
Spatial and contextual factors in human performance on the travelling salesperson problem

James N MacGregor

School of Public Administration, University of Victoria, Victoria, British Columbia V8W 2Y2, Canada

Thomas C Ormerod, Edward P Chronicle[¶]

Department of Psychology, Lancaster University, Lancaster LA1 4YF, UK;

e-mail: E.Chronicle@lancaster.ac.uk

Received 2 November 1998, in revised form 28 September 1999

Abstract. The travelling salesperson problem (TSP) provides a realistic and practical example of a visuo-spatial problem-solving task. In previous research, we have found that the quality of solutions produced by human participants for small TSPs compares well with solutions from a range of computer algorithms. We have proposed that the ability of participants to find solutions reflects the natural properties of human perception, solutions being found through global perceptual processing of the problem array to extract a best figure from the TSP points. In this paper, we extend the study of human performance on the task in order to understand further how human abilities are utilised in solving real-world TSPs. The results of experiment 1 show that high levels of solution quality are maintained in solving larger TSPs than had been investigated previously with human participants, and that the presence of an implied real-world context in the problems has no effect upon performance. Experiment 2 demonstrated that the presence of regularity in the point layout of a TSP can facilitate performance. This was confirmed in experiment 3, where effects of the internality of point clusters were also found. All three experiments were consistent with a global, perceptually based approach to the problem by participants. We suggest that the role of perceptual processing in spatial problem-solving is an important area for further research in both theoretical and applied domains.

1 Introduction

In the Euclidean version of the travelling salesperson problem (TSP), the shortest tour is found around a set of locations on a two-dimensional plane in which each location is visited only once, and the starting location is returned to. The TSP is ubiquitous in commercial and industrial contexts where near-optimal (ie shortest distance) solutions are desirable to minimise the use of resources. In addition to its practical importance, the general form of the TSP has been of considerable theoretical importance in computer science and operations research.

Considerable efforts have been made in the field of operations research to identify efficient heuristic algorithms that provide TSP solutions of acceptable quality whilst minimising computational cost. However, even the most efficient heuristic algorithms still require in the order of P^3 computations to compute solutions to within 2%–3% of optimality, where P is the number of points (Golden et al 1980). Operations researchers have, from time to time, attempted to examine human performance on TSPs. One significant attempt was that of Krolak et al (1971). In Krolak et al's procedure, an algorithm generated local clusters of points (eg reducing a random array of 100 points to 14 clusters). These were then connected by human participants to produce a tour, which was then compared against computer-generated tours and subsequently revised in an iterative cycle between human and computer. Krolak et al reported a number of comparisons of human–computer and computer-only solution performance, and demonstrated that human–computer generated tours are at least as short, and often shorter, than those of computer-based algorithms.

[¶] Author to whom all correspondence and requests for reprints should be addressed.

Krolak et al did not attempt to characterise the nature of human-only solutions to TSPs, nor the cognitive processes by which they might be generated. MacGregor and Ormerod (1996) argued that psychological research into human performance was not only warranted, but overdue, for four reasons: first, the potential ability of human cognitive processes to solve such problems in reasonable times is of fundamental interest. Second, the simplicity and flexibility of the TSP as a problem offers an attractive tool to investigate aspects of spatial cognition more generally. Third, knowledge of the cognitive processes utilised during performance on the TSP might permit explanation of performance on other spatial problems which is not well accounted for by existing cognitive architectures. For example, TSP solutions have recently been found to be important in the high-level planning of eye-movements during visual-search tasks (Pomplun and Ritter 1999). Fourth, human-only solutions to real-world examples of TSPs might outperform heuristic algorithms, which would offer potential benefits in applied contexts.

MacGregor and Ormerod (1996) demonstrated that, with 10-point and 20-point problems, human solvers were exceptionally good at finding the optimum tours around a set of problems that varied in complexity, outperforming all of the heuristic algorithms against which their performance was measured. Average solution time was less than 40 s, including drawing the route through the nodes. In a subsequent study employing the same TSPs as MacGregor and Ormerod (1996), Ormerod and Chronicle (1999) provided evidence strongly suggesting that human TSP performance is based upon global perceptual processing of the problem array. Ormerod and Chronicle used a solution-identification task in which participants had to judge from a rapid (2 s) stimulus exposure whether a given solution to a 10-point TSP was optimal or sub-optimal. They found that identification was systematically dependent on the optimality of the given solution, even under forced-pace response conditions, and hypothesised that participants matched the given solution against a figure extracted by global perceptual processing of the problem array. This hypothesis was supported by data from a further experiment, in which the presentation of short (100 ms) 'primes' of the problem array alone prior to presentation of the problem plus solution stimulus significantly facilitated identification performance relative to no prime or to longer primes (above 500 ms). Ormerod and Chronicle interpreted their results in terms of Sanocki's (1993) 'global-to-local contingency' theory of the perceptual processing of objects. In this theory, a 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. In the context of the TSP, this account implies that a global hypothesis about a best figure might be available early in processing, and could potentially inform a decision about optimal route. In other words, participants may solve TSPs by identifying a best figure [in the sense used by the Gestalt school, cf Attneave (1982)] that fits the dot array and then utilising this figure to generate a tour. It may be that the provision of point clusters, as in Krolak et al's approach, aids in the extraction of a best figure. Alternatively, it may be that a best figure can be extracted from a global scan of the array, without requiring local processing to extract clusters of points. As MacGregor and Ormerod (1996) 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" (page 537). It is therefore conceivable that, at least with the TSP exemplars used by MacGregor and Ormerod (1996) and Ormerod and Chronicle (1999), the interaction between the figural (Gestalt) properties of the problem and the natural inclinations of the visual system fortuitously allows efficient problem-solving.

Before one can fully capitalise upon human capabilities with TSPs, as suggested above, it is necessary to explore performance on a wider range of problems. Three specific issues need to be tackled, and these are investigated in the experiments reported here. First, is a high standard of human performance limited to only small (10-point to 20-point) problems, such as those investigated by MacGregor and Ormerod? We have anecdotal evidence that the quality of human solutions is maintained with 40-point and 60-point problems. The use of 48-point problems in experiment 1 allowed a proper empirical investigation of this issue. Second, is a high standard of human performance restricted to TSPs with specific spatial properties? The layout of problems explored by MacGregor and Ormerod corresponded to an approximate convex hull with varying numbers of internal points. It may be that high levels of human performance are restricted to types of display such as these. In experiments 2 and 3 we explored human performance with problems in which the display characteristics were manipulated systematically. Third, does an implied real-world context alter human performance relative to context-free TSPs as used in previous research? Buckmaster (1992) reports data from a pilot study suggesting that the presence of real-world thematic context (namely, the superimposition of a familiar geographical map over the TSP problem) actively impairs human solution-finding performance. If this result were replicated, it would have important and damaging implications for the utilisation of human capacities in solving real-world TSPs, that by their very nature embody realistic context. This issue was explored as an additional manipulation in experiment 1.

2 Experiment 1

Experiment 1 had two main purposes: to measure human performance on a TSP larger than previously used, and to investigate the effects on performance, if any, of providing different contexts and implied thematic content to the dot arrays. There are many conceivable ways in which context could be manipulated, and the two chosen here—a circuit-board diagram and a map of the USA—were chosen because (a) they provide seemingly very different contexts and (b) they offer plausible and realistic practical examples of TSPs. In all, three identical 48-point TSPs were used in the experiment. The first was presented as a circuit-board diagram, and the dots were surrounded by a rectangular frame to represent the ‘board’. The second was presented as a map of USA cities, and the dot array was framed with an outline of the United States. The third was presented with no context and no frame around the display.

2.1 Method

2.1.1 *Participants.* The participants were 141 undergraduate student volunteers, of whom 103 provided usable responses. Since the experiment was conducted in a group setting, the gender and age distributions are unknown.

2.1.2 *Materials.* The dot array was created by randomly selecting 48 cities from the 100 cities of Problem 24 reported by Krolak et al (1971). The Cartesian coordinates for the points will be available on the Perception World Wide Web site (<http://www.perception.com/10ggmacgregor/html/>). The points were scaled so that when printed on regular A4 paper they were contained within a rectangle 140 mm wide and 70 mm high. The points were printed as filled circles, approximately 3 mm in diameter. For the circuit-board condition, the dot array was surrounded by a 160 mm × 90 mm rectangular frame. For the USA-cities condition, the array was framed by a sketch of boundaries of the United States, scaled so that a number of points fell at seemingly appropriate locations on the ‘coasts’. The dot array was identical for each participant, in order to avoid the possibility of figural variation confounding the effect of context, and is shown in figure 1.




Figure 1. The 48-point problem used in experiment 1. In the USA-cities condition the points were surrounded by an outline of the USA, and in the circuit-board condition they were surrounded by a rectangular frame (note that the point size was smaller in the stimuli, removing overlaps between points).

2.1.3 Procedure. Participants completed the task following another, unrelated, paper-and-pencil task. Booklets containing the tasks were arranged in random order of experimental condition and passed out to participants. The instructions for the no context condition were to "... draw the shortest continuous path you can through all of the dots, visiting each dot once and once only and returning to the point where you began ...". Instructions were altered for the other conditions by describing the array as a map of the United States or the layout of a circuit board, and using "city" or "transistor" instead of "dot" and "point". Participants were given 3 min to complete the task.

2.2 Results and discussion

The number of responses for the USA-cities, circuit-board, and no-context conditions were 47, 46, and 48, respectively; of these the numbers of unscorable responses were 11, 13, and 14, respectively. Unscorable responses occurred when either the path was left incomplete or when a point was revisited. They were eliminated from further analysis. Path lengths were measured in millimetres for each of the remaining solutions. The means (and standard deviations) for the three conditions were: 607 (30), 603 (31), and 608 (31) for the USA-cities, circuit-board, and no-context conditions, respectively. The conditions did not differ significantly, $F_{2,100} = 0.30$.

Heuristic procedures were applied, with the TRAVEL program (Boyd et al 1987), in an attempt to find the shortest path for the problem. These procedures can mathematically establish the lower bound of the optimal solution. The shortest path cannot be lower than the lower bound: it is either equal to or higher than this. If a path is found whose length equals the lower bound, then it must be optimal. If the path length found is higher than the lower bound, then it may or it may not be optimal. The best path found had a length of 543, which is 1% above the lower bound of an optimal solution. The best human solution was 553, 2% above the best path found. On average, the participants' solutions were 10% above the best-known solution. The worst was 26% above the best-known solution.

With respect to the effect of context, from the present results it appears to be nil. The results for the two conditions in which contexts were provided are almost identical to those of the control condition. Admittedly, there are numerous other contexts which have not been tested here, but the wide dissimilarity between the contexts suggested by a map of US cities and a circuit board indicates that the result is likely to be generalisable. The results also indicate that placing frames around TSPs in close proximity to the perimeter of the dot array also has no effect on performance.

Regarding the quality of human solutions, the results appear to support and extend our previous findings with smaller problems (MacGregor and Ormerod 1996) where for 10-node and 20-node problems participants' path lengths were on average 3.1% and

6.3%, respectively, above the optimal (or best-known) solutions. For the 10-node problems, the best participant solutions were optimal, while for the 20-node problems, the best participant solutions were on average 1.5% above optimal. The average participant solution to the present 48-node problem was 10% above the best-known solution, while the best participant solution was 2% above. The decrease in performance with increasing number of points appears from the present results to be, at worst, linear. This is an important outcome, since the complexity of a TSP, in terms of the number of possible solutions, increases as a *factorial* function of the number of points.

The operating characteristics of a number of heuristic procedures for solving TSPs are known, and provide some benchmarks against which to compare the participants' performances. For example, the worst-case behaviour for the nearest-neighbour procedure, measured as the ratio of the heuristics solution to the optimal solution, is $\frac{1}{2}\log_2 P + \frac{1}{2}$. For a 48-node problem, this gives a worst-case solution of 3.3 times the optimal. The equivalent human result here was 1.3 times the optimal. Golden et al (1980) report the worst-case behaviour for six more 'tour construction' heuristics, and they range from 1.5 to 7.9 for a 48-node problem. The observed worst-case of human solvers falls comfortably below this range.

3 Experiment 2

The results of experiment 1 suggest that the standard of human performance with larger TSPs, although not as high in absolute terms as that found by MacGregor and Ormerod (1996) with small problems, is nonetheless impressive. A linear decrease in performance in the face of a factorial increase in problem complexity suggests that, even with the random-point arrays of experiment 1, there are facets of human performance that can contribute usefully to the solution of large-scale problems. It is interesting to consider the effects on human performance of introducing nonrandom-point positions into a TSP. For example, in some real-world applications such as in the planning of circuit-board wiring, a set of points might be expected to be organised around a regular grid-like structure. The effects on human performance of introducing regularity into TSPs have not been investigated previously. It is possible that regularity may aid the production of solutions, for example by limiting the space of possible solutions to those suggested by regular figural components. On the other hand, regularity may impair performance, for example in situations where more than one good Gestalt figure is perceivable in a regular array, or by encouraging the drawing of solutions to conform to the regular structure which contain lines that cross over or that have acute angles.

Therefore in experiment 2 we compared responses to a randomly generated TSP with those to one having a regular structure. The random condition was obtained by combining the responses obtained in the previous experiment. These were compared with the responses from an additional sample of participants to a 48-node TSP, constructed from the vertices of a 6×8 matrix.

3.1 Method

3.1.1 *Participants*. The participants were 189 undergraduate volunteers (141 from experiment 1), of whom 137 provided usable responses. As before, gender and age distributions were unknown.

3.1.2 *Materials*. The stimulus for the structured TSP consisted of 48 dots arranged in a 6×8 matrix approximately 100 mm high and 140 mm wide printed on A4 paper. Dots were filled circles 3 mm in diameter, as before. The structured TSP is shown in figure 2.




Figure 2. The ‘regular’ 48-point problem used in experiment 2, with a possible optimal tour shown on the right.

3.1.3 *Procedure.* The instructions and procedure for participants in the regular condition were identical to those of the control condition in experiment 1.

3.2 Results and discussion

There were 48 responses to the regular condition, 14 of which were unscorable and eliminated from further analysis. Path lengths were measured in millimetres. The optimal solution for the regular TSP is a path of length 960 mm (the lateral distance between adjacent points times the number of points). The participants’ mean path length was 976, 1.7% above the optimal, with a standard deviation of 67. One participant produced a path of length 1348, more than 5 standard deviations above the mean. If this result is excluded, the average performance was 0.005% above the optimal. 29 of the 34 subjects produced optimal paths, compared to 0 out of 103 responses to the random-dot array. Clearly, the difference is significant, $\chi_1^2 = 111.40$.

The results suggest that, for the arrays shown in figure 2, regularity in a TSP can facilitate the production of optimal solutions. It is interesting to note that the 34 participants with a scoreable solution to the regular problem produced only 11 different solutions. Of these, 27 participants produced a variant of the solution shown in figure 2 (10 identical, the remainder producing a solution rotated through 90° or 180°). This can be compared with the control condition of experiment 1, in which no pair of participants produced an identical solution. It is interesting to note that performance was not impaired by the regularity of the array, and that participants were not obviously misled by the regularity of the array into producing unscorable solution attempts at a greater rate than with the random arrays of experiment 1.

4 Experiment 3

Previous research on TSPs of 10 and 20 nodes has suggested in preliminary fashion that participants’ solutions are affected not so much by the total number of nodes as by how those nodes are distributed. Specifically, the greater the number of interior points falling within the perimeter of the dot array, the more complex the problem appears to be for participants (MacGregor and Ormerod 1996). Experiment 2, above, suggests the additional possibility that arrays that have a regular distribution may be more simple than random arrays.

In our previous research we found that, to a very high degree, participants produce paths that connect adjacent points on the perimeter in order (though interior points may intervene). In other words, there is little uncertainty regarding the order of connection of boundary points. This means that a TSP consisting only of boundary points would be trivial for human participants. It also explains why it is the number of interior

points, rather than the total number of points, that determines the difficulty of a problem. However, what of points that lie close to the boundary? Are they treated in the same manner as boundary points? In other words, there may be an element of degree to the boundary/interior distinction. This suggests that the proximity of nodes to the perimeter may be a factor in problem difficulty. Experiment 3 was designed to explore further the factors of internality and regularity.

4.1 Method

4.1.1 *Participants.* The participants were 175 undergraduate volunteers enrolled in a psychology course. Again the experiment was conducted in a group setting, so that age and gender distributions are unknown.

4.1.2 *Materials.* Six 12-node TSPs were used, all constructed from the same basic design consisting of 8 perimeter nodes located at the vertices of a regular octagon and 4 interior nodes, located at points of intersection of the two vertical and two horizontal arcs that would connect opposite pairs of boundary points. Two other regular TSPs were created from this structure, one by 'shrinking' the central 'square' of nodes, the other by transposing this smaller 'square' to lie closer to the perimeter. This provided three regular TSPs whose interior points varied in terms of proximity to the boundary. Next, irregular versions of each of these were created by applying small random distortions to the coordinates (see figure 3). Each problem was drawn within a




Figure 3. The six problems used in experiment 3. Regular figures are shown in the left column, irregular figures in the right. The top row shows patterns with intermediate internal points, the middle row central, and the bottom peripheral.

120 mm × 120 mm area, and the approximate diameter of each point was 3 mm. The six problems were printed on separate A4 pages and combined in booklets with a covering instruction page. The order of presentation of problems was counterbalanced.

4.1.3 Procedure. The booklets containing the tasks were arranged in random order and passed out to participants. The instructions were similar to those used in experiment 2, except that participants were instructed to attempt all six problems. Thus, the experiment employed a 2 × 3 within-participants design, the factors being location of interior points (peripheral, intermediate, and central) and pattern (regular and distorted).

4.2 Results and discussion

The optimal solution for each of the six TSPs was found and performance was measured in terms of the number of participants finding the optimal tour for each. These numbers are shown in table 1. The results were analysed by repeated-measures analysis of variance, which indicated significant effects of the interaction of location with pattern, $F_{2,348} = 12.55$, $MSE = 0.14$; the pattern factor, $F_{1,174} = 29.18$, $MSE = 0.15$; and the location factor, $F_{2,348} = 162.60$, $MSE = 0.17$. The three effects accounted for 7%, 14%, and 48% of their respective variance components. Overall, the mean path length of participants' solutions was 6% above the optimal, the median was 3% above. The worst solution was 68% above the optimal.

Table 1. The number (percentage) of participants finding the optimal solution for each of the six experimental conditions in experiment 3.

Pattern	Location		
	peripheral	intermediate	central
Regular	117 (67)	58 (33)	37 (21)
Distorted	116 (66)	9 (5)	19 (11)

The results illustrated that, although participants' solutions to the TSPs of experiment 3 were generally very good, their quality varied systematically with differences in the layout of the displays. The significant interaction between the location of interior nodes and their pattern seems to have arisen, in large measure, because of the fact that, for the two peripheral locations the problems were extremely easy, and no differences arose due to variations in pattern for these locations. However, at the intermediate and central locations, pattern differences appear to have had an effect, perhaps more so in the former than the latter. The effect was in the direction of the regular displays being simpler than the distorted. While the interaction was significant, it did not account for a very high percentage of variance. The main effect of location was by far the strongest effect, and generally appears to indicate that, when interior points are located near the boundary, TSPs are simpler for human participants. This reinforces and extends the previous suggestions about the role of internal nodes (MacGregor and Ormerod 1996). The main effect of pattern, while smaller, was still substantial, and, while the results indicate that the effects of regularity may not be straightforward, more-regular displays seem to produce better performances. This is consistent with the findings of experiment 2. A potential caveat, stemming from the within-subjects experimental design, was that the order of problems encountered might allow the development and use of strategies for tackling irregular patterns. For example, participants may have learned from the regular patterns to make a single indent from the convex hull to visit all interior points, then used that strategy for the irregular patterns. Although the risk of a confound is low given the counterbalancing of presentation order, a subset of the data was reanalysed to check for systematic differences in performance

depending on whether the particular problem was encountered first or last in the series. For the irregular/peripheral pattern, the number of participants finding the optimal solution was not significantly higher when the problem was encountered last as compared to when it was encountered first ($\chi^2 = 1.24$, $N = 75$, $p = 0.26$). Furthermore, MacGregor and Ormerod (1996) reported no order effects in their original study of human performance with TSPs. It therefore seems unlikely that any systematic distortion of the data would have arisen in this fashion.

One limitation of the experiment that will be apparent from figure 3 is that the degree of clustering of the internal nodes varied in a way that was partially confounded with their location. For the peripheral and central conditions the dots were more tightly clustered than for the intermediate location. The large effect of location in the former case indicates that, whatever effects clustering may have, they do not eliminate the effects of location. However, the differences between the intermediate and other locations could be due in part to effects of clustering as well as of location and pattern. In particular, the apparently better performance of the distorted pattern in the central over the intermediate location may reflect the influence of clustering. This is clearly a factor that should be incorporated into future experiments of this kind.

A second limitation arises because, when nodes were randomly shifted to produce the distorted conditions, some of the points on the boundary inadvertently became, to a small degree, interior points. Thus, strictly speaking, the distorted patterns contained more interior points. However, the results for the two peripheral conditions were virtually identical, indicating that this minor difference had no effect. This is what would be expected, since the points were so close to the boundary as to act, effectively, like boundary points.

5 General discussion

The experiments reported in this paper were conducted to investigate three issues concerning human performance with TSPs: the nature of human performance with larger TSPs than have previously been reported, the role played by spatial variations such as figural regularity in TSP arrays, and the effect of realistic context on performance. The results of these experiments underline the role of perceptual processes in solving problems of this nature, as we discuss below. Furthermore, we argue that there is much to be gained in optimising the solution of large-scale realistic TSPs by capitalising upon the unaided performance of human solvers.

With respect to human performance with larger TSPs, in experiments 1 and 2 we investigated performance on TSPs with 48 points, compared with MacGregor and Ormerod's (1996) study of human performance with 10-point and 20-point problems. The results of experiment 1 indicate that human participants are capable of producing high-quality solutions to larger problems. Whilst performance was not quite as good with 48-point problems as that observed by MacGregor and Ormerod with smaller problems, the results are nonetheless very encouraging. Decline in performance with increasing problem size appears to be roughly linear, in the face of a factorial increase in problem complexity (in terms of the number of possible solutions). One of the most interesting aspects of the data from experiment 1 from an applied perspective is that the worst human participant's solution was still better than the worst case behaviour of all the computational algorithms tested by Golden et al (1980). This is a useful result, since it suggests that unaided human performance may not be as fallible as Krolak et al (1971) assumed. In all of our experiments, participants produced only a single solution to each problem. Whether they are able to improve on their solutions with repeated attempts at the same problem is an interesting avenue for further research. What is important, in respect of Krolak et al's (1971) approach, is that in all our experiments participants produced their solutions within a 5 min time limit. The presentation

conditions did not allow for the recording of individual solution times, though it was apparent to the experimenters that many participants completed their tours within 2 min. Thus, our participants' unaided performance seems to compare favourably on measures of both time and path length with the procedure of Krolak et al.

Regarding the effects of spatial variation in TSP arrays, the results of experiments 2 and 3 indicate that a high level of human performance with TSPs is not solely contingent on the specific problem layouts used in previous research by our group. Human participants are able to capitalise upon the presence of regularity in a point array, and—given the findings of Ormerod and Chronicle (1999)—they probably use perceptual hypotheses about good figure in order to do this. One should perhaps not read too much into the frequent finding of optimal tours with the regular TSP in experiment 2 relative to performance with random arrays of experiment 1, since it is feasible that the regular array simply offers a large proportion of optimal tours out of the total set of possible tours (a possibility that is difficult to establish given the vast number of possible tours around a 48-point problem). What is clear, however, is that humans can utilise the regularity of the array in a way that would not be possible for computational methods without the development of highly constrained and problem-specific algorithms. Experiment 3 demonstrated further that problem arrays with an obvious Gestalt figure gave rise to better solutions than problems in which there was either no obvious Gestalt figure or several alternative interpretations. Further support for the importance of initial perceptual information in solving TSPs comes from a recently developed computational model for TSPs (MacGregor et al 1999) where local point-by-point decisions are made in the context of, and guided by, a global framework. The model has proved to be highly successful in simulating human performance on randomly generated TSPs. Taken together with the results of experiments 2 and 3 this indicates, first, that global-to-local contingency might operate with more complex stimuli than previously investigated by Sanocki and others, and, second, that the generation of a global hypothesis can be aided by the presence of figural regularity. The procedures used in the experiments reported here did not permit the detailed examination of strategy development and use by individual participants, and it remains an interesting question whether participants may develop and use idiosyncratic methods for completing TSPs. It is notable, for example, in experiment 3, that two kinds of suboptimality may be distinguished: those where the global figure is incorrectly drawn (leading to major path-length variations), and those where the tour of interior points is distorted (leading to minor path-length variations). It remains possible that alternative experimental procedures, perhaps without time constraints, would permit the emergence of complementary strategies for examining local figural features.

Our exploration of context effects in experiment 1 was limited to the presentation of contexts implied through visual and instructional cues. The results indicate that contextual cues such as these do not affect TSP-solving performance, which does not confirm the data from Buckmaster's (1992) pilot study. This is, in one sense, good news, since a detrimental effect of context would have had serious consequences for attempts to apply human performance in practical settings. On the other hand, context often plays a more direct role in realistic TSPs, by changing the priority, availability, or salience of subsets of points relative to others in a problem array (eg in a realistic 'distribution' context, it may be necessary to deliver goods to some sites before others). Whether the high standard of human performance can be maintained in the presence of this kind of contextual factor remains to be explored.

It is interesting to consider the implications of the findings reported in this paper for theories of spatial cognition and perception more generally. It seems plausible to argue that global perceptual processing may be used by humans in solving problems that are homomorphic to the TSP, for example, in route-finding (eg Anderson et al 1993)

and statistical assessment of data distributions (eg Lloyd and Steinke 1977; Walter 1993). Cognitive architectures such as ACT-R (Anderson 1993) and SOAR (Newell 1991) do not address the influence that global perceptual processing might have upon problem-solving performance. Others (eg Larkin 1989; Zhang and Norman 1994; Zhang 1997) have recognised the importance of external displays in mediating cognition. However, they focus upon explaining how external displays constrain the operation of local processing strategies. We suggest that recognising 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 and spatial cognition.

References

- Anderson J R (Ed.), 1993 *Rules of the Mind* (New York: Lawrence Erlbaum Associates)
- Anderson J R, Kushmerick N, Lebiere C, 1993 "Navigation and conflict resolution", in *Rules of the Mind* Ed. J R Anderson (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 93–120
- Atneave F, 1982 "Prägnanz and soap-bubble systems: A theoretical exploration", in *Organization and Representation in Perception* Ed. J Beck (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 11–29
- Boyd S C, Pulleyblank W R, Cornuéjols G, 1987 "TRAVEL—An interactive travelling salesman problem package for the IBM personal computer" *Operations Research Letters* **6** 141–143
- Buckmaster J K, 1992 *Evaluating Human Performance on the Travelling Salesman Problem Using a Hypercard Based System* MSc thesis, Loughborough University, Loughborough, UK
- Golden B, Bodin L, Doyle T, Stewart W, 1980 "Approximate traveling salesman algorithms" *Operations Research* **28** 694–721
- Krolak P, Felts W, Marble G, 1971 "A man-machine approach toward solving the traveling salesman problem" *Communications of the ACM* **14** 327–334
- Larkin J, 1989 "Display-based problem solving", in *Complex Information Processing* Eds D Klahr, J Kotovsky (Boston, MA: MIT Press) pp 319–341
- 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 C, 1996 "Human performance on the travelling salesman problem" *Perception & Psychophysics* **58** 527–539
- MacGregor J N, Ormerod T C, Chronicle E P, 1999 "A model of human performance on the travelling salesperson problem", in *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* Eds M Hahn, S C Stoness (Mahwah, NJ: Lawrence Erlbaum Associates) p. 804
- Newell A, 1991 *Unified Theories of Cognition* (Cambridge: Cambridge University Press)
- Ormerod T C, Chronicle E P, 1999 "Global perceptual processes in problem-solving: The case of the traveling salesperson" *Perception & Psychophysics* in press
- Pomplun M, Ritter H, 1999 "A three-level model of comparative visual search", in *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* Eds M Hahn, S C Stoness (Mahwah, NJ: Lawrence Erlbaum Associates) pp 543–548
- Sanocki T, 1993 "Time course of object identification: Evidence for a global-to-local contingency" *Journal of Experimental Psychology: Human Perception and Performance* **19** 878–898
- Walter S D, 1993 "Visual and statistical assessment of spatial clustering in mapped data" *Statistics in Medicine* **12** 1275–1291
- Zhang J, 1997 "The nature of external representations in problem-solving" *Cognitive Science* **21** 179–218
- Zhang J, Norman D A, 1994 "Representations in distributed cognitive tasks" *Cognitive Science* **18** 87–122

