

The roles of the convex hull and the number of potential intersections in performance on visually presented traveling salesperson problems

DOUGLAS VICKERS, MICHAEL D. LEE, MATTHEW DRY, and PETER HUGHES
University of Adelaide, Adelaide, Australia

The planar Euclidean version of the traveling salesperson problem requires finding the shortest tour through a two-dimensional array of points. MacGregor and Ormerod (1996) have suggested that people solve such problems by using a global-to-local perceptual organizing process based on the convex hull of the array. We review evidence for and against this idea, before considering an alternative, local-to-global perceptual process, based on the rapid automatic identification of nearest neighbors. We compare these approaches in an experiment in which the effects of number of convex hull points and number of potential intersections on solution performance are measured. Performance worsened with more points on the convex hull and with fewer potential intersections. A measure of response uncertainty was unaffected by the number of convex hull points but increased with fewer potential intersections. We discuss a possible interpretation of these results in terms of a hierarchical solution process based on linking nearest neighbor clusters.

In both the everyday world and the psychological laboratory, the ability to solve practical problems is generally regarded as a centrally important cognitive activity and a direct expression of intelligence (Resnick & Glaser, 1976; Sternberg, 1982). Over the years, however, research perspectives on problem solving have undergone considerable change (Lovett, 2002). For example, early work in the Gestalt tradition focused on problems that required a single insightful step leading to a restructuring of the representation of the problem, resembling perceptual reorganization (Duncan, 1959). Later work continued to study *knowledge-lean* problems, requiring little or no background knowledge, but placed more emphasis on sustained cognitive strategies in multistep problems (Newell & Simon, 1972). Such problems typically have a unique answer that can be arrived at by a systematic, finite series of steps. More recently, interest has shifted to *knowledge-rich* problem domains, such as chess and medical diagnosis, and has paid particular attention to the nature and acquisition of expertise (e.g., Chi, Glaser, & Farr, 1988).

There is some tension between these different perspectives. On the one hand, for example, problems that have received the most attention in knowledge-lean laboratory studies (such as the missionaries and cannibals and the tower of Hanoi problems) have little claim to being representative of real-world problem solving. On the other hand,

research into real-world problem solving has emphasized the specificity of expertise, the limited transfer to a different domain, and the differences in problem-solving performances across different domains (Ericsson & Lehmann, 1996).

This tension brings out the potential importance of an important class of hitherto neglected problems. These are computationally intractable problems for which there is no algorithm that can be guaranteed to produce a definitive solution within a practicable time. Such problems lend themselves to laboratory investigation because they require no specific background knowledge, they can often be stated simply, and they are easily understood. At the same time, however, they have some claim to being representative of many real-world practical situations, such as those encountered in operations research. In particular, there has recently been a growing interest in human performance on a well-known example of such computationally infeasible problems, the so-called traveling salesperson problem (TSP), traditionally referred to as the traveling salesman problem (Graham, Joshi, & Pizlo, 2000; M. D. Lee & Vickers, 2000; MacGregor & Ormerod, 1996; MacGregor, Ormerod, & Chronicle, 1999, 2000; Vickers, Butavicius, Lee, & Medvedev, 2001).

The classic planar Euclidean TSP can be formulated as follows: Given n interconnected cities (represented by points in the plane), devise an itinerary that will visit each city exactly once, return to the initial city, and keep the total distance traveled to a minimum. To arrive at an optimal solution by complete enumeration means considering $(n - 1)!/2$ itineraries. Although this is feasible when n is a small number, such as 5, and only 12 different routes have to be considered, it rapidly becomes impossible as n

The research reported in this paper was supported by Australian Research Council Grant DP0211150 to D.V. and M.D.L. We thank Michael Brusco, Evan Heit, and several anonymous reviewers for helpful comments. Correspondence concerning this article should be addressed to D. Vickers, Department of Psychology, University of Adelaide, Adelaide, SA 5005 Australia (e-mail: douglas.vickers@psychologyadelaide.edu.au).

becomes even moderately large. The TSP is one of an important class of so-called NP-complete problems, for which it is believed there is no algorithm that can be guaranteed to arrive at an optimal solution within a reasonable, polynomial time (Lawler, Lenstra, Rinooy Kan, & Schmoys, 1985; Wilf, 1986).

In this article, we compare two approaches to understanding how people solve visually presented TSPs. The first assumes a global-to-local strategy, in which solutions are guided by the convex hull¹ of the stimulus array (MacGregor & Ormerod, 1996; MacGregor et al., 2000). We examine theoretical arguments and empirical evidence for and against this hypothesis, before describing an alternative, local-to-global hypothesis that takes into account the observation that observers appear to avoid intersecting paths (MacGregor et al., 2000; van Rooij, Stege, & Schactman, 2003). We then compare these two approaches in an experiment in which the different effects the number of points on the convex hull and the number of potential intersections have on human performance are measured.

A Global-to-Local Process Based on the Convex Hull

Despite the difficulty of TSP problems, most studies have reported that human performance on visual versions compares well with computational procedures (Graham et al., 2000; M. D. Lee & Vickers, 2000; MacGregor & Ormerod, 1996; MacGregor et al., 1999, 2000; Polivanova, 1974; Vickers et al., 2001). For example, MacGregor and Ormerod found that the performance of observers was so close to the best-known solutions that there were no individual differences and there was zero correlation between performance across different problems. On these grounds, MacGregor and Ormerod concluded that the solution processes employed correspond to spontaneous, uniformly present organizing tendencies of the visual system. More specifically, they went on to suggest that observers make use of an innate perceptual preference for organizing stimulus elements into convex figures, as is suggested by Gestalt analyses of figure-ground differentiation in perception (Palmer, 1999, pp. 282–283). MacGregor et al. (2000) implemented this proposal in a computational model that uses the convex hull to establish a subtour, into which the remaining nodes (cities) are then inserted serially. This strategy, they claimed, would allow observers to exploit the principle, stated by Flood (1956), that the optimal solution will always connect adjacent nodes on the convex hull in sequence, even though it also passes through interior nodes in the process.

At first sight, there seems to be considerable evidence that observers solve TSPs by using the convex hull to establish a subtour and then serially inserting the interior points. For example, MacGregor and Ormerod (1996) found high correlations between the number of interior points (which varies inversely with the number of points on the convex hull) and both the number of different solutions generated and a response uncertainty measure that sum-

marized the frequency with which different observers connected the same nodes in their solutions. Observers' solutions had significantly fewer indentations than would be expected, and almost all connected adjacent boundary points in order. This was interpreted as evidence that not only do observers have a preference for connecting points on the convex hull in *sequence*, but they also have a preference for connecting adjacent points on the boundary *to each other*. Reanalyzing results for 13 problems, obtained by R. K. L. Lee (1985), MacGregor and Ormerod also found that there was a significant correlation between the number of interior points and the mean percentage by which observers' total path lengths exceeded the optimal. In addition, MacGregor et al. (2000) found that only 11 out of 455 solutions had lines that intersected. According to MacGregor et al. (2000), "[t]hese consistent elements argue that a relatively uniform process underlies performance. Furthermore, the results suggest the general form of that process to be similar to a convex hull approach" (p. 1184).

Despite the apparent generality of these conclusions, M. D. Lee and Vickers (2000) point to certain consequential restrictions. As MacGregor and Ormerod (1996) acknowledged, the circular form of their problems might "implicitly have prompted subjects to pay attention to the boundary of the convex hull" (p. 537). The likelihood of this is increased by the fact that the probability that a random configuration of nonboundary points would fall entirely within the prescribed areas is negligibly small in most cases. The importance of the circular form is further increased by M. D. Lee and Vickers's demonstration that the probability of having a large number of points on the convex hull of a 10- or 20-point array is also extremely small. In the view of M. D. Lee and Vickers, these combined constraints would have increased the saliency of the convex hull, thereby artifactually biasing the results in favor of MacGregor and Ormerod's hypothesis.

In reply to this criticism, MacGregor and Ormerod (2000) point out that MacGregor and Ormerod (1996) also reanalyzed previously unpublished data from R. K. L. Lee (1985) for 13 TSPs, ranging from 10 to 60 nodes, that had been randomly generated and that "the pattern of results was consistent with those obtained with [the] constrained TSPs" (p. 1502). However, what MacGregor and Ormerod (1996) actually reported was that, for R. K. L. Lee's data, there was a correlation of $r = .84$ between the number of interior points and the percentage over the optimal for the human solutions.

There are two problems with this. First, if these arrays were indeed randomly generated, the number of interior points would be confounded with the total number of points in each array. The finding that performance declines as a function of the total number of points would be predicted by most models of human performance on TSPs. Second, as was implied by MacGregor and Ormerod (1996) and as was confirmed by subsequent analysis by the present authors, there was no significant correlation between the number of interior points and the mean per-

centage by which total path length exceeded the optimal in either of their own two experiments. Thus, the two patterns of results may be consistent (if we assume that MacGregor & Ormerod's, 1996, results were constrained by ceiling effects but that those of R. K. L. Lee, 1985, were not). However, the two patterns are certainly not identical.

A further underlying problem is illustrated by the finding that a very high proportion of observers "adhered to [the] principle" of connecting the points on the convex hull in order, which is cited by MacGregor and Ormerod (1996, p. 533) and MacGregor et al. (2000, p. 1184) as "consistent with the hypothesis" that observers follow this as a procedure for arriving at a solution. However, the observers produced optimal or near-optimal performance on all problems. Thus, the finding that observers tended to follow the points on the convex hull sequentially is tantamount to an empirical verification of the general principle that the optimal path for every Euclidean TSP will connect the convex hull points in order. Exactly the same result is produced by a number of algorithms that produce optimal or near-optimal solutions but that incorporate very different processes, both from each other and from algorithms based on the convex hull.

This problem is most directly evident in the relationship between intersections and convex hull points. Quintas and Supnick (1965) have provided proofs of the two "folk theorems" that (1) in the Euclidean plane, the optimal tour cannot intersect itself and (2) if a tour does not visit the vertices of the convex hull in order, it is not optimal. Their proof of the second theorem (Quintas & Supnick, 1965, p. 979) is essentially that a tour not following the convex hull has an intersection and so, by the result established in the first theorem, cannot be optimal. In this way, Quintas and Supnick showed that, by negation of the consequent, if a tour does not intersect itself, it must follow the convex hull. This means that the avoidance of intersections is not a further independent piece of evidence that supports the convex hull hypothesis of MacGregor and Ormerod (1996). Instead, as has been explained by van Rooij et al. (2003), the fact that observers generally follow the convex hull nodes in sequence logically follows from their avoidance of intersections.

A Local-to-Global Process Based on Nearest Neighbors

In view of the above strictures, it is desirable to formulate alternative approaches to understanding human performance on visually presented TSPs. One possibility is that observers do indeed exploit spontaneous perceptual organization, as was suggested by MacGregor and Ormerod (1996), but that this is a local-to-global process, based on the identification, by an automatic parallel process, of links between nodes that are nearest neighbors (NNs).

A hypothesis of this kind has a number of advantages. First, because NN links constitute least distances and cannot intersect, a solution procedure based on NNs would result in few, or no, intersections. This is consistent with the finding by MacGregor and Ormerod (1996) and van

Rooij et al. (2003) that only a very small proportion of human solutions to TSP problems contain intersections. Second, as has been argued by van Rooij et al., it is more elementary than the convex hull hypothesis because, if a tour does not intersect itself, it must follow the convex hull (Quintas & Supnick, 1965). Third, as has been pointed out by Graham et al. (2000), a convex hull strategy may be potentially useful in the Euclidean TSP, but it has no obvious application to other visually presented optimization problems. Such problems include *open* TSP tasks, in which the aim is to find the shortest path through a number of nodes between two distinct starting and finishing nodes. They also include the minimum spanning tree (MST) and the Steiner tree (ST) problems. The MST is a spanning tree (a structure that contains no loops and links all n nodes of a pattern with $n - 1$ edges), for which the sum of the edge lengths is a minimum (Johnsonbaugh, 2001). The ST similarly connects all of the nodes, using a minimal sum of edge lengths, but allows additional nodes to be introduced into the graph (Hwang, Richards, & Winter, 1992). In contrast to the convex hull strategy, it is arguable that a procedure based on the detection of NNs is potentially useful in all of these problems.

Within the field of visual perception, there is a growing amount of evidence in support of a hypothesis of this kind. For example, Caelli (1981) was the first to argue that the rated quality of perceived structure in Glass patterns (Glass, 1969) could be explained in terms of differences between the distributions of NN distances of corresponding and noncorresponding dots. Pomerantz (1981) suggested that simple dot patterns were seen as connected by the shortest path and presented several examples, all of which corresponded to MSTs composed solely of NN links. Similarly, Dawson (1991) identified NN vectors as a main factor determining apparent motion in the Ternus display and its variants. More recently, Vickers, Dry, Lee, and Hughes (2003) have identified first and second NN links as accounting for around 90% of the variance in discriminative responses to Glass patterns. Vickers, Preiss, and Hughes (2003) have found that the number and magnitude of links and clusters seen in random-dot patterns is closely predicted by the number and magnitude of NN links and clusters and that the extent of motion perceived in random-dot kinematograms varies with the size of the mean NN distance. Vickers and Preiss (2003) have shown that apparently paradoxical impressions of coherent motion, reported by Ross, Badcock, and Hayes (2000), from sequences of independently generated Glass patterns can also be explained in terms of constraints on the vectors linking NNs in successive patterns.

Of immediate relevance in the present context, Vickers, Bovet, Lee, and Hughes (2003) found that solution length for both the open and the closed versions of the TSP was reliably correlated and, in both cases, was well predicted by the proportion of NN links that a standard NN algorithm (Golden, Bodin, Doyle, & Stewart, 1980) shared with the benchmark (best-known) solution. Further evidence that a common process, based on NN links, under-

lies the solution processes to a range of visually presented optimization problems is provided by the finding that solution lengths for each of a TSP, an MST, and an ST task were all three highly correlated with a measure of path complexity that indexed whether the nodes of each link in a participant's solution were connected to the nearest, second nearest, or n th nearest node (Vickers, Mayo, Heitmann, Lee, & Hughes, in press).

The perception of Glass patterns and of motion in random-dot kinematograms takes place very quickly and automatically, and it seems unlikely that the detection of NNs is the result of a serial, cognitively directed process. Instead, it seems more plausible to suppose that NN detection is a parallel process made possible by the wide interconnection of elements in the visual field. Because the principal parameter determining NN distances is dot density and because dot density has been suggested as the property that the visual system uses to implement scale invariance in the perception of symmetry and Glass structure, it is highly probable that NN relations are a critical factor in receptive field plasticity (Motter & Belky, 1998; Rainville & Kingdom, 2002). For the present, the precise mechanism underlying NN detection remains open to speculation. However, one tentative conception is that it might take the form of a neurophysiological realization of effective regions of influence around each node in a pattern, similar to those defined by Voronoi diagrams (O'Rourke, 1994).

The important point in the present context is that the rapid automatic identification of NNs may serve to exclude substantial numbers of solution paths. In particular, because no two edges connecting NNs can intersect, a reliance on NNs should restrict the search for TSP solutions to pathways that contain no crossings. On this view, the almost universal absence of intersections in human TSP solutions is a necessary consequence of an automatic local-to-global perceptual organization process, rather than a deliberate avoidance. A close parallel is the absence of intersecting motion paths between random-dot patterns presented in rapid succession (Attneave, 1974; Kolers, 1972; Navon, 1976; Ullman, 1979; Vickers, Preiss, & Hughes, 2003).

EXPERIMENT

In addition to the strictures on the evidence cited in favor of the convex hull hypothesis, M. D. Lee and Vickers (2000) have drawn attention to the constraint that, when TSP configurations are constructed randomly, the variation in the number of nodes that fall on the convex hull is relatively small and both the variance and the mean number of such nodes exhibit asymptotic behavior. As a result, as configurations of increasing size are considered, the concept *number of interior or nonboundary nodes* becomes equivalent to *total number of nodes*. However, as MacGregor and Ormerod (2000) correctly state in their reply, "it is possible to construct problems of any size in

which the total number of points and the number of non-boundary points vary orthogonally. For example, this could be done by selecting problems from the upper and lower ranges of the distributions of nonboundary points shown in Lee and Vickers's Figure 1" (p. 1052).

According to MacGregor and Ormerod's (1996) convex hull hypothesis, "the complexity of a TSP should depend not so much on the total number of points, but on the number of nonboundary, or interior, points" (p. 528). Therefore, the convex hull hypothesis predicts that a manipulation of this kind, in which the number of boundary points is increased (and the number of interior points decreased), should facilitate the solution process and lead to shorter path lengths and to reduced response uncertainty.

On the other hand, if human TSP solutions depend on the automatic detection of NN links, arrays that contain a large number of potential intersections will impose more constraint on the solution process. Because intersections identify longer pathways and any process that effectively excludes intersections reduces the number of permissible alternative pathways, an increased number of potential intersections should lead to shorter solutions and to reduced response uncertainty.

In order to investigate the relative importance of these two factors and to evaluate these two hypotheses, an experiment was carried out, closely conforming to the above suggestion by MacGregor and Ormerod (2000), in which the two factors, identified by the convex hull and the NN hypotheses as important determinants of performance, were varied in a factorial design. The first factor was the number of nodes on the convex hull. The second factor was the number of potential intersections in an array. Following MacGregor and Ormerod's (2000, p. 1502) suggestion, these factors were varied by selecting problem instances from a randomly generated sample of arrays.

Method

Participants. There were 20 participants, the majority of whom (19 out of 20) either were studying or had studied at a tertiary level. Their ages ranged from 17 to 40 years, with a mean age of 23.

Stimuli. The test arrays were drawn from a pool of 20,000 randomly generated 40-node arrays and were selected according to the number of nodes they had on the convex hull and their number of potential intersections, as calculated by counting the number of intersections over all possible pairs of line segments.

Figure 1 shows each of the 20,000 randomly generated arrays in terms of their number of nodes on the convex hull and number of intersections, together with the marginal distributions for these two measures. By using this information, nine problem types were defined, following a 3 (high, medium, and low number of points on the convex hull) \times 3 (high, medium, and low number of intersections) factorial design.

The mean number of intersections was 63,542, and this was designated as the medium level, with high (66,341 intersections) and low (60,743 intersections) selected to be 1.5 standard deviations (2,799 intersections) above and below the mean, respectively (tolerance for selection was $\pm 5\%$ of the standard deviation, or ± 93 intersections). The mean number of nodes on the convex hull was calculated in a similar manner ($M = 9.61$, $SD = 1.587$). The medium number of nodes on the convex hull was selected to be either 9 or 10

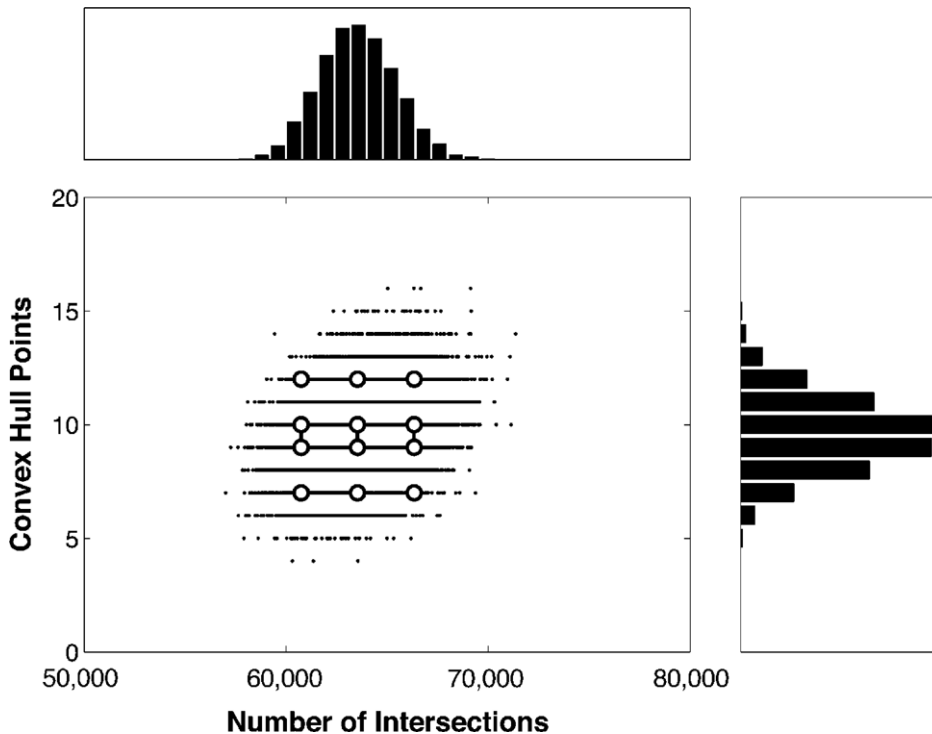


Figure 1. The selection of experimental stimuli. The black points represent 20,000 randomly generated arrays, characterized in terms of the number of points on the convex hull and the number of intersections. The histograms show the marginal distributions for these two measures. The 3×3 grid of white markers indicate the low, medium, and high convex hull and intersection values used to select the experimental stimuli.

points (alternated randomly), with a low number of nodes as 7, and a high number as 12. The locations of the nine problem types, showing the 3×3 grid, are displayed as white markers (or a pair of white markers, in the case of medium convex hull points) in Figure 1. Eight arrays from the pool of 20,000 that met the criteria were chosen at random for each of the nine problem types.

Procedure. The participants attended two test sessions. Each test session ran for approximately 2.5 h. The participants were given an hour to complete 18 problems. After a half-hour break, the observers completed another 18 problems. Each test session contained two of each of the nine types of problem, and within each session, the arrays were presented to the observers in a quasirandom manner.

The arrays were presented on computer color monitors. In each array, the nodes were distributed within a square field 15×15 cm. The participants could begin at any point by left-clicking on a node (or on a location where they wished to establish a new node) with the computer mouse. Then, while holding down the mouse button, they drew a path by dragging the mouse cursor to a subsequent node and releasing the button, causing a straight line to be drawn between that node and the previously visited node. By right-clicking on a link (or created node) to select it and then pressing the Delete key on the keyboard, the participants could undo any links or nodes they had drawn. The participants were thus free to connect the nodes in any order, to work alternately from two nodes, or to work on several separated clusters of nodes.

Immediately prior to commencing each test session, the participants were presented with two examples of possible solutions to a TSP, one of which contained an intersection. To maintain motivation and to encourage the participants to produce their best solutions, after each problem was completed, they were presented with the

benchmark solution and a score of the percentage difference between their own solution and the benchmark.

Results

Deviations from optimality. Benchmark optimal or near-optimal solutions to the test arrays were calculated by means of a simulated annealing heuristic (Press, 1992; Reinelt, 1994). To permit comparison across the different problem types, the observers' pathway lengths were expressed as percentages above the benchmark score. Thus, an observer's score for a perfect pathway would be zero, and any pathway that exceeded the benchmark would have a positive score.

Figure 2 shows the change in the deviation from optimality measure, averaged across all eight arrays for each problem type, in terms of the convex hull and the intersection factors. Surprisingly and contrary to both the findings and the convex hull hypothesis of MacGregor and Ormerod (1996), mean scores increased as a direct function of the number of nodes on the convex hull. Conversely, mean scores decreased as the number of intersections increased.

Because of the shortcomings of null hypothesis significance testing (NHST) (e.g., Cohen, 1994; Edwards, Lindman, & Savage, 1963; Howson & Urbach, 1993; Hunter, 1997; Lindley, 1972), we used the Bayesian approach to

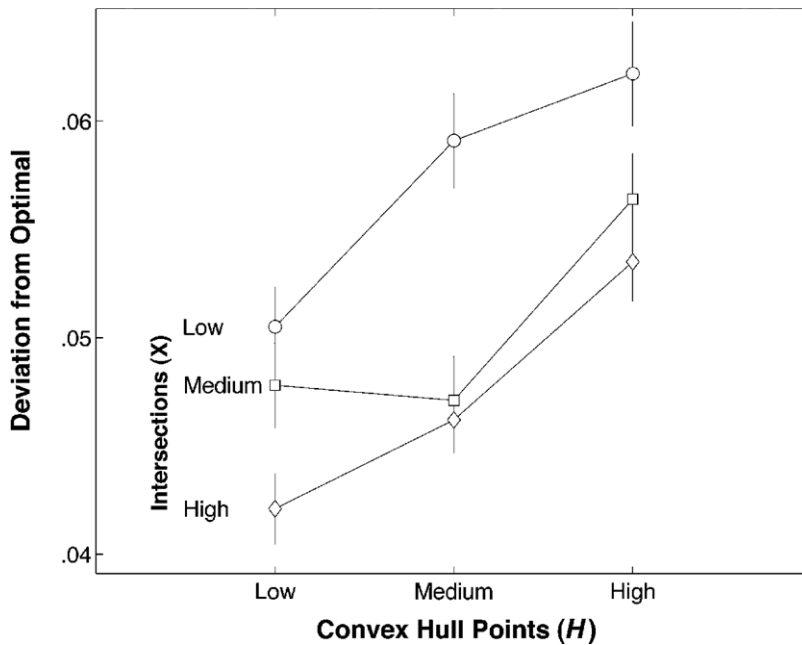


Figure 2. The mean deviation from optimality, taken across all the participants' solutions to all eight problems, for each combination of points on the convex hull and number of intersections. Deviation from optimality is measured as the proportion by which a solution length exceeds the benchmark optimal value. The error bars show one standard error in each direction.

statistical inference (e.g., Carlin & Louis, 2000; Gelman, Carlin, Stern, & Rubin, 1995; Kass & Raftery, 1995; Leonard & Hsu, 1999; Sivia, 1996) to reach conclusions about the influence that the number of points on the convex hull and the number of potential intersections have on performance. This analysis compared six competing models that made different assumptions about the effects of the factors and the form of their interaction. The first model, called *full*, assumes that both factors affect performance and that they interact in a completely unconstrained way. This model requires that each of the $3 \times 3 = 9$ means in the data requires a separate parameter. At the other end of the spectrum, the second model, called *constant*, assumes that neither factor has any effect on performance, and so all of the means are captured by the same single parameter. The third model, *intersect*, assumes that only changes in the number of intersections affect performance, whereas the fourth model, *hull*, assumes that only changes in the number of points on the convex hull affect performance. Both of these models require three parameters, one for each level of the factor that is assumed to affect performance. The fifth model, *one-way*, assumes that both factors affect performance in an entirely independent way, so that the change in one does not depend on the value of the other. This model involves four parameters: two to describe the rates of change for each factor and two to describe the differences in their initial values. The sixth and final model, *two-way*, extends one-way by assuming

that there is some limited dependency. In particular, it assumes that each change in a factor affects performance by a constant amount but that the value of this constant depends on the level of the other factors. The two-way model is more constrained than the full model, requiring six parameters: three to describe the rate of change for each of the two factors.

Figure 3 shows the best fits to deviation from optimality under each of the six models. From these data fits and the known parametric complexity of the models, it is possible to calculate the Bayesian information criterion (BIC: Schwarz, 1978) for each model. The BIC values, in turn, allow Bayes factors between each pair of models to be estimated (e.g., Kass & Raftery, 1995), which quantify how much more likely one model is than another. Table 1 summarizes these results, showing the log-likelihood measure of data fit, the number of parameters, the BIC, and the Bayes factors for each model. The Bayes factor is taken with respect to the most likely model (i.e., the one with the minimum BIC value), which is the one-way model.

What the Bayesian analysis shows is that the data provide the most evidence for the one-way model, with the hull model being 3.97 times less likely, the intersect model being 5.43 times less likely, and so on. The data provide the least evidence for those models that assume some form of interaction between the two factors. Our conclusion is that both the number of points on the convex hull and the number of intersections influence performance

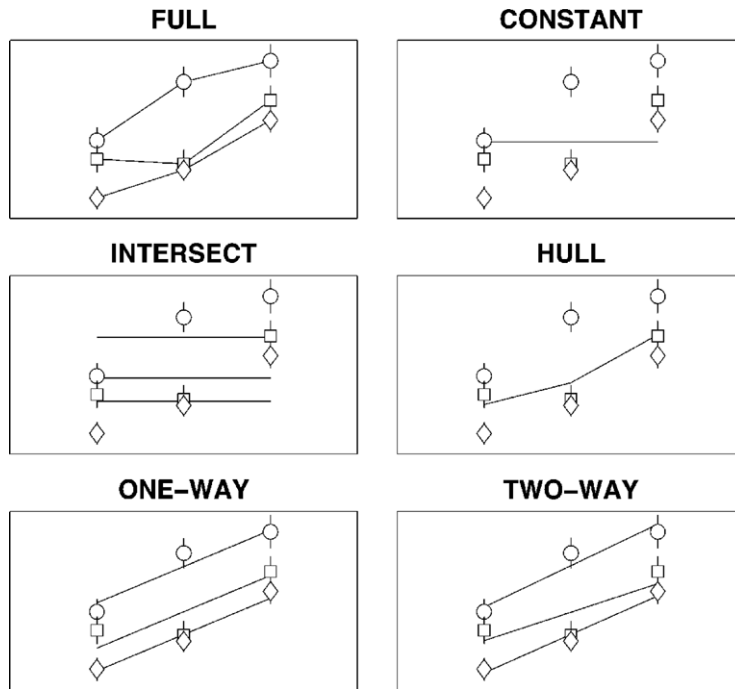


Figure 3. The best fits to deviation from optimality for the six competing models. For each model, the best fit is shown by the lines overlaid on the display of the data used in Figure 2.

but do so independently of each other. This conclusion is consistent with NHST results, which found two significant one-way effects but no significant interaction.

Response uncertainty. In addition to deviations from optimality in the case of individual observers, a group measure of response uncertainty was calculated for each array. The frequency with which each point in the array was connected with each other point was counted and an $n \times n$ matrix was constructed. Probabilities (p_i) were obtained by dividing the frequencies by the number of participants (20) and then applying Shannon's (1948) standard information theory formula, given in Equation 1, to calculate the total response uncertainty (H), with k being the total number of connections made by the participants:

$$H = \sum_{i=1}^k p_i \log_2 p_i. \tag{1}$$

Figure 4 shows the change in the response uncertainty measure, averaged across all eight arrays for each problem type, in terms of the convex hull and intersection factors. Figure 5 shows the best fits to response uncertainty under each of the six models. Table 2 gives the various Bayesian measures for each model against response uncertainty. In this case, the intersect model has the minimum BIC value and so is the basis for the Bayes factors. Our conclusion is that increasing the number of intersections does decrease response uncertainty, but there is no clear evidence that the number of points of the convex hull has an effect. This conclusion is also consistent with NHST methods, which

found a significant one-way effect for intersections but none for the convex hull and no significant interaction.

Discussion

The present study has a number of implications. These concern the effects of the number of convex hull points and the number of potential intersections on deviations from optimality and response uncertainty and the development of a locally based model of human performance on visually presented TSPs. We discuss these in turn.

Convex hull and nearest neighbor processes in solving TSPs. The analyses of our results support the interpretation that both the number of points on the convex hull and the number of potential intersections have independent ef-

Table 1
The Log-Likelihood Data Fit, Parametric Complexity, Bayesian Information Criterion (BIC) Value, and Bayes Factor for Each of the Six Models in Relation to the Deviation From Optimality Score

Model Name	Log-Likelihood Data Fit	Parametric Complexity	BIC Value	Bayes Factor
Full	10.55	9	-1.32	26.63
Constant	4.96	1	-7.72	6.55
Intersect	7.27	3	-7.95	5.43
Hull	7.65	3	-8.70	3.97
One-way	10.12	4	-11.46	1.00
Two-way	10.18	6	-7.19	7.86

Note—The Bayes factors are taken in relation to the most likely model, which is the one-way model.

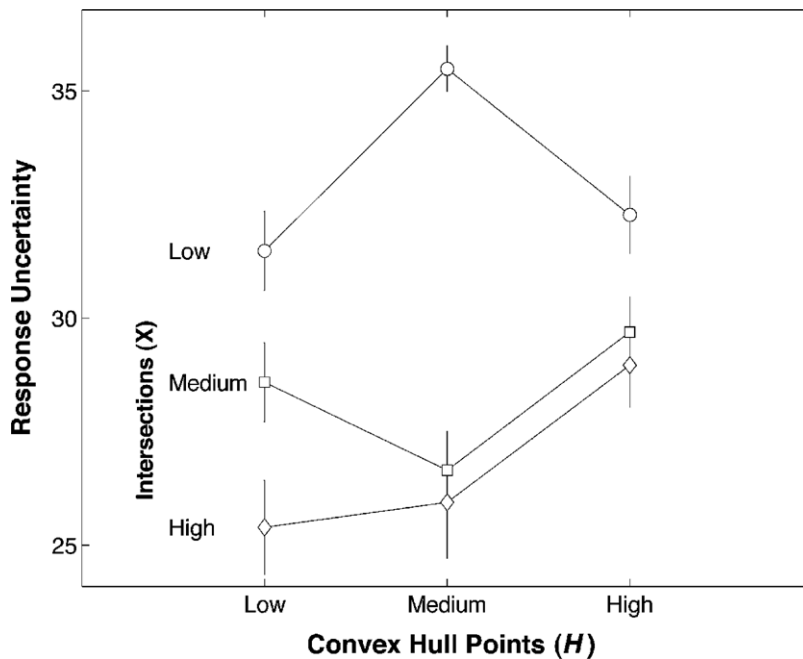


Figure 4. The mean response uncertainty, taken across all the participants' solutions to all eight problems, for each combination of points on the convex hull and number of intersections. The error bars show one standard error in each direction.

fects on deviations from optimality. However, the direct relation between the number of points on the convex hull and the extent of deviations from optimality contradicts the convex hull hypothesis of MacGregor and Ormerod (1996). In contrast, the inverse relation between number of potential intersections and deviation from optimality is consistent with a local-to-global processing procedure based on the automatic detection of NNs.

In view of previous interpretations in favor of the convex hull, the direct relation, shown in this experiment, between the number of points on the convex hull and the extent of deviations above optimality in solutions was unexpected. The configurations selected as problems were statistically representative of the population of randomly constructed 40-node arrays and it seems unlikely that an opposite outcome could be obtained with any other method of selecting configurations that would retain orthogonality in the values of the two independent variables.

One possible explanation of the apparent conflict with previous results is that potential intersections that are *internal* (i.e., those that do not involve any of the nodes on the convex hull) may be particularly important for solution quality. When 40-node arrays are generated by randomly selecting coordinates according to a uniform distribution, the proportion of potential internal intersections is around 23%. For the arrays we employed, this increases slightly when arrays are selected that have a higher total of potential intersections. However, when the number of nodes on the convex hull is manipulated (while holding the total number of potential intersections constant), the

number of internal intersections is halved (going from 29,231 to 14,173) as the number of convex hull nodes is increased from 7 to 12. As would be predicted by the hypothesis that observers employ some procedure based on NN links, percentage scores above the benchmark increase from 4.7%, through 5.1%, up to 5.7% as the number of convex hull nodes is increased from 7, through 9 or 10, up to 12, respectively.

The reason that internal intersections may be particularly important may be that a smaller number of internal intersections may indicate greater uncertainty in identifying NN links. In contrast, NN links that involve points on the convex hull are identified from among a much smaller set of possibilities. Thus, increasing the number of points on the convex hull constitutes a way of manipulating the relative distribution of internal and hull-related intersections, as opposed to varying the overall number of intersections, while keeping their relative distributions constant. In this way, the number of potential internal intersections and the number of potential intersections that involve points on the convex hull may have independent and unequal effects on performance, as was found.

The analyses also agree in indicating that number of potential intersections has an inverse effect on response uncertainty and that there is no evidence that number of points on the convex hull has an effect or that these two factors interact. This outcome fails to support MacGregor and Ormerod's (1996) interpretation of their finding of an inverse relation between number of points on the convex hull and response uncertainty. Again, however, it is consis-

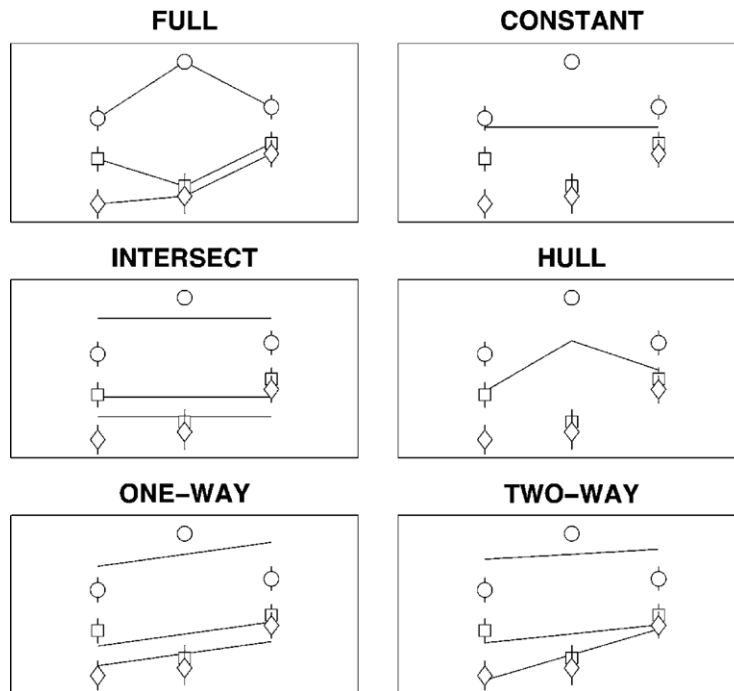


Figure 5. The best fits to response uncertainty for the six competing models. For each model, the best fit is shown by the lines overlayed on the display of the data used in Figure 4.

tent with the hypothesis that observers use NN links to arrive at solutions and effectively exclude paths that intersect.

The finding that an increase in the number of convex hull points (with number of potential intersections held constant) did not produce an increase in response uncertainty (paralleling the decrease in solution optimality) may be because two opposing influences are at work. As was just argued, the concomitant decrease in the number of potential internal intersections may be associated with an increase in response uncertainty (because NNs are harder to identify). On the other hand, because an increase in intersections involving the convex hull may be associated with a decrease in response uncertainty and because near-optimal solutions are in any case constrained to follow the convex hull, it may be that, when the number of points on the convex hull is manipulated, while holding the total number of potential intersections constant, these two opposing influences effectively cancel each other. The potential for differences in the distribution of internal and hull-related intersections to produce different effects on the two response measures is clearly an important question for further research.

Taken in conjunction with the interpretative weaknesses in the evidence cited earlier, these findings suggest that there is no evidence that can be adduced as unqualified support for the global-to-local convex hull hypothesis and there appears to be some clear evidence that runs counter to it. In contrast, the view that observers arrive at solutions on the basis of detecting NN links that, in effect, exclude intersections is well supported by the main results of the

present experiment, although the relative effects of internal and hull-related intersections need clarification.

A hierarchical nearest neighbor process for solving TSPs. We have argued that the inverse relation found between the number of potential intersections and both deviation from optimality and response uncertainty is consistent with the hypothesis that observers exploit spontaneous perceptual organization based on the identification of links between NNs. We have further suggested that such links could be identified automatically by a visual system that processed relational information in parallel.

One of the simplest procedures for achieving reasonably good solutions to the TSP is the standard NN algorithm (SNN; Golden et al., 1980). Despite its simplicity, the SNN algorithm has some attractions as a model of human TSP performance. First, unlike the convex hull approach, it can

Table 2
The Log-Likelihood Data Fit, Parametric Complexity, Bayes Information Criterion (BIC) Value, and Bayes Factor for Each of the Six Models for Response Uncertainty

Model Name	Log-Likelihood Data Fit	Parametric Complexity	BIC Value	Bayes Factor
Full	-1.70	9	23.17	19.17
Constant	-6.97	1	16.14	5.33
Intersect	-2.82	3	12.23	1.00
Hull	-6.02	3	18.62	9.44
One-way	-2.64	4	14.07	2.51
Two-way	-2.55	6	18.29	8.27

Note—The Bayes factors are taken in relation to the most likely model, which is the intersect model.

apply to both open tours, with different starting and finishing nodes, and closed tours, with the same start and finish. Second, some form of NN procedure is likely to be useful in arriving at good solutions to other optimization problems, such as the MST and ST tasks mentioned earlier. Third, the production of a human TSP solution takes place in a serial link-by-link order, similar to the operation of the SNN algorithm.

Nevertheless, few researchers would suggest that observers follow a simple SNN procedure. MacGregor and Ormerod (1996) found that mean solutions, averaged across observers and across problems, in both their experiments were consistently superior to the mean solutions produced by the SNN algorithm. Similarly, Graham et al. (2000) found that, when the problems were larger (20 or 30 nodes), the SNN algorithm produced substantially longer tours. If these results had been the reverse, it would be possible to explain them by supposing that human performance was also limited by some other factor. However, the fact that human performance is superior to that of the SNN procedure means that the standard algorithm is seriously deficient.

An alternative but closely related possibility is to suppose that observers first establish clusters of two or more nodes, based on NN distances, and then sequentially link these up into a tour, possibly using some version of a *nearest cluster* algorithm. The first process of cluster organization seems likely to be an automatic parallel process. However, the second stage of linking the clusters is more naturally conceived as a sequential cognitively penetrable one, in which performance is influenced by the choice of starting point. A hierarchical nearest neighbor (HNN) hypothesis of this kind would have a number of implications.

First, as is shown by our results, the number of potential intersections should determine human TSP performance. None of the NN clusters can contain intersections. It also seems possible that a sequential process of linking nearest clusters might be guided by a conscious avoidance of intersections. A hierarchical HNN procedure of this kind would be consistent with the almost universal avoidance of intersections in human TSP solutions. Such a procedure would restrict the number of alternatives that have to be considered. Therefore, an HNN hypothesis would predict that a greater number of potential intersections would facilitate performance, both in terms of producing shorter solutions and in terms of reducing uncertainty concerning which node to link next.

Second, performance on TSPs should be affected by any manipulation that changes the distribution of NNs, such as the degree of regularity in an array (MacGregor et al., 1999). In particular, a major factor is dot density (or number of nodes). It is to be expected that the number of nodes should influence the time taken to solve TSPs. However, whereas most algorithms require an accelerating increase in time as the number of nodes becomes greater, observers appear to take a time that is a linear function of the number of nodes (Graham et al., 2000; Vickers, Bovet, et al., 2003). As has been argued by Graham et al., this suggests that observers employ a simple procedure that does not necessitate searching through a substantial por-

tion of the space of potential TSP solutions. Because the number of NN clusters is a linear function of the number of nodes in an array (Vickers, Preiss, & Hughes, 2003), and because each decision about linking NN clusters is similar in form, the HNN would seem to be one example of such a procedure.

At this stage, we are unsure of whether the avoidance of intersections in TSP pathways is solely a necessary consequence of following a solution procedure based on the detection of NN clusters or whether it is also in part a probable consequence of deliberately avoiding intersections in the process of connecting the NN clusters into a solution pathway. However, it does seem to be the case that, by following a locally focused process, observers are able to exploit spontaneous organizing tendencies of the visual system to arrive at TSP solutions that approximate closely to a globally defined optimality. On a more general level, the hypothesis that observers can exploit locally focused processes that use relational information to achieve a level of performance that is not far short of some globally specified optimum is consistent not only with current theorizing that is based on some form of optimization as an explanatory principle (e.g., Chater & Vitányi, 2003), but also with recent work that emphasizes the simplifying heuristics that human beings use to make decisions that are near-optimal (e.g., Gigerenzer & Todd, 1999).

REFERENCES

- ATTNEAVE, F. (1974). Apparent movement and the what-where connection. *Psychologia*, **17**, 108-120.
- CAELLI, T. M. (1981). Some psychophysical determinants of discrete Moiré patterns. *Biological Cybernetics*, **39**, 97-103.
- CARLIN, B. P., & LOUIS, T. A. (2000). *Bayes and empirical Bayes methods for data analysis* (2nd ed.). New York: Chapman & Hall.
- CHATER, N., & VITÁNYI, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, **7**, 19-22.
- CHI, M. T. H., GLASER, R., & FARR, M. J. (1988). *The nature of expertise*. Hillsdale, NJ: Erlbaum.
- COHEN, J. (1994). The earth is round ($p < .05$). *American Psychologist*, **49**, 997-1003.
- DAWSON, M. R. W. (1991). The how and why of what went where in apparent motion: Modeling solutions to the motion correspondence problem. *Psychological Review*, **98**, 569-603.
- DUNCAN, C. P. (1959). Recent research on human problem solving. *Psychological Bulletin*, **56**, 397-429.
- EDWARDS, W., LINDMAN, H., & SAVAGE, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, **70**, 193-242.
- ERICSSON, K. A., & LEHMANN, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, **47**, 273-305.
- FLOOD, M. M. (1956). The travelling salesman problem. *Operations Research*, **4**, 61-75.
- GELMAN, A., CARLIN, J. B., STERN, H. S., & RUBIN, D. B. (1995). *Bayesian data analysis*. London: Chapman & Hall.
- GIGERENZER, G., & TODD, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- GLASS, L. (1969). Moiré effect from random dots. *Nature*, **243**, 578-580.
- GOLDEN, B., BODIN, L., DOYLE, T., & STEWART, W. (1980). Approximate travelling salesman algorithms. *Operations Research*, **28**, 694-711.
- GRAHAM, S. M., JOSHI, A., & PIZLO, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & Cognition*, **28**, 1191-1204.
- HOWSON, C., & URBACH, P. (1993). *Scientific reasoning: The Bayesian approach*. La Salle, IL: Open Court.

- HUNTER, J. E. (1997). Needed: A ban on the significance test. *Psychological Science*, **8**, 3-7.
- HWANG, F. K., RICHARDS, D. S., & WINTER, P. (1992). *The Steiner tree problem*. Amsterdam: Elsevier.
- JOHNSONBAUGH, R. (2001). *Discrete mathematics*. Upper Saddle River, NJ: Prentice-Hall.
- KASS, R. E., & RAFTERY, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, **90**, 773-795.
- KOLERS, P. A. (1972). *Aspects of motion perception*. London: Pergamon.
- LAWLER, E. L., LENSTRA, J. K., RINOOY KAN, A. H. G., & SCHMOYS, D. B. (1985). *The traveling salesman problem: A guided tour of combinatorial optimization*. Chichester, U.K.: Wiley.
- LEE, M. D., & VICKERS, D. (2000). The importance of the convex hull for human performance on the traveling salesman problem: A comment on MacGregor and Ormerod (1996). *Perception & Psychophysics*, **62**, 226-228.
- LEE, R. K. L. (1985). *A heuristic approach to the traveling salesman problem* [Unpublished Management Rep.]. Victoria, BC: University of Victoria, School of Public Administration.
- LEONARD, T., & HSU, J. S. J. (1999). *Bayesian methods: An analysis for statisticians and interdisciplinary researchers*. New York: Cambridge University Press.
- LINDLEY, D. V. (1972). *Bayesian statistics: A review*. Philadelphia: Society for Industrial and Applied Mathematics.
- LOVETT, M. C. (2002). Problem solving. In H. Pashler & D. Medin (Eds.), *Stevens' Handbook of experimental psychology: Vol. 2. Memory and cognitive processes* (pp. 317-362). New York: Wiley.
- MACGREGOR, J. N., & ORMEROD, T. C. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics*, **58**, 527-539.
- MACGREGOR, J. N., & ORMEROD, T. C. (2000). Evaluating the importance of the convex hull in solving the Euclidean version of the traveling salesperson problem: Reply to Lee and Vickers (2000). *Perception & Psychophysics*, **62**, 1501-1503.
- MACGREGOR, J. N., ORMEROD, T. C., & CHRONICLE, E. P. (1999). Spatial and contextual factors in the traveling salesperson problem. *Perception*, **28**, 1417-1428.
- MACGREGOR, J. N., ORMEROD, T. C., & CHRONICLE, E. P. (2000). A model of human performance on the traveling salesperson problem. *Memory & Cognition*, **28**, 1183-1190.
- MOTTER, B. C., & BELKY, E. J. (1998). The zone of focal attention during active visual search. *Vision Research*, **38**, 1007-1022.
- NAVON, D. (1976). Irrelevance of figural identity for resolving ambiguities in apparent motion. *Journal of Experimental Psychology: Human Perception & Performance*, **2**, 130-138.
- NEWELL, A., & SIMON, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- O'ROURKE, J. (1994). *Computational geometry in C*. New York: Cambridge University Press.
- PALMER, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
- POMERANTZ, J. R. (1981). Perceptual organization and information processing. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 141-180). Hillsdale, NJ: Erlbaum.
- PRESS, W. H. (1992). *Numerical recipes in FORTRAN: The art of scientific computing* (2nd ed.). Cambridge: Cambridge University Press.
- QUINTAS, L. V., & SUPNICK, F. (1965). On some properties of shortest Hamiltonian circuits. *American Mathematical Monthly*, **72**, 977-980.
- RAINVILLE, S. J. M., & KINGDOM, F. A. A. (2002). Scale invariance is driven by stimulus density. *Vision Research*, **42**, 351-367.
- REINELT, G. (1994). The traveling salesman: Computational solutions for TSP applications. In *Lecture notes in computer science* (Vol. 840). Berlin: Springer-Verlag.
- RESNICK, R., & GLASER, L. B. (1976). Problem solving and intelligence. In L. B. Resnick (Ed.), *The nature of intelligence* (pp. 205-230). Hillsdale, NJ: Erlbaum.
- ROSS, J., BADCOCK, D. R., & HAYES, A. (2000). Coherent global motion in the absence of coherent velocity signals. *Current Biology*, **10**, 679-682.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, **6**, 461-464.
- SHANNON, C. E. (1948). The mathematical theory of communication. *Bell Systems Technical Journal*, **27**, 379-243, 623-656.
- SIVIA, D. S. (1996). *Data analysis: A Bayesian tutorial*. Oxford: Oxford University Press, Clarendon Press.
- STERNBERG, R. J. (1982). Reasoning, problem solving, and intelligence. In R. J. Sternberg (Ed.), *Handbook of intelligence* (pp. 225-307). New York: Cambridge University Press.
- ULLMAN, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.
- VAN ROOIJ, I., STEGE, U., & SCHACTMAN, A. (2003). Convex hull and tour crossings in the Euclidean traveling salesperson problem: Implications for human performance studies. *Memory & Cognition*, **31**, 215-220.
- VICKERS, D., BOVET, P., LEE, M. D., & HUGHES, P. (2003). The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean travelling salesperson problems. *Perception*, **32**, 871-886.
- VICKERS, D., BUTAVICIUS, M. A., LEE, M. D., & MEDVEDEV, A. (2001). Human performance on visually presented traveling salesman problems. *Psychological Research*, **65**, 34-45.
- VICKERS, D., DRY, M., LEE, M. D., & HUGHES, P. (2003). *The role of nearest neighbours in the perception of Glass patterns*. Manuscript submitted for publication.
- VICKERS, D., MAYO, T., HEITMANN, M., LEE, M. D., & HUGHES, P. (in press). Intelligence and individual differences in performance on three types of visually presented optimisation problems. *Personality & Individual Differences*.
- VICKERS, D., & PREISS, A. K. (2003). *Impressions of motion from Glass patterns*. Manuscript submitted for publication.
- VICKERS, D., PREISS, A. K., & HUGHES, P. (2003). *The role of nearest neighbours in the perception of structure and motion in dot patterns*. Manuscript submitted for publication.
- WILF, H. S. (1986). *Algorithms and complexity*. Englewood Cliffs, NJ: Prentice-Hall.

NOTE

1 The convex hull is a boundary, so that no line joining any two nodes in the array can fall outside it.

(Manuscript received June 7, 2002;
revision accepted for publication May 13, 2003.)