

# A model of human performance on the traveling salesperson problem

J. N. MACGREGOR

University of Victoria, Victoria, British Columbia, Canada

and

T. C. ORMEROD and E. P. CHRONICLE

Lancaster University, Lancaster, England

A computational model is proposed of how humans solve the traveling salesperson problem (TSP). Tests of the model are reported, using human performance measures from a variety of 10-, 20-, 40-, and 60-node problems, a single 48-node problem, and a single 100-node problem. The model provided a range of solutions that approximated the range of human solutions and conformed closely to quantitative and qualitative characteristics of human performance. The minimum path lengths of subjects and model deviated by average absolute values of 0.0%, 0.9%, 2.4%, 1.4%, 3.5%, and 0.02% for the 10-, 20-, 40-, 48-, 60-, and 100-node problems, respectively. Because the model produces a range of solutions, rather than a single solution, it may find better solutions than some conventional heuristic algorithms for solving TSPs, and comparative results are reported that support this suggestion.

The Euclidean form of the traveling salesperson problem (TSP) consists of finding the shortest path (*tour*) that passes through a set of points and returns to the origin. Any particular instance of a TSP obviously has a finite set of solutions, so there is a guaranteed method of solving any problem by exhaustive search of the solution space. However, exhaustive search becomes increasingly impractical as the number of points ( $n$ ) increases. This is because the number of possible solutions increases as  $(n - 1)!/2$  (ignoring direction of travel). Thus, although a computer generating 1,000 solutions per second would require only 3 min to find all the solutions to a 10-node problem, a 20-node problem would occupy it for almost 4 million years. The increased capacity of computers over the past 10 years has made it possible to find guaranteed optimal solutions to fairly large TSPs within reasonable time periods. In 1986, a new record was set (since broken) when a 2,392-city problem was solved, breaking the previous record of 532 cities and requiring 27 h (Sangalli, 1992). There has been a long-term interest in the field of operations research in devising economical procedures for finding solutions that approximate the optimal, and many such heuristic techniques have been developed and compared. To achieve "reasonable" approximations to an optimal solution to within a few percentage points above the shortest path, such procedures generally need to perform on the order of  $n^3$  calculations (Golden, Bodin, Doyle, & Stewart, 1980).

One approach to improving heuristic procedures has been to enlist the assistance of human operators (Krolak, Felts, & Marble, 1971; Michie, Fleming, & Oldfield, 1968). Krolak et al. compared computer-generated solutions with solutions produced by a human-computer interactive approach. A number of comparisons were reported that indicated that human-computer solutions were at least as short, and often shorter, than those generated by computer alone. The results imply that human operators respond to TSPs in a manner that is not only *different* from, but more *effective* than, the heuristics used by Krolak et al. Otherwise, the human-computer solutions should have been the same as, or poorer than, those generated by computer only. This inference is supported by results reported by MacGregor and Ormerod (1996), indicating that the best human solutions to a range of 10-city and 20-city problems were often superior to those generated by several heuristic procedures. The results of these researches suggest that it may be worthwhile to investigate how humans solve TSPs and, if feasible, to incorporate elements of the human approach into heuristic algorithms. The present article represents a step in this direction. The article draws on characteristics of human solutions suggested by MacGregor and Ormerod to identify a model of human TSP-solving behavior. The model is tested, using previously published and unpublished human data.

MacGregor and Ormerod (1996) reported the results of two experiments with small TSPs (10 and 20 nodes) that indicated the following characteristics of human performance: (1) people produce "good" solutions, with the best human performances being, on average, within 0.75% of the optimal solutions; (2) differences, if any, in skill levels across individuals are slight, in that the average correlation of the rank order of individuals' solutions across

---

Correspondence concerning this article should be addressed to J. N. MacGregor, School of Public Administration, University of Victoria, P. O. Box 1700, Victoria, BC V8W 2Y2, Canada (e-mail: jmacgreg@uvic.ca).

problems is virtually zero; (3) the complexity of problems varies directly with the number of “internal” nodes falling within the perimeter of a problem, not with the total number of nodes; (4) people produce paths that connect boundary points in sequential order of adjacency (449 of 455 solutions adhered to this principle); (5) few solutions have lines that cross (11 of 455 solutions); and (6) solutions have significantly fewer “indentations” than the average allowed by the structure of the problems. Ormerod and Chronicle’s (1999) results further suggest that some form of global processing may underlie human solutions, and taken together, the results provide evidence as to the kinds of processes that generate human solutions to TSPs. Below, we discuss the nature of those processes and propose a model of human TSP performance. The model is tested, using data from previous experiments.

### HEURISTIC PROCEDURES


The model is a synthesis of features taken from a number of approximate heuristic procedures that have been developed in the field of operations research. Many such procedures have now been designed and compared (Golden et al., 1980), and there are several that might conceivably serve as models of human performance. One of the simplest is the *nearest neighbor* (NN) approach, which is the apocryphal model of human performance. For the NN method, an initial point is selected as the start of the path (the leading point). The unconnected point closest to the leading point is then added to the path and becomes the leading point. This repeats until all the nodes have been added; then, the last and first points are joined to complete the tour. Another type of heuristic follows a *convex hull* approach, in which an initial subtour is obtained by connecting all the points lying on the perimeter of the array of nodes (the convex hull) and, thereafter, interior points are included in the subtour sequentially in an order that adds minimally to the incremental distance. A simple analogue of the process is to think of the nodes as pegs in a board, with the initial subtour formed by stretching a rubber band around their perimeter. Interior pegs are then incorporated by stretching the segment of the band closest to an interior peg over that peg and then repeating this until all of the pegs have been included. Convex hull heuristics differ in their choice of *insertion criterion*—that is, in how they define *closest*. For example, the *cheapest insertion* criterion finds the points  $i, j$  in the subtour and  $k$  not in the tour for which the distance  $i, k + k, j - i, j$  is minimal, and  $k$  is inserted between  $i$  and  $j$  (Golden et al., 1980). Norback and Love (1977) used a largest interior angle criterion, finding the unconnected point  $k$  for which the angle  $i, j, k$  is maximal and inserting  $k$  between  $i$  and  $j$ . In general, the convex hull approach, as a heuristic for finding good solutions, has the advantage of capitalizing on one of the well-known characteristics of an optimal solution. That is, the optimal path connects adjacent points on the boundary of the convex hull in sequence, although it may pass through interior points between adja-

cent boundary points (Flood, 1956). Convex hull heuristics conform to this principle. A corollary is that solutions that do not adhere to this principle will result in crossed arcs, which is clearly nonoptimal.

### ANALYSIS

The experimental findings summarized previously indicated that human solutions show a high degree of consistency in a number of respects. Almost invariably, subjects connected boundary points in the order of adjacency and, equally invariably, produced no crossed arcs. The quality of their solutions was influenced by the number of interior points to be incorporated. They generated solutions with few indentations. These consistent elements, together with the fact that proficiency appeared to be equally distributed across individuals, 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. Like convex hull heuristics, human solvers join points on the boundary in order of adjacency, avoid crossing lines, and seem to be capable of producing optimal or near-optimal solutions with a high degree of regularity. The fact that they produce solutions with few indentations and that the quality of their solutions erodes with an increasing number of interior points further suggests that, in generating solutions, they may like to stay close to the convex hull. In addition, the idea that human solutions may be governed by the outline of the array is consistent with a global-processing approach to the problem (Ormerod & Chronicle, 1999). Nevertheless, there are elements of performance that are inconsistent with convex hull heuristics. Subjects appear to complete problems in a sequential way, with the terminal node of one connection becoming the origin node of the next. Of the three heuristics described above, only the NN approach operates in this way. In contrast, the convex hull heuristics may jump between nonadjacent arcs, seeking the next best move in a manner seemingly quite uncharacteristic of human performance. The NN approach is similar to human performance in another respect, in that it can produce a range of solutions to a single problem, whereas the other two heuristics typically produce a single solution. It appears from these observations that human solution processes may incorporate characteristics of each of these heuristics while being identical to none. That is, a previously unidentified heuristic may guide human solutions.

MacGregor and Ormerod (1996) compared the performances of these three heuristics with human solutions, using six new 10-node problems, a standard 10-node problem, and seven 20-node problems. The results were consistent with the foregoing analysis. For the six 10-node problems, the two convex hull heuristics generated only 1 solution each per problem. The NN procedure was repeated for each possible starting point and, therefore, could produce up to 10 different solutions per problem. In terms of the best solutions found, the averages for the



**Figure 1.** Illustration of the model's solution to a traveling salesperson problem when starting from Point 1 and progressing in a counterclockwise direction.

NN, cheapest insertion, and largest interior angle procedures were, respectively, 1.9%, 0.9%, and 4.3% above the optimal. This compared with 0% for the best human solutions. Only the human and the NN approaches produced a range of solution path lengths. With respect to these distributions, the mean of human performances on each problem, averaged across problems, was 3.8%, as compared with 6.0% for the NN approach. The experiment also employed the 10-node problem described by Dantzig, Fulkerson, and Johnson (1959). For this problem, the best human solution was optimal, as compared with solutions of 0.6%, 2.7%, and 3.0% above optimal for the NN, cheapest insertion, and largest interior angle approaches, respectively. The average of the human solutions was 3.8% above optimal, as compared with 10.4% for the NN approach. Similar results were obtained with the seven 20-node problems. Averaged across the problems, the best solutions for the NN, cheapest insertion, and largest interior angle procedures were, respectively, 3.6%, 4.1%, and 11.0% above optimal, as compared with 1.5% for the best human solutions. The mean solutions for humans and the NN procedure were 6.3% and 11.5%, respectively. The results seem to indicate that, for these small problems at least, the best human performances were consistently superior to any of the heuristics and that the average for the human solutions was substantially superior to the average for the NN heuristic (the only heuristic to provide a range of solutions).

### TOWARD A MODEL OF HUMAN PERFORMANCE

As they stand, none of the three heuristics appears able to provide an adequate model of the human solution process. As a first step to creating a more adequate model,

an effort was made here to incorporate the psychologically attractive elements of each of the heuristics. It appeared that there were three such elements. First, the model should retain a convex hull basis, since this ensures that boundary points are connected in order. Second, the procedure should be sequential, since this conforms to what subjects do. Third, it should have the potential to produce different paths when starting on different points, like the NN procedure.

The model<sup>1</sup> described below may be used with different insertion criteria (e.g., cheapest insertion or largest angle) and is described initially in a general way. For convenience, the term *close* is used to refer to the decision produced by the insertion rule. Informally, the model operates as follows. Initially, the arcs between adjacent boundary points are *sketched*—not connected, but used as a guide and reference for subsequent judgments. Next, a starting point and a direction (clockwise or counterclockwise) are selected randomly. If the start is an interior point, it is connected immediately to the closest point on the boundary in the direction of travel. From here, using the arc sketched between this point and the adjacent boundary point as a reference, the closest interior point is identified. If this point is closer to any other sketched or connected arc, it is *passed*, and the path moves to the next adjacent node on the boundary in the direction of travel. This continues until an interior point is found that is not closer to any other arc. The current node is then connected to this point but remains the current node. Using this newly created arc as reference, the process is repeated until a complete tour is obtained.

To illustrate, consider how the procedure performs on the problem shown in Figure 1. First, the arcs joining the points on the boundary of the hull are sketched; then, a starting point and a direction are selected at random, say starting with Point 1 and moving in a counterclockwise direction. The current arc is therefore 1–2, and Point 7 would be considered for inclusion. However, the test would reveal that Point 7 is closer to another arc, 5–1, and would be passed. Point 6 would then be considered but, being closer to Arc 2–3, would also be passed. This exhausts candidates for inclusion between Points 1 and 2, and the algorithm would move on to Point 2, and Arc 2–3 would become the current arc. Point 7 would be considered for inclusion but, being closer to Arc 5–1, would be passed over. Next Point 6 would be considered, would pass the test, and would be connected between Points 2 and 3. That is, Arc 2–3 would be erased, and two new arcs, 2–6 and 6–3, would be created. Arc 2–6 now becomes the current arc, and Point 7 would be considered for inclusion between Points 2 and 6. It would fail the test, and having exhausted all free points from this arc, the procedure will move on to Point 6, and Arc 6–3 becomes the current arc. This would proceed until Arc 5–1 becomes the current arc and Point 7 is inserted between 5 and 1. In this case, this completes the tour, with the path 1, 2, 6, 3, 4, 5, 7, 1. A more detailed description is given in the steps below.

*Step 1:* Sketch the connections between adjacent boundary points of the convex hull.

**Table 1**  
**Minimum, Mean, and Maximum Path Lengths as a Percentage**  
**Deviation From the Comparable Experimental Results for the**  
**Cheapest Insertion (CI) and Largest Angle (LA) Versions of the Model**

	Number of Nodes	Problem Number	Path Lengths					
			Minimum		Mean		Maximum	
			CI	LA	CI	LA	CI	LA
Experiment 1	10	1	0.0	0.0	-0.8	-0.8	-8.8	-8.8
		2	0.0	0.0	-3.4	-3.4	-18.8	-18.8
		3	0.0	5.2	-1.0	3.1	-11.2	-7.7
		4	0.0	0.0	1.0	6.4	-0.5	2.9
		5	0.0	1.0	-0.5	0.5	-7.3	-5.8
		6	0.0	1.0	-3.0	-1.4	-12.1	-8.2
		Mean absolute difference	0.0	1.2	1.6	2.6	9.8	8.7
Mean difference	0.0	1.2	-1.3	0.7	-9.8	-7.7		
Experiment 2	20	1	-0.5	0.0	-2.6	-2.2	-13.3	-13.0
		2	0.4	1.9	-1.9	3.1	-10.4	2.6
		3	0.0	1.1	-0.2	3.2	-5.9	0.9
		4	-1.4	-1.4	4.8	6.6	7.2	19.3
		5	3.3	8.9	2.3	11.2	-6.0	1.1
		6	3.2	6.1	-0.5	2.7	-10.2	-1.9
		7	-0.7	1.0	-6.2	-0.3	-16.1	-5.2
Mean absolute difference	1.4	2.9	2.6	4.2	9.9	6.3		
Mean difference	0.6	2.5	-0.6	3.5	-7.8	3.8		

*Step 2:* Select a starting point and a direction randomly.

*Step 3:* If the starting point is on the boundary, the starting node is the *current node*. The arc connecting the current node to the adjacent boundary node in the direction of travel is referred to as the *current arc*. Proceed immediately to Step 4. If the starting point is not on the boundary, apply the insertion rule to find the closest arc on the boundary. Connect the starting point to the end node of the closest arc which is in the direction of travel. This node becomes the current node.

*Step 4:* Apply the insertion criterion to identify which unconnected interior point is closest to the current arc. Apply the insertion criterion to check whether the closest node is closer to any other arc. If not, proceed to Step 5. If it is, move to the end node of the current arc. This becomes the current node. Repeat Step 4.

*Step 5:* Insert the closest node. The connection between the current node and the newly inserted node becomes the current arc. Retaining the current node, return to Step 4 and repeat Steps 4 and 5 until a complete tour is obtained.

## TESTS OF THE MODEL

### Comparisons With Data From 10- and 20-Node Traveling Salesperson Problems

Two versions of the model were tested initially. One used the cheapest insertion criterion, the other the largest interior angle criterion. The first tests used the stimuli and results from MacGregor and Ormerod (1996), who reported results for 45 subjects in their first experiment and for 20 in the second. Both experiments used a within-subjects design, and in both cases the subjects were in-

structed to select a starting point and then draw the shortest path that passed through each point and returned to the start. The subjects were tested in a group setting and were instructed to take no more than 5 min per problem. In testing the model, 45 replications were conducted with each basic problem from Experiment 1 to give an equal number of model and human solutions. Similarly, 20 replications were conducted for the problems from Experiment 2. The main results are presented in Table 1, which shows the minimum, mean, and maximum path lengths obtained with both versions of the model as a percentage deviation from the comparable experimental results (i.e., model results minus experimental results divided by experimental results, expressed as a percentage). Summary statistics are provided in two forms, the mean absolute values of the deviations and the mean algebraic values. The former provides a measure of goodness of fit, the latter an indication of whether the model tends to overestimate (positive values) or underestimate (negative values) the corresponding human results.

For the 10-node problems, it can be seen that both versions of the model produced minimum path lengths similar to the experimental results, although the cheapest insertion criterion appeared to give a more consistent performance, with a maximum deviation of 0.0%, as compared with 5.2% for the largest angle criterion. Looking next at the mean path lengths, a similar picture emerges. Overall, both criteria produced solutions close to the experimental results, but more consistently with the cheapest insertion criterion than with the largest angle criterion. Finally, the maximum path lengths resulted in the poorest fit for both criteria, with deviations in both cases as

**Table 2**  
**Minimum Path Lengths for the Cheapest Insertion (CI)**  
**and Largest Interior Angle (LA) Versions of the Model**  
**and Their Percentages of Deviation From Human**  
**Performance for 10-, 20-, 40-, and 60-Node Problems**

Problem	Number of Nodes	CI Model		LA Model	
		Best Path	Deviation (%)	Best Path	Deviation (%)
1	10	541	0.0	541	0.0
2	10	495	0.0	495	0.0
3	10	541	0.0	543	0.2
Mean			0.00		0.07
4	20	650	0.0	650	0.0
5	20	708	-0.1	745	5.10
6	20	701	0.0	711	1.4
7	20	596	0.0	606	1.7
Mean			-0.03		2.05
8	40	1,000	0.5	1,106	11.2
9	40	919	1.7	1,014	12.2
10	40	1,006	3.4	1,076	10.5
11	40	906	4.0	990	13.7
Mean			2.39		11.91
12	60	1,114	1.2	1,248	13.4
13	60	1,143	5.8	1,171	8.5
Mean			3.53		10.92

extreme as -19%. In general, the models produced maximum path lengths that were shorter than the worst human solutions.

For the 20-node problems, a similar pattern emerged. For minimum and mean path lengths, the cheapest insertion criterion produced solutions closer to the human results than did the largest angle criterion, although with slightly less accuracy than with the 10-node problems. Again, the poorest fit occurred for maximum path lengths, with both insertion criteria producing worst cases that were better than the corresponding worst human solutions.

Overall, the results suggest that the model may provide a promising first step to modeling human performance. From the results so far, the cheapest insertion criterion appears to be somewhat better than the largest angle criterion for approximating human solutions, although it does not appear to provide as wide a range of path lengths at the upper end.

#### Comparisons With Data From 10-, 20-, 40-, and 60-Node Traveling Salesperson Problems

MacGregor and Ormerod (1996) reported results, using data for 13 TSPs from Lee (1985). The data were collected from 50 subjects (graduate and undergraduate students), who produced solutions to a variety of randomly generated TSPs, ranging from 10 to 60 nodes. Several different random problems were used at each level of number of nodes. This was employed as a within-subjects factor, whereas number of nodes was a between-subjects factor. The subjects were tested in group settings and were instructed to draw the shortest path passing through all of the points and returning to the starting point. In his paper, Lee reported only the best human solutions, and so the fol-

lowing tests are necessarily restricted to comparing the best model solutions with these. In conducting the tests, 50 replications of both versions of the model were employed. The results are summarized in Table 2. For each problem, the table shows the number of nodes and the best paths produced by both versions of the model as (1) a distance and (2) a deviation from the corresponding best human solution.

For the 10- and 20-node problems, the results were very similar to those reported above. In both cases, both insertion criteria produced best solutions close to the human results, with mean percentage deviations of 0.00% and 0.07% for the 10-node problems and -0.03% and 2.05% for the 20-node problems for the CI and LA criteria, respectively.

For the larger problems, the cheapest insertion criterion continued to provide the better fit. On average, the deviations were 2.39% and 3.53% for the 40- and 60-node problems, respectively, whereas the corresponding results for the LA procedure were 11.91% and 10.92%, an order of magnitude poorer.

Two additional trends are discernible in the results. The goodness of fit of the models seems to decrease slightly as the size of problem increases, and the direction of error (for minimum path lengths) becomes consistently positive, indicating that the model's best solutions may become poorer than the best human performances. Nevertheless, the fit for minimum path lengths remains relatively good for problems up to 60 nodes. Combined with the trends observed above, the indications are that the model will tend to produce a relatively good fit at the lower path lengths, but a narrower range of solutions than human performances. It appears that this tendency may be-

come more pronounced as the size of problem increases. This was examined below, using data for a 100-node problem. Because the cheapest insertion criterion seems to result in a better approximation to human performance, subsequent tests are limited to this version.

#### **Comparisons With Data From a 100-Node Traveling Salesperson Problem**

The following test used previously unpublished data from 8 undergraduate students who produced solutions to a single randomly generated 100-node TSP described in Krolak et al. (1971). The subjects were tested in a group setting and instructed to choose a starting point and to draw the shortest path from that point, passing through each point and returning to the start. For comparison, 40 model solutions were generated.

The paths generated by the subjects ranged in length from 23,788 to 53,932, with a mean of 31,700. The corresponding model values ranged from 23,784 to 31,211, with a mean of 26,940. The model's best path was virtually identical to the best human solution (0.02% shorter), whereas the mean and the worst solutions were 15.0% and 42.1% shorter, respectively. The results were fairly consistent with the trends identified above, although the model's fit to the human solutions was better than would be expected from the previously observed trend for the minimum path length, although being somewhat poorer than would be expected for the mean and the maximum. The fact that the model's mean and maximum path lengths were considerably shorter than the corresponding human figures may have arisen because of the extremely poor performance of 1 subject, whose solution of 53,932 was 5.9 standard deviations above the overall subject mean of 31,700. If this person's solution is excluded, the deviations of the model from minimum, mean, and maximum human path lengths were  $-0.02\%$ ,  $-5.6\%$ , and  $-16.8\%$ , respectively. The pattern was similar to the trend identified above, with the model providing a very good fit to the minimum score and a good fit to the average, but overall providing a narrower range than the human solutions.

#### **Comparisons With Data From a 48-Node Traveling Salesperson Problem**

Data from MacGregor, Ormerod, and Chronicle (1999) from a single TSP of 48 nodes provided a more detailed comparison of human and model performances. The data consisted of 103 human solutions to a problem created by randomly selecting 48 points from a 100-node problem described by Krolak et al. (1971). The subjects were tested in a group setting, using the same instructions as those in MacGregor and Ormerod (1996). A set of 100 model solutions was generated by randomly selecting starting points and direction of travel, clockwise or counterclockwise. Model solutions were generated, using only the cheapest insertion version, which appears to provide the better approximation.

The results were as follows. The best human solution had a path length of 553. The best model solution was

561, 1.4% above this. The mean of the human solutions was 606, and that of the model was 602, 0.7% below the human mean. The worst human and model solutions were 697 and 644, respectively, the model worst case being 7.6% better than the human. The results were very similar to those reported earlier in the article, with the model providing a good fit to average and best human solutions but underestimating the worst human solution (in this case, by 7.6%).

The subjects in this experiment were asked to indicate the point at which they started, and 99 of the 103 did so. This permits a more sensitive test of the model, since model solutions can vary, depending on the starting point selected. The subjects indicated a total of 22 different starting points. Seven were unique, whereas as many as 19 of the subjects (20%) chose the most popular. (Incidentally, 72% of the human solutions started on the hull, as compared with the 29% that would be expected on the basis of the availability of hull points.) Model solutions were generated from each of the 22 starting points indicated by the subjects, for each of the two possible directions of travel. A model score for each starting point was obtained by averaging these two scores. The mean subject path lengths across the 22 starting points ranged from 567 to 628. The differences between the subject and the model path lengths for each starting point were expressed as a percentage of the subject path length. The absolute values of these percentage deviations between human and model results ranged from 0% to 5.6%, with a weighted average of 1.1%.

The goodness of fit of the model seems to vary, depending on whether solutions started on hull or interior points, as may be seen by the pattern of outcomes in Table 3. In the case of hull starts, the model provides a very close fit to the human path lengths, ranging from 0% to 3%, with an average absolute deviation of 0.6%. In the case of interior starts, the fit is consistently poorer, ranging from  $-6\%$  to  $+5\%$ , with an absolute average of 2.4%. These differences in model fit perhaps arise because the human solutions cover a greater range for interior than for hull starts, and it may be that the model simply provides a better fit over the mid-range. Alternatively, the model's way of dealing with interior starts—going directly to the hull on the next move—is likely to be too simple and may fail to reflect what people do after starting on an interior point.

## **DISCUSSION**

The present article proposed and tested a model of human performance on TSPs. Because of empirical evidence that people are influenced by the convex hull in generating solutions to TSPs (MacGregor & Ormerod, 1996; Ormerod & Chronicle, 1999), the model was designed to conform in a general way to a convex hull approach. However, it differs from conventional convex hull heuristics by generating solutions from a given starting point and progressing in a specified direction (clock-

**Table 3**  
**A Comparison of Human and Model Path Lengths for a 48-Node Problem,**  
**for Starting Points on the Hull Versus Interior Starting Points**

Hull Start			Interior Start		
Human Path Length	Model Path Length	% Difference	Human Path Length	Model Path Length	% Difference
595	597	0	567	599	-6
598	598	0	589	597	-1
603	605	0	590	606	-3
607	607	0	590	589	0
607	607	0	591	610	-3
607	607	0	595	602	-1
607	603	1	608	595	2
609	607	0	609	599	2
610	610	0	616	602	2
616	597	3	619	608	2
620	603	3	628	598	5

wise or counterclockwise). This allows the model to generate a range of solutions for a given problem, since the solution generated can vary, depending on starting point or direction. In this respect the model is unlike other convex hull approaches, in particular, and heuristic algorithms, in general, which typically generate a single solution only. An exception is the NN procedure, which may generate different solutions, depending on starting point. However, previous results with human subjects indicated that their solutions differed qualitatively and quantitatively from the NN approach (MacGregor & Ormerod, 1996).

The proposed model was tested by using previously reported and unreported data on human performance with a variety of problems ranging from 10 to 60 nodes, plus one 48-node problem and one 100-node problem (MacGregor & Ormerod, 1996; MacGregor et al., 1999). To a large extent, the results of the model conformed closely to those of the human subjects, both quantitatively and qualitatively. Two versions of the model were tested, one using a cheapest insertion criterion, the other a largest interior angle criterion, and although both versions performed reasonably well, the results favored the former. Nevertheless, there were consistent indications that the model produced a narrower range of solutions than did the human subjects, particularly in the upper range of path lengths. That is, the worst solutions generated by the model tended to be better than the worst human solutions.

There are several reasons why this may have occurred. At one extreme, it may indicate that subjects use a different approach than the model. However, the overall goodness of fit between the model and the human solutions argues against this. Perhaps people use the same general approach as the model, but with a different insertion criterion. The present model results indicate that the insertion criterion can make a considerable difference to the outcome.

Another potential difference between the model and human performance is that the latter will be subject to a

degree of perceptual error in the judgment of distances, angles, or whatever other factors provide the basis for decisions. This suggests that introducing a degree of random (or systematic) error into the model's judgments may produce a wider range of solutions, which may correspond even more closely to the range of human solutions than do the present results.

Another qualification to the model's generally good performance arises because of human solvers' sensitivity to pattern and regularity. For example, figural factors, such as the proximity of interior points and the regularity of their arrangement, appear to influence human solutions (MacGregor & Ormerod, 1996; MacGregor et al., 1999). There are no mechanisms within the present model to respond to factors such as these, so the model is unlikely to produce a particularly good fit to human solutions to highly patterned TSPs. However, for random or relatively irregular TSPs, the present results suggest that the model provides a very reasonable approximation to the solutions that people produce.

Finally, the present approach of modeling human performance could be of value in identifying means of improving the solutions generated by conventional heuristic algorithms. For example, MacGregor and Ormerod (1996) tested a convex hull heuristic and the largest interior angle method (Norback & Love, 1977) on their 13 random problems and reported that, on average, the best path lengths produced by these procedures were, respectively, 2.5% and 7.6% longer than the optimal solutions. By comparison, the best solutions of the model described here were 1% above optimality across the same 13 test problems. The experiment also employed the highly structured 10-node problem described by Dantzig et al. (1959), and for this problem the best solutions were 2.7% and 3.0% above optimal for the cheapest insertion and the largest interior angle approaches. By comparison, the present procedure found the optimal solution. These improvements over conventional procedures arise because the present approach produces a range of solutions, rather

than a single solution, and although many of these may be worse than those produced by conventional heuristics, some may be better.

#### REFERENCES

- DANTZIG, G. B., FULKERSON, D. R., & JOHNSON, S. M. (1959). On a linear-programming, combinatorial approach to the travelling-salesman problem. *Operations Research*, *7*, 58-66.
- FLOOD, M. M. (1956). The travelling salesman problem. *Operations Research*, *4*, 61-75.
- GOLDEN, B., BODIN, L., DOYLE, T., & STEWART, W. (1980). Approximate travelling salesman algorithms. *Operations Research*, *28*, 694-711.
- KROLAK, P., FELTS, W., & MARBLE, G. (1971). A man-machine approach toward solving the travelling salesman problem. *Communications of the ACM*, *14*, 327-334.
- LEE, R. K. L. (1985). *A heuristic approach to the travelling salesman problem* (Unpublished Management Report). University of Victoria, School of Public Administration.
- MACGREGOR, J. N., & ORMEROD, T. (1996). Human performance on the travelling salesman problem. *Perception & Psychophysics*, *58*, 527-539.
- MACGREGOR, J. N., ORMEROD, T. C., & CHRONICLE, E. P. (1999). Spatial and contextual factors in the travelling salesperson problem. *Perception*, *28*, 1417-1428.
- MICHIE, D., FLEMING, J. G., & OLDFIELD, J. V. (1968). A comparison of heuristic, interactive and unaided methods of solving a shortest-route problem. In D. Michie (Ed.), *Machine intelligence 3* (pp. 245-255). New York: Elsevier.
- NORBACK, J. P., & LOVE, R. F. (1977). Geometric approaches to solving the travelling salesman problem. *Management Science*, *23*, 1208-1223.
- ORMEROD, T. C., & CHRONICLE, E. P. (1999). Global perceptual processes in problem solving: The case of the traveling salesperson. *Perception & Psychophysics*, *61*, 1227-1238.
- SANGALLI, A. (1992, December 12). Why sales reps pose a hard problem. *New Scientist*, *136*, 24-28.

#### NOTE

1. Tests of a similar model are reported by Lee (1985), but as a method of generating solutions to TSPs, not as a model of human performance. Also involved in the development of the model reported by Lee were N. Lam, E. Lee, and J. MacGregor.

(Manuscript received March 31, 1999;  
revision accepted for publication January 20, 2000.)