3. The main effect of problem (means shown in Figure 3) was not significant [$F(5,65) = 1.086$, $MS_e = 4.924 \times 10^6$, $p = .37$, $\omega^2 < .001$]. The interaction between optimality and problem (means shown in Figure 4) was also significant [$F(20,260) = 2.804$, $MS_e = 2.612 \times 10^6$, $p = .014$, $\omega^2 = .010$]. No other interactions were significant.

Errors. Significant main effects were found of optimality [$F(4,52) = 12.425$, $MS_e = 0.655$, $p = .01$, $\omega^2 = .130$] and problem [$F(5,65) = 12.34$, $MS_e = 0.130$, $p = .0001$, $\omega^2 = .032$]; means are shown in Figure 3]. The main effect of block was not significant [$F(3,39) = 0.715$, $MS_e = 0.077$, $p = .495$, $\omega^2 < .001$]. The two-way interaction between optimality and problem was significant [$F(20,260) = 8.41$, $MS_e = 0.100$, $p = .0001$, $\omega^2 = .064$] and is illustrated in Figure 4. The three-way interaction between optimality, block, and problem was also significant [$F(60,780) = 2.18$, $MS_e = 0.048$, $p = .031$, $\omega^2 = .014$]. However, the effect size was small, as compared with the main effects of optimality and problem and the interaction between these factors; it was, thus, not analyzed further. No other interactions were significant.

Discussion

The main effect of optimality, found in both identification time and error data, shows that solution identification performance is a function of this factor. As Figure 3 illustrates, the participants were able to make accurate judgments about optimum solutions and about the suboptimality of solutions in the 45% range. With solutions around 15%–35% longer than the optimum, performance declined considerably. Comparing the curves for identification times and errors in Figure 3, there is no evidence of a speed–accuracy tradeoff in performance.

The data for the problem factor are slightly more complicated than those for optimality. Although there was no main effect of problem in the identification time data, this main effect was significant in the error data, although there is no evidence of a linear trend. Instead, it appears that performance was significantly more accurate with the one-internal point, four-internal point, and Dantzig problems than with the other problems.

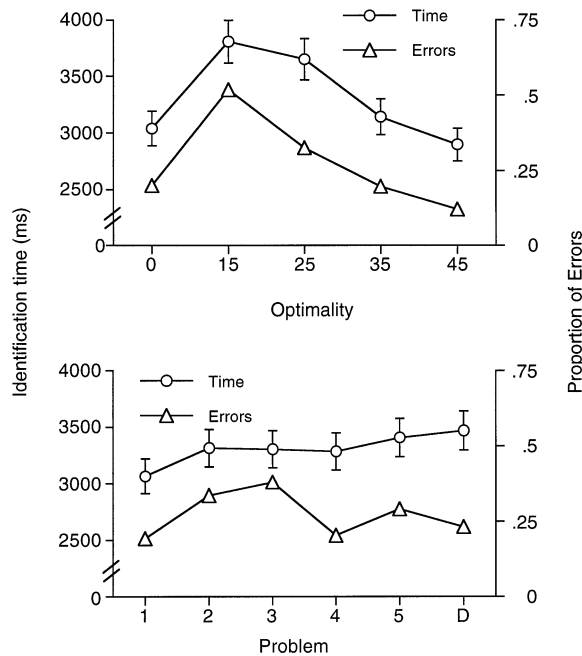


Figure 3. Mean identification times (in milliseconds) and proportion of errors for each level of optimality and problem in Experiment 2. Error bars are omitted for the error data, since they fall within the space taken by the plotting point in all cases.

The significant main effect of block in the identification time data is thought to reflect simple familiarization with the experimental task, rather than learning of the stimuli. If individual TSP stimuli had been learned by the participants, it would have been expected that the factor of block would interact with the other manipulated factors, so that the effects of optimality and problem lessened or disappeared across the levels of block. However, the two-way interactions between block and optimality and between block and problem were not significant. Although the three-way interaction between block, optimality, and problem was significant in the error data, visual inspection of the data does not indicate any consistent reduction in the effects of optimality and problem across levels of block, and in any case, the effect size is small, as compared with other significant effects. Further support for the notion that the differently oriented TSPs found across the levels of block were unlikely to have been seen as identical comes from the work of Rock, Schreiber, and Ro (1994), who found, using both recognition and forced-choice methods, that two-dimensional shape perception depended on orientation.

Taken together, the significant main effects of optimality are supportive of the hypothesis that participants derive a best figure for the presented TSP and compare it with the presented route to make a judgment about optimality. The interaction in the response time data between optimality and problem reveals a broadly similar effect of optimality across levels of problem: Each individual

problem gives rise to response times that are relatively low with both optimal and highly suboptimal solutions but that increase with solutions of intermediate optimality. However, it was noticeable that optimal solutions tended to be identified more quickly with simpler problems than with complex problems; this pattern was reversed for highly suboptimal solutions, which appeared to be identified more quickly with complex problems.

Although these findings are consistent with the hypothesis concerning comparison of the best figure with the given solution, they do not speak to the processes and mechanisms involved. Judgment of best figure in typical Gestalt examples can be considered to involve a combination of both global and local processes—thus, for example, in a circle whose perimeter consists of dots, both the global continuation and closure and the local proximity of dots are presumably important. It is possible, therefore, that the participants used either or both kinds of processes in deriving best figures from the presented TSPs, particularly given that they were able to spend as long as they wanted in making a response, the stimuli remaining visible until a response was selected. It is, hence, interesting to note an apparent bimodal distribution in the par-

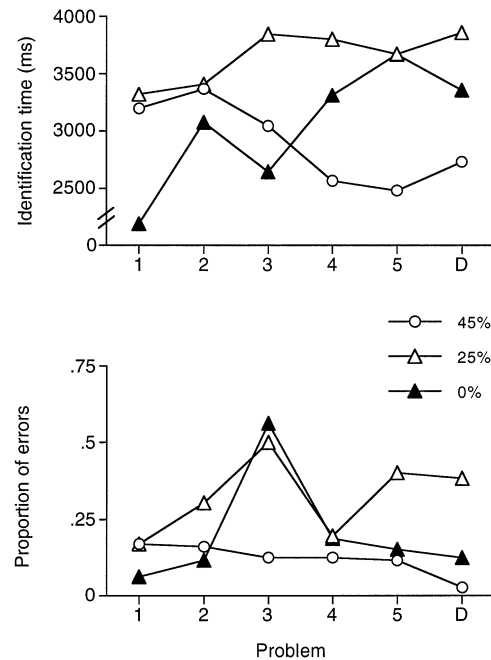


Figure 4. Mean identification times (in milliseconds) and proportion of errors for the optimality \times problem interaction in Experiment 2. The legend gives the approximate percentage distance over the optimal solution for each level of optimality. Optimality levels of 15% and 35% are omitted for clarity. Identification times for the 15% level follow those for the 25% level, although at a higher level, whereas identification times for the 35% level lie between those for the 25% and 45% levels. The error data for 15% and 35% levels follow those for the 25% level, with more errors for stimuli 15% above optimum and fewer for stimuli 35% above optimum.

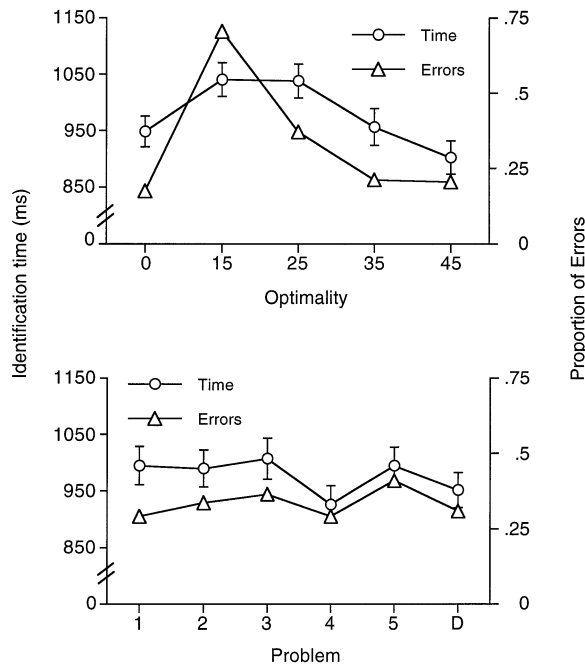


Figure 5. Mean identification times (in milliseconds) and proportion of errors for each level of optimality and problem factors in Experiment 3.

participants' overall mean identification times. Of 14 participants, 6 had overall mean identification times of between 1.6 and 2.3 sec, whereas the remaining 8 had overall mean identification times of between 3.8 and 5.7 sec. It is, therefore, plausible to suggest that the participants used different strategies to cope with the task: Perhaps, for example, a more global judgment of good figure by the faster participants, and a more detailed and local assessment and check of route segments by the slower participants. The following experiment took as its rationale the general finding of global precedence in the apprehension of complex forms (Kimchi, 1992) to suggest that, if only the former, global strategy were necessary and sufficient for the task, similar effects of optimality and problem should pertain even with the additional constraint of a forced-pace response.

EXPERIMENT 3

Method

Participants. Twenty-two undergraduate students at Lancaster University (9 women and 13 men, mean age 25.5 years) were paid £4 to participate.

Materials. The TSP stimuli were identical to those used in Experiment 1.

Design. There were two experimental factors, both repeated measures: optimality, with five levels from 0%–45%, and problem, with six levels (1 through 5 internal points, plus the Dantzig). The experimental measures were the same as those in Experiment 2.

Procedure. The procedure was exactly the same as that for Experiment 2, except that each stimulus was presented for only 2 sec, during which time the participant was required to make a response.

During the intertrial interval of 2 sec, the screen displayed the words *Break—no response* and keypresses made after removal of the stimulus were not recorded.

Results

Mean identification times and proportion of errors for each level of optimality and problem were calculated for each participant, and two-way ANOVAs were calculated for these data.

Identification times. A significant main effect of optimality was found [$F(4,84) = 3.976, MS_e = 120,547, p = .025, \omega^2 = .018$]. The main effect of problem was not significant [$F(5,105) = 1.51, MS_e = 71,627.3, p = .21, \omega^2 = .002$]. The mean identification times for each level of optimality and problem are shown in Figure 5. The interaction between optimality and problem (shown in Figure 6) was also significant [$F(20,420) = 3.099, MS_e = 75,574.5, p = .017, \omega^2 = .039$].

Errors. A significant main effect of optimality was found [$F(4,84) = 27.026, MS_e = 0.239, p = .001, \omega^2 = .168$]. The main effect of problem was not significant [$F(5,105) = 1.807, MS_e = 0.132, p = .136, \omega^2 = .004$]. The mean proportion of errors for each level of optimality and problem are shown in Figure 5. The two-way interaction between optimality and problem was not significant [$F(20,420) = 1.498, MS_e = 0.136, p = .152, \omega^2 = .009$].

Discussion

The results of Experiment 3 are remarkably consistent with those found in Experiment 2, despite the imposition of a forced-pace response regime. Again, the participants' performance is clearly a function of solution optimality, with high- and low-optimality solutions being identified faster and more accurately than solutions of intermediate optimality. The interaction between optimality and problem was similar to that found in Experiment 2, in that the optimal solutions were identified faster with simple problems, whereas suboptimal solutions were identified faster with complex problems.




Figure 6. Mean identification times (in milliseconds) for the optimality × problem interaction in Experiment 3. Legend and levels of optimality factors are as described in Figure 4.

The consistency and quality of performance in Experiment 3 under a forced-pace response regime is suggestive of global processing in the derivation of best figures for the presented TSPs. It is certainly unlikely that the participants could have judged the quality of solutions to 10-point TSPs through the application of a serial point-by-point strategy in identification times that were, on average, less than 1 sec. Nonetheless, the simple ability of the participants to make judgments about TSPs within the time limit, although consistent with the processing hypothesis under consideration, does not exclude the possibility that the efficient allocation of attentional resources or other strategic shifts by the participants might have permitted this pattern of performance, even with the time constraint.

Experiment 4 therefore adopted a different method, in order to examine more closely the processing involved in this TSP task. Sanocki (1993) presented a series of studies using a part-whole priming method, in which the prime stimuli consisted of subsets of features that were common to two alternative targets. Under a general feature model of object identification (e.g., Massaro & Friedman, 1990), such a prime should not facilitate identification and/or discrimination of the targets, as compared with a neutral prime, since the information contained in the prime is common to both targets. Sanocki (1993), however, found a consistent response time advantage for targets primed in this manner and interpreted these findings as evidence in favor of his contingency model of object identification, discussed earlier in this paper.

Applied to the present question relating to the TSP, such a model of processing allows predictions in terms of the presentation time of a prime stimulus. With no prime, judgment of TSP solution quality should proceed as in Experiments 2 and 3. With a short prime consisting simply of the problem (without a superimposed solution), a hypothesis about the complete solution could only be derived through global processing of the stimulus. If TSP solution judgments are based on a rapid global assessment of TSP best figure, followed by a comparison with the presented figure, it is expected that short prime durations will reduce the time taken to judge routes, since the global best figure is available before the route is presented. With increasing prime length, however, an increasing amount of contingent local processing can be undertaken. By analogy with Sanocki's (1993) Experiment 4, this is likely to interfere with a rapid decision, since slower local comparisons will be made between an hypothesis generated from the prime stimulus and the presented route. Alternatively, the long prime stimulus may allow sufficient time for an elaborated hypothesis about a solution to be developed, which then has to be discarded once the solution is presented for judgment and the original global figure recovered for comparison. In either scenario, it is likely that primes of long duration will give rise to an increase in the response times for decisions. Note that this goes strictly against an alternative local processing model

in which participants attempt to compute optimum TSP solutions on a point-by-point serial route, finding a basis through the application of some heuristic (such as nearest neighbor)—here, it would be predicted that the longer the time available at prime presentation, the more processing would be possible and the *faster* the decision.

EXPERIMENT 4

Method

Participants. Twenty-two undergraduate and postgraduate students volunteered to participate, and were paid £5 for their time. There were 13 men and 9 women, whose ages ranged from 19 to 26 years.

Materials. The six 10-point TSPs used for this experiment were identical to those used in Experiment 3. Prime stimuli consisted of the TSP nodes without superimposed solutions.

Design. There were three repeated measures experimental factors: prime, with seven levels (0, 100, 250, 500, 750, 1,000, and 2,000 msec); problem, with six levels (one through five internal points, plus Dantzig); and optimality, with five levels (path lengths 0%, 15%, 25%, 35%, and 45% longer than optimal). The experimental measures were (1) the solution identification times, measured from solution presentation onset to keypress, and (2) error percentage (*no* responses to 0% above-optimal stimuli and *yes* responses to any other level of optimality were counted as errors).

Procedure. The laboratory software package SuperLab was used to present stimuli and record response times and errors. Each participant was presented with 210 trials. In each trial, the prime stimulus was presented for an appropriate duration, followed immediately by the problem plus solution, which remained visible either until the participant's response or until a time limit of 2,000 msec was reached. However, in trials in which the level of prime was 0 msec, only the problem plus solution was presented. For each trial, the task of the participant was to decide whether or not the solution presented was an optimal tour round the points and to press a *yes* or *no* key, as appropriate. After each trial, a blank screen was presented for 2,000 msec before the subsequent trial. The trials were divided into seven blocks of 30; all the permutations of problem and optimality were presented in pseudorandom order within each block, and prime was varied across block, with the order of blocks being randomized anew for each participant. The participants were offered the opportunity to rest for up to 2 min between each block; the experiment took approximately 30 min to complete.

Results

Identification times were separated according to the correctness or incorrectness of the response. This process resulted in a large number of missing data points for some subjects; consequently, all the identification times (correct and incorrect together) were initially subjected to a three-way ANOVA, over the repeated measures factors of prime, problem, and optimality. There was no significant main effect of problem [$F(5,105) = 1.318$, $MS_e = 1.275 \times 10^5$, $p = .275$, $\omega^2 < .001$] nor any significant interaction of problem with prime [$F(30,630) = 1.288$, $MS_e = 6.906 \times 10^4$, $p = .232$, $\omega^2 < .001$] or with prime and optimality together [$F(120,2520) = 1.143$, $MS_e = 6.845 \times 10^4$, $p = .313$, $\omega^2 = .002$]. Subsequent analyses were, therefore, collapsed over the problem factor.

Correct identification times. Mean identification times associated with correct responses, pooled over the

problem factor, were subjected to a two-way ANOVA, with prime and optimality as repeated measures factors. Outliers more than 2.5 standard deviations from the mean were removed prior to analysis. The interaction between prime and optimality was not significant [$F(24,504) = 1.615, MS_e = 3.822 \times 10^4, p = .11, \omega^2 = .004$]. There were significant main effects of both prime [$F(6,126) = 3.858, MS_e = 1.224 \times 10^5, p < .01, \omega^2 = .026$] and optimality [$F(4,84) = 44.381, MS_e = 5.153 \times 10^4, p < .001, \omega^2 = .111$].

Mean identification times for each level of prime, both overall and by type of response, are shown in Figure 7. Post hoc pairwise comparisons, using Tukey's HSD test, showed that the following primes gave rise to significantly ($p < .05$) different identification times: 0 and 100 msec, 100 and 2,000 msec, 500 and 2,000 msec. Identification times varied with optimality, as is shown in Figure 8.

Errors. The number of errors, defined as either rejections of optimal routes or acceptances of suboptimal routes, made by each participant in each condition of the experiment were subjected to a two-way ANOVA, with prime and optimality as repeated measures. There was no significant interaction of prime and optimality [$F(24,504) = 1.111, MS_e = 0.023, p = .36, \omega^2 = .001$] and no significant main effect of prime [$F(6,126) = 2.033, MS_e = 0.038, p = .10, \omega^2 = .003$]. There was a significant main effect of optimality [$F(4,84) = 89.330, MS_e = 0.103, p < .001, \omega^2 = .512$], the nature of which is represented in Figure 8.

Discussion

Identification times varied with prime, so that they were shortest when the prime was 100 and 500 msec and longest when the prime was 0 and 2,000 msec. This variation in identification time was not connected with variability in the accuracy of performance: The number of errors was not significantly different across all levels of prime. This pattern is broadly consistent with the prediction that, if TSP solution judgments are based on a rapid global assessment of best figure, short primes will reduce identification times, since the global best figure is made available immediately before the route is presented. This provides strong evidence in favor of global processing

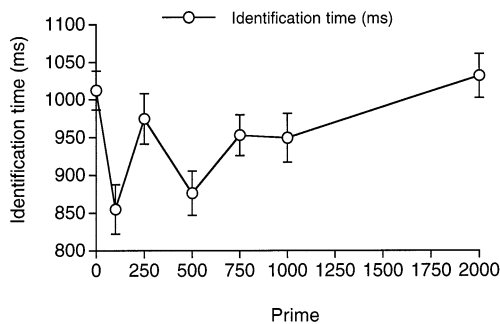


Figure 7. Mean identification times (in milliseconds) for each level of the prime factor in Experiment 4.

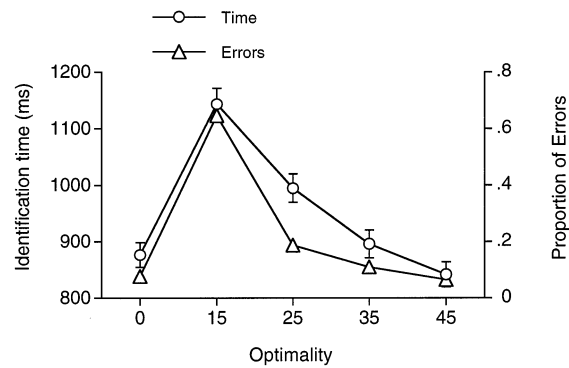


Figure 8. Mean identification times (in milliseconds) and proportion of errors for each level of optimality in Experiment 4.

and is not consistent with the alternative hypothesis that TSP solution judgments are based on the computation of a tour through the application of a point-by-point serial route-finding heuristic.

An unexpected discontinuity was found in the identification time data concerning the effect of prime length: Whereas primes of 100 and 500 msec showed facilitation of solution identification, relative to no prime or long primes, the 250-msec prime showed reduced facilitation. A number of explanations for this discontinuity are possible, but the simplest of these seems to be that it is a result of saccadic eye movements. Reaction times for saccadic eye movements of around 200 msec (see, e.g., Vaughan, 1983) are consistent with a pattern in which, after an initial fixation on the prime stimulus in which a global figure is extracted, a saccadic eye movement is made, followed by refixation on the prime stimulus in which the global figure is refreshed. Primes of 100 msec would be unaffected by saccadic eye movements, and primes of 500 msec would allow for an eye movement followed by refixation, whereas primes of 250 msec would include a saccadic eye movement without refixation, which might disrupt the maintenance of a global figure that was initially extracted from the prime.

The main effect of optimality replicated the findings of Experiments 2 and 3, in that identification times for optimal and 45% suboptimal solutions were short; for 15%, 25%, and 35% suboptimal solutions, identification times were longer, reaching a maximum at 15% suboptimal. Accuracy of identification followed a similar pattern, with near-perfect identification when optimality was optimal and 45% suboptimal, and chance levels of identification with 15% suboptimal solutions.

GENERAL DISCUSSION

Experiment 1 demonstrated a clear relationship between TSP solution optimality and judgments of good figure: The more optimal the TSP route, the better a figure it was judged to have. In Experiments 2 and 3, the participants quickly and reliably identified optimal so-

lutions and rejected those of low optimality; judgments about solutions with intermediate optimality took longer and were less accurate. Experiment 3 also demonstrated the participants' ability to make these judgments within a 2-sec time limit. The ability to differentiate reliably between levels of optimality in such short identification times is suggestive of global perceptual processing of figural properties in TSPs. Experiment 4 demonstrated that short prime stimuli, consisting of the problem alone, gave rise to facilitation of solution identification, as compared with both the no prime condition and conditions with longer prime stimuli. These findings are supportive of a model of TSP judgment in which a global good figure derived from the problem alone is judged against the presented route and a decision is made on the basis of the degree of match or mismatch.

The significant interactions found in all the experiments between optimality and problem factors indicate that it is somewhat easier to identify optimal solutions with simple TSPs and suboptimal solutions with more complex TSPs. This is also consistent with the proposed model: Optimal solutions to simple problems have the least disrupted convex hull of all the possible TSP solutions, and best figure may, therefore, be easily discovered. As problem complexity increases, the number of disruptions to the convex hull also increases for optimal problems; hence, the best figure may be more difficult to identify. Suboptimal solutions for simple TSPs have a disrupted convex hull, but major variations in optimality can be achieved without adding features, such as line cross-overs. In more complex problems, because of the increased number of internal points, it is much more common to find crossed lines and other indicators of solution suboptimality. This relates well to Sanocki's (1993) contingency model: Optimal solutions to simple problems are highly consistent with the emergence of a good figure that conforms to an uninterrupted convex hull, whereas suboptimal solutions to complex problems may be identified quickly via a clear mismatch between a global reference frame and the presence of figure-violating characteristics, such as crossed lines and sharp changes in line angle.

Note that, although the data from our experiments appear to be consistent with Sanocki's (1993) global-to-local contingency model, we do not make any specific claims about the precise nature of global processing that underlies TSP identification. In particular, the experiments were not designed to distinguish between the contingency and the general feature models reviewed by Sanocki, since variants of both models might, in principle, allow global processing of perceptual stimuli. In Experiment 4, the prime stimulus consisted of the complete set of points for each TSP, so that any model predicting the global extraction of a figure could account for our data. To test between contingency and general feature models in accounting for TSP solution identification would require the use of component parts of TSP problems as primes,

whose figure is consistent with optimal tours, suboptimal tours, or both.

Notwithstanding our caution in suggesting a specific global processing mechanism, it is interesting to speculate on the relationship between one of the characteristic properties of good TSP solutions, that of conforming to a convex hull, and the nature of contour detection in human vision. For example, Field, Hayes, and Hess (1993) have shown how the detection of paths within random Gabor patch arrays is a function of the orientation and position of elements, so that paths are detected where elements lie on a simple smooth curve passing between elements. The extraction of global structures based on the conforming of sets of local elements to curvilinear trajectories may be related to the key role of the convex hull, which typically offers the most obvious global curvilinear structure in a random dot TSP array. Not only are convex hulls strongly correlated with good TSP solutions, but they are also likely to dominate the elicitation of figures from TSP arrays, because of their significance in contour detection in human vision. One might predict that, where TSP arrays allow the extraction of curvilinear structures that do not conform to a convex hull, it will be these that influence TSP performance. For example, MacGregor and Ormerod (1996) suggest that an array consisting of two perpendicular curvilinear sets of dots is unlikely to yield high-quality TSP solutions.

Some evidence of localized figural effects was found across the different problem types. The error data for Experiments 2 and 3 show a similar pattern across different problems, performance on 1- and 4-point and Dantzig problems being more accurate than on 2-, 3-, and 5-point problems. In discussing the results of Experiment 2, it was suggested that the degree of clustering of internal points in each problem appears to be related to accuracy in identifying solution optimality. It is interesting to note that MacGregor and Ormerod (1996) also found evidence for effects of internal point proximity. They conducted a post hoc regression analysis of the response uncertainty data from the 20-point problems used in their second experiment and found that addition of interior point proximity as a second independent variable increased the variance accounted for by approximately 17%. Clusters of internal points may detract from the global figure of the convex hull, thereby impairing the identification of TSP solution quality. However, this clustering hypothesis should be treated with caution: Only a small set of problems were used in the experiments, and also a number of competing hypotheses remain untested. For example, it may be that the regularity of internal point patterns, rather than their degree of clustering, is more important in determining problem complexity. Experiments are currently in progress to test these competing hypotheses concerning localized figural effects.

The outcomes of this research have two main practical implications. First, they may usefully inform the efforts

of operations researchers in the development of practical methods for solving large-scale TSPs. In particular, the allocation of tasks adopted by Krolak et al. (1971), in which they used computer-based methods to find approximate tours of TSPs whose details were subsequently completed by human problem-solvers, may be in precisely the opposite direction to the most effective allocation. Our results suggest that the unique capability of humans at solving TSPs lies in the rapid identification of good figures. It may even be that requiring humans to adjust partially completed tours encourages them to adopt local processing strategies that interfere with, or nullify, the identification of a TSP's global figural properties.

Second, global perceptual processing may be used by humans in solving problems that are homomorphic to the TSP—for example, in route finding (see, e.g., Anderson, Kushmerick, & Lebiere, 1993) and the statistical assessment of data distributions (see, e.g., Lloyd & Steinke, 1977; Walter, 1993). For example, Anderson et al. investigated how humans choose between routes that vary according to known and future costs, as a test of Anderson's (1993) ACT-R theory. In their experiments, participants navigated a spatial array from start to target locations via a number of intermediate locations. The cost of alternative routes was manipulated, so that fast travel was possible on direct routes, whereas slow walking could be undertaken where no direct route existed. Unlike a realistic map-reading task in which all the routes between locations are displayed, participants saw only the direct routes radiating from their present location at any stage. This enabled Anderson et al. to examine the factors underlying decisions to choose one route over another at intermediate locations in a tour. Their data suggest the application of a local processing strategy, in which participants evaluated the cost of taking alternative routes between pairs of locations as a function of the Euclidean distances between present, future, and target locations.

Our TSP stimuli did not include cost information about each location and route. Yet, real-world TSPs often require cost-benefit analyses. For example, in planning a trip around a number of outlets, a distributor needs to be aware that some outlets' needs are more urgent than others or that some outlets may be closed at certain times. Pilot data from a master's dissertation (Buckmaster, 1992) showed that irrelevant contextual information (in this case, superimposing TSP points on a background outline of the United States) reduced the quality of participants' TSP solutions to approximately the level of the nearest neighbor heuristic, suggesting a switch from global to local processing strategies. Providing relevant contextual information, such as costs associated with each location or route, might be expected to impair TSP performance still further. We are presently exploring a possible remedy to this problem, in which the costs associated with each location or route in a TSP are abstracted from the topographical display of locations and are themselves represented as a TSP in which location in a spatial array is a function of cost. In an inversion of Krolak et al.'s (1971) approach, we conceive that humans may produce solu-

tions to multiple TSPs (one for each representation of cost or topography), which can then be combined computationally to determine the most efficient overall tour.

Notwithstanding the importance of cost information, an essential characteristic of TSPs and their homomorphs is that they have visual displays showing the relative spatial distribution of points. Although ACT-R gives a good account of performance in Anderson et al.'s (1993) unusual route-finding paradigm, we argue that the restricted displays that their participants saw precluded the application of strategies for extracting figural information. We suggest that perceptual processing of the figural properties inherent in a visual display may complement route-finding performance when complete visual displays of routes between locations are provided. Unitary cognitive architectures, such as ACT-R and SOAR (Newell, 1990), do not address the influence that global perceptual processing has on problem-solving performance. Others (e.g., Larkin, 1989; Zhang & Norman, 1994) have recognized the importance of external displays in mediating cognition. However, they focus on explaining how external displays constrain the operation of local processing strategies. We suggest that recognizing the role of global perceptual processing in extracting figural properties from stimulus displays will form an important extension to existing theories of human problem solving.

REFERENCES

- ANDERSON, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- ANDERSON, J. R., KUSHMERICK, N., & LEBIERE, C. (1993). Navigation and conflict resolution. In J. R. Anderson (Ed.), *Rules of the mind* (pp. 93-120). Hillsdale, NJ: Erlbaum.
- BUCKMASTER, J. K. (1992). *Evaluating human performance on the travelling salesman problem using a HyperCard based system*. Unpublished master's thesis, Loughborough University.
- DANTZIG, G. B., FULKERSON, D. R., & JOHNSON, S. M. (1959). On a linear-programming, combinatorial approach to the travelling-salesman problem. *Operations Research*, 7, 58-66.
- DODD, D. H., & SCHULTZ, R. F. (1973). Computational procedures for estimating magnitude of effect for some analysis of variance designs. *Psychological Bulletin*, 79, 391-395.
- FIELD, D. J., HAYES, A., & HESS, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field." *Vision Research*, 33, 173-193.
- GAREY, M. R., & JOHNSON, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. San Francisco: Freeman.
- GOLDEN, B., BODIN, L., DOYLE, T., & STEWART, W. (1980). Approximate travelling salesman algorithms. *Operations Research*, 28, 694-721.
- KIMCHI, R. (1992). Primacy of wholistic processing and the global/local paradigm: A critical review. *Psychological Bulletin*, 112, 24-38.
- KROLAK, P., FELTS, W., & MARBLE, G. (1971). A man-machine approach toward solving the travelling salesman problem. *Communications of the ACM*, 14, 327-334.
- LARKIN, J. H. (1989). Display-based problem solving. In D. Klahr & J. K. Kotovsky (Eds.), *Complex information processing* (pp. 319-341). Hillsdale, NJ: Erlbaum.
- LLOYD, R., & STEINKE, T. (1977). Visual and statistical comparison of choropleth maps. *Annals of the Association of American Geographers*, 67, 429-463.
- MACGREGOR, J. N., & ORMEROD, T. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics*, 58, 527-539.
- MASSARO, D. W., & FRIEDMAN, D. (1990). Models of integration given multiple sources of information. *Psychological Review*, 97, 225-252.
- NEWELL, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

- POLIVANOVA, N. I. (1974). On some functional and structural features of the visual-intuitive components of a problem-solving process. *Voprosy Psikhologii* [Questions of Psychology], **4**, 41-51.
- POMERANTZ, J. R. (1981). Perceptual organization and information processing. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 141-180). Hillsdale, NJ: Erlbaum.
- ROCK, I., SCHREIBER, C., & RO, T. (1994). The dependence of two-dimensional shape perception on orientation. *Perception*, **23**, 1409-1426.
- ROSENKRANTZ, D., STERANS, R., & LEWIS, P. (1974). Approximate algorithms for the travelling salesperson problem. In *15th Annual IEEE Symposium of Switching and Automata Theory* (pp. 33-42). New York: IEEE Computer Society, Publications Office.
- SANGALLI, A. (1992, December 12). Why sales reps pose a hard problem. *New Scientist*, pp. 24-28.
- SANOCKI, T. (1991). Effects of early common features on form perception. *Perception & Psychophysics*, **50**, 490-497.
- SANOCKI, T. (1993). Time course of object identification: Evidence for a global-to-local contingency. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 878-898.
- VAUGHAN, J. (1983). Saccadic reaction time in visual search. In K. Rayner (Ed.), *Eye movements in reading* (pp. 397-412). New York: Academic Press.
- WALTER, S. D. (1993). Visual and statistical assessment of spatial clustering in mapped data. *Statistics in Medicine*, **12**, 1275-1291.
- WINER, B. J., BROWN, D. R., & MICHELS, K. M. (1991). *Statistical principles in experimental design* (3rd ed.). New York: McGraw-Hill.
- ZHANG, J., & NORMAN, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, **18**, 87-122.

(Manuscript received September 17, 1996;
revision accepted for publication June 23, 1998.)