

A Comparison of Heuristic and Human Performance on Open Versions of the Traveling Salesperson Problem

James N. MacGregor,¹ Edward P. Chronicle,² and Thomas C. Ormerod³

Abstract

We compared the performance of three heuristics with that of subjects on variants of a well-known combinatorial optimization task, the Traveling Salesperson Problem (TSP). The present task consisted of finding the shortest path through an array of points from one side of the array to the other. Like the standard TSP, the task is computationally intractable and, as with the standard TSP, people appear to be able to find good solutions with relative ease. The three heuristics used mechanisms that have been cited as potentially relevant in human performance in the standard task. These were: convex hull, nearest neighbor, and crossing avoidance. We compared heuristic and human performance in terms of lengths of paths, overlap between solutions, and number of crossings. Of the three heuristics, the convex hull appeared to result in the best overall fit with human solutions.

Introduction

Over the last decade, a new area of interest has emerged in the study of problem solving, concerned with how people perform on instances of mathematically intractable problems known as optimization problems. An example is the classic Traveling Salesperson Problem (TSP), where the task is to find the shortest closed path that passes through a set of points and returns to the origin. Despite a sustained research effort from the fields of operations research and computer science, no general method for solving this and related problems in reasonable (polynomial) time has been found, and there is a growing suspicion that none may exist. At the same time, many ingenious heuristic procedures have been devised that are capable of finding close approximations to optimal solutions at reasonable computational costs but cannot guarantee optimality (Golden, Boden, Doyle, & Stewart, 1980).

¹School of Public Administration, University of Victoria, British Columbia, V8W 2Y2, Canada

²Department of Psychology, University of Hawaii at Manoa, 2430 Campus Road, Honolulu, HI 96822 U.S.A.

³Psychology Department, Lancaster University, Fylde College, Lancaster, LA1 4YF, U.K.

Correspondence concerning this article should be addressed to: James MacGregor, Telephone: 250 721 6435, Email: jmacgreg@uvic.ca

An intriguing aspect of human solutions is that they tend to be very good and frequently outstrip the simpler heuristics' solutions. This appears to be true whether people are given very little time and have to quickly sketch a route (MacGregor & Ormerod, 1996) or they are given as much time as they wish and are permitted to erase and redraw path segments until satisfied (Vickers, Bovet, Lee, & Hughes, 2003). The fact that human solutions tend to be good is the one finding on which there is general agreement (e.g., Graham, Joshi, & Pizlo, 2000; MacGregor, Ormerod, & Chronicle, 1999). On other aspects of human performance, the research findings diverge. On the issue of whether there are individual skill differences, for example, some studies have failed to find any supportive evidence (Chronicle, MacGregor, Ormerod, & Burr, 2006). Others have reported significant correlations in performance, not only across instances of TSPs but across different optimization problems (Vickers, Mayo, Heitmann, Lee, & Hughes, 2004). It has been conjectured that the inconsistency in findings may have arisen from procedural differences (Chronicle et al., 2006) or from the presence of ceiling effects (Vickers et al., 2003).

There are also marked differences in the kinds of mechanisms proposed to underlie human performance. These fall into two broad theoretical positions characterized as global-to-local and local-to-global. The global-to-local approach comprises two models, the pyramid algorithm, based on hierarchical clustering (Graham et al., 2000) and the convex-hull model, a boundary-following heuristic (MacGregor, Ormerod, & Chronicle, 2000). (For a set of points in a two-dimensional plane, the convex-hull is the minimum polygon that contains the entire set.) The local-to-global position is represented by the work of Vickers and colleagues (Vickers, Lee, Dry, & Hughes, 2003), based on a mechanism of connecting nearest-neighbor links. A further independent process has been suggested that people follow a strategy of producing solutions that have no crossed lines (van Rooij, Stege, & Schactman, 2003). Crossing-avoidance could conceivably be implemented in a form that emphasizes either global or local processing.

Each of the positions has strengths and weaknesses. Both of the global-to-local approaches have computational models, whereas neither Vickers et al. (2003) nor van Rooij et al. (2003) propose specific mechanisms. In the case of nearest-neighbor links, it is unclear how it differs from the simple nearest-neighbor heuristic of moving to the nearest available unconnected point. Nearest-neighbor solutions are known to be inferior to subjects' (Graham et al., 2000; MacGregor & Ormerod, 1996) and can readily produce crossed lines, which are rare in human solutions (MacGregor & Ormerod, 1996; MacGregor et al., 1999; van Rooij et al., 2003).

The pyramid algorithm has the additional advantage of having been successfully applied to different problem-solving and perceptual tasks (Pizlo & Li, 2005), whereas the convex-hull approach as a specific model has been limited to the TSP, although its basic mechanism of boundary perception is of broad generality. On the other hand, the pyramid approach as a theory of TSP performance has limited empirical support, stemming

from a single study of six participants (Graham et al., 2000). The convex hull has a broader base of empirical support (Chronicle et al., 2006; MacGregor & Ormerod, 1996; MacGregor et al., 1999, 2000; Ormerod & Chronicle, 1999), including the findings that: (a) solutions tend to visit boundary points (points on the convex hull) in order of adjacency, (b) problem complexity increases as a function of the number of interior points rather than the total number of points, (c) solutions rarely have crossed lines, (d) people can judge the optimality of solutions at brief exposures, (e) solutions have few indentations (an *indentation* arises when one or more interior points is visited between adjacent boundary points), and (f) performance is better in instances where interior points are closer to the boundary. However, van Rooij et al. (2003) questioned whether this body of evidence supports a convex-hull approach and instead have proposed that avoiding crossings is the primary underlying factor. Crossing-avoidance alone would account for findings (a) to (c) and might successfully account for the remaining findings if combined with other mechanisms. However, it has been shown that a strategy of crossing-avoidance alone produces solutions that are significantly longer than subjects' (MacGregor, Chronicle, & Ormerod, 2004).

Recently, the psychology of optimization problem solving has expanded its horizons by investigating problems other than the TSP. These have included a variant of the TSP where the task is to find the shortest continuous path through a set of points from a start to an end point without the requirement to complete a circuit (Chronicle et al., 2006; Vickers et al., 2003).¹ Like the standard version of the task, this "open" version is computationally intractable. Also, the open version is particularly relevant to the theoretical issues concerning human performance discussed previously. We have argued that if performance is governed by local processes then the task of finding an open path should not differ from that of finding a closed tour. Conversely, if the solution process follows a circuit guided by the convex hull then breaking that circuit in open problems can be expected to affect performance. The balance of the experimental results support the latter position, in that three experiments have found performance to be reliably poorer in open versions than in closed versions of the same problems, whereas one experiment found no significant difference.

Nevertheless, although performance was poorer in the open versions, it was still relatively good. Across the four experiments, average lengths for open solutions ranged from 6% to 10% above the optimal, compared with 3% to 5% above the optimal for corresponding closed solutions. The results raise the questions of how people are able to generate such relatively good solutions to open versions and whether the strategies they employ are related to those that they apply to the standard TSP. This article presents an initial exploration of these issues by comparing the performance of three heuristics

¹The Vickers procedure instructed participants to start at any node and end at a different node. The Chronicle procedure provided a start node (at the extreme left or right of the array) and an end node (at the opposite extreme). Both versions are intractable (Iris van Rooij and Ulrike Stege, personal communication).

on open paths with that of subjects, on two sets of experimental results. The following section describes the heuristics and reports the two sets of comparisons with human performance. The section after that discusses the results and suggests additional mechanisms that could improve the heuristics as models of human performance.

Comparison of Heuristics with Empirical Results

Description of Heuristics

The heuristics are described as they apply to the test problems taken from Chronicle et al. (2006), in which subjects attempted to find the shortest path through a set of points, starting on one side of the display (at the leftmost or rightmost point) and proceeding to the opposite side.

Nearest-neighbor Heuristic

From the start point, connect the path to the nearest point that is not the end point. Repeat until only the end point remains; then connect to it.

Crossing-avoidance Heuristic

From the start, connect the path to the closest point on the x -axis. Repeat until reaching the end point.

Convex-hull Heuristic

From the start, connect the path to the closest point that is not the end point that can be reached without crossing the convex hull around the points not yet on the path. Repeat until only the end point remains; then connect to it.

The first two heuristics represent local-to-global approaches in that each decision requires examining a local region of the display only. In the third, processing is global-to-local because each decision depends on considering a wider characteristic of the array. We note that the crossing-avoidance heuristic is only one of many ways in which crossing avoidance might be implemented.


Comparison 1

The first comparison used the data from Experiment 1 of Chronicle et al. (2006), in which 25 subjects generated solutions to 12 different 12-node open problems. We first compared the path lengths of the heuristic and subject solutions across the 12 problems. These results are illustrated in [Figure 1](#), which shows the path lengths for the three heuristics together with the mean and maximum subject path lengths. As can be seen from the figure, the crossing-avoidance solutions exceeded the range of subjects'

solutions in a majority of cases, whereas convex-hull and nearest-neighbor solutions fell within that range, with one exception in the latter case. Compared with the mean subject path length, convex-hull and nearest-neighbor solutions appear to provide reasonable first approximations, with the former tending to undershoot the mean and the latter to overshoot.

Figure 1.

Maximum and mean subject path lengths and convex-hull, nearest-neighbor, and crossing-avoidance path lengths for 12-point problems.



Although path length captures an important aspect of a solution, paths of similar lengths may nevertheless have quite different configurations. For example, although the majority of nearest-neighbor solutions fell within the range of subjects' solutions, 75% of them had crossings, compared with 2% of subjects' solutions. (The convex-hull and crossing-avoidance heuristics had 0 crossings.)

Because path length reflects only one characteristic of a solution, we examined the similarity of paths in an alternative way. For these relatively simple problems, subjects often produced identical solutions, allowing us the expedient of comparing the heuristics' solutions with the modal human solution for each problem. To do so, we counted the overlap between solutions in terms of the number of identical arcs. That is, for each arc i,j on the modal subject path, if a heuristic connected the arc i,j or j,i , overlap was incremented by 1. The maximum possible overlap between two identical paths is 11. The expected overlap, based on 10,000 randomly generated paths, is 2, with a standard deviation of 1.29, making overlap levels of 5 or more in the present case significantly greater than chance ($p < .05$). The results of this comparison are summarized in [Table 1](#), which shows the total overlap between each of the heuristic solutions and the corresponding modal subject path for the twelve problems. [Figure 2](#) illustrates the performance of the three heuristics compared with participants' modal performance on Problem 8.


Table 1

Overlap Between Subject Modal Paths and Convex-Hull, Nearest-Neighbor and Crossing-Avoidance Solutions to 12 Open Problems.

Problem Number	Percentage of subjects with modal solution	Convex-Hull	Nearest-Neighbor	Crossing-avoidance
1	36%	11	11	6
2	68%	11	9	6
3	44%	6	8	5
4	32%	11	9	4
5	43%	11	9	4
6	29%	11	7	6
7	25%	7	6	5
8	20%	8	6	3
9	38%	6	5	4
10	38%	8	6	9
11	32%	7	7	8
12	30%	7	7	4

Figure 2.

Participants' modal path for Problem 8 and paths for the convex-hull, crossing-avoidance, and nearest-neighbor heuristics (clockwise from top left). Dotted lines show deviations from modal path.



The convex-hull solution exactly matched the modal subject path on 5 of the 12 problems and the nearest neighbor on 1, whereas none of the crossing-avoidance solutions were exact matches. In all cases, the overlap with the modal solution for the convex-hull and nearest-neighbor solutions was significantly greater than chance, whereas for the crossing-avoidance heuristic, overlap was at chance levels in 5 of the 12

instances. Finally, the mean overlap for the convex-hull heuristic (8.67), was significantly greater than for the nearest neighbor (7.5), $t(11) = 2.65, p < .05$, which in turn was significantly greater than for crossing-avoidance (5.33), $t(11) = 3.03, p < .01$.


Comparison 2

We repeated the test using the data from the second experiment of Chronicle et al. (2006). The stimuli were six randomly generated 15-node problems. Twenty-four participants solved each problem, half starting from the leftmost node, half from the right. Because there were relatively few identical subject paths in this case, we limited the test to a comparison of path lengths only. The results are summarized in Figure 3, which shows the maximum and mean subject path lengths and path lengths for the three heuristics. The results are shown for each of the six problems where subjects started from the left, followed by the same six problems starting from the right.

Overall, the convex-hull heuristic fell beyond the range of subjects' paths in 1 of the 12 cases, the nearest-neighbor in 8, and the crossing-avoidance in 11. The incidence of crossings was again high in the nearest-neighbor solutions at 58% of solutions, compared with 3% for subjects and 0 for the other two heuristics.

Figure 3.

Maximum and mean subject path lengths and convex-hull, nearest-neighbor, and crossing-avoidance path lengths for 15-point problems.



Discussion


The article explored the potential of three heuristics to account for human performance on open versions of TSPs. Of the three, the convex-hull heuristic provided the best fit to human solutions. In all but 1 of 24 comparisons, it generated path lengths that fell within the range of human solutions and produced no crossed arcs, and in 12 comparisons it

had a significantly greater overlap with the subjects' modal solution than the other procedures. The nearest-neighbor heuristic generated paths that fell beyond the range of subject solutions in 9 of 24 comparisons and produced many more crossings than subjects (appearing in more than 50% of solutions, compared with 3% of subjects' solutions). The crossing-avoidance heuristic performed most poorly of all, having the lowest overlap with subjects' modal solutions and path lengths that fell beyond the range of subjects' solutions in a majority of the 24 comparisons. The comparative success of the convex-hull approach is important because it extends its range of application and raises the possibility that it may further extend to other spatial problems.

Although these preliminary results tend to support the convex-hull approach, we did see evidence supporting the other two. For example, some subject solutions were identical to nearest-neighbor solutions, and although none exactly matched a crossing-avoidance solution, some bore a strong family resemblance. These trends are nicely illustrated with our more recent, unpublished data for a 20-point stimulus. In this case, 7 of 30 human solutions were identical to the convex-hull solution, and 1 exactly matched the nearest-neighbor solution (including a crossing), and although there were no exact matches with a crossing-avoidance solution, a family resemblance is apparent in Figure 4. The results suggest the possibility that each of the heuristics may apply to some extent, and the question arises how they might be adapted to better represent what people do? We consider this question next.

Figure 4.

Crossing-avoidance solution (left) and subject's solution (right) to a 20-point problem.




The heuristics might be made more realistic in a number of ways. Here we consider two, "lookahead" and just noticeable difference (jnd). Lookahead describes the capacity of people to consider future moves in a problem-solving context. In the present case, the heuristics were limited to the upcoming level of the decision tree only, whereas people may analyze problems to a greater depth. The jnd concept refers, of course, to how large a physical difference is required to register as a reliable perceptual difference. In the present case, the heuristics operated under perfect information about physical

distances and coordinate positions, whereas people will be insensitive to small variations in these factors.

The heuristics could be provided with greater lookahead in several ways. One would be to allow the heuristic to consider moves that incorporate more than one point at a time, analogous to increasing processing capacity. Another would be to define a moving spatial area within which any points have a probability of being considered for the path, analogous to focal attention (Pizlo, Stefanov, Saalweachter, Haxhimusa, & Kropatsch, 2005). Allowing a heuristic to “break the rules” within these limits could avoid some decisions that are uncharacteristic of people. These two approaches, increasing capacity or focal attention, can have similar or different effects, depending on context. For example, when faced with a path segment consisting of the three points, A, B, and C, shown in the left panel of Figure 5, the current crossing-avoidance heuristic would connect them as ACB, which is substantially longer than the more human-like ABC, shown in the next panel. If the heuristic were able to look two points ahead on the x-axis and choose how best to connect the two to the current point, then it would avoid this error. A similar effect would result from providing an area of visual attention of sufficient diameter. However, if the next and final point were in position D, then the heuristic would be committed to ABCD, incurring a crossing, as shown in the next panel. If the heuristic began with a lookahead of three points on the x-axis or a wide enough area of focal

Figure 5.

How line segments might be connected under different assumptions about lookahead or focal attention.




attention, it might connect the path as ADBC and avoid the crossing. However, increasing lookahead or focal attention may not necessarily make heuristics produce more human-like solutions. For example, one of the weaknesses of the nearest-neighbor heuristic is that decisions made early in the path may create crossings much later in the process, involving spatially distant points so that increasing either lookahead capacity or focal attention by any reasonable amount is not likely to provide a remedy.

Incorporating a finite jnd into a heuristic's judgments is another mechanism that could lead to more human-like paths. For example, in the convex-hull heuristic, the imaginary lines that represent the boundary of the hull have no thickness. Allowing them a finite thickness and defining any point that touches the hull to be on it would prevent the heuristic from making some decisions that people rarely make. (A similar effect could be achieved by providing the points with a finite radius.) The present heuristic creates the path on the left in Figure 6, while under the implementation of a jnd, the path on the right would also become possible because the interior point would then be considered to lie on the convex hull for the first decision. The path on the right is more typical of what people do than the path on the left. In addition, if either the jnd or lookahead mechanisms were implemented with a probabilistic element, then the heuristics would be capable of producing a range of solutions, rather than just one.

Figure 6.

Solutions produced by the convex-hull heuristic under perfect information (left) and with a finite jnd (right).



In closing, although the three heuristics have been contrasted here, there is no reason that they could not be used in combination. Indeed, they already share common features. Both the convex-hull and crossing-avoidance heuristics, for example, seek nearest neighbors under certain constraints. (The former seeks nearest neighbors on the convex hull, whereas the latter seeks nearest neighbors on the x-axis.) The convex-hull heuristic, although not designed to do so, may have the effect of avoiding crossings.

References

- Chronicle, E. P., MacGregor, J. N., Ormerod, T. C., & Burr, A. (2006) It looks easy! Heuristics for combinatorial optimisation problems. *Quarterly Journal of Experimental Psychology*, *59*, 783–800.
- Golden, B., Bodin, L., Doyle, T., and Stewart, W. (1980). Approximate Travelling Salesman Algorithms. *Operations Research*, *28*, 694–711.

- Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The Travelling Salesman Problem: A hierarchical model. *Memory & Cognition*, *28*, 1191–1204.
- MacGregor, J. N., Chronicle, E. P., & Ormerod, T. C. (2004). Convex hull or crossing avoidance? Solution heuristics in the Travelling Salesperson Problem. *Memory and Cognition*, *32*, 260–270.
- 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 human performance on the Travelling Salesman Problem. *Perception*, *28*, 1417–1427.
- MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2000). A model of human performance on the Travelling Salesperson Problem. *Memory and Cognition*, *28*, 1183–1190.
- Ormerod, T. C., & Chronicle, E. P. (1999). Global perceptual processing in problem-solving: The case of the Travelling Salesperson. *Perception and Psychophysics*, *61*, 1227–1238.
- Pizlo, Z., & Li, Z. (2005) Solving combinatorial problems: 15-puzzle. *Memory and Cognition* (in press).
- Pizlo, Z., Stefanov, E., Saalweachter, J., Haxhimusa, Y., & Kropatsch, W. G. (2005). *Adaptive pyramid model for the Traveling Salesman Problem*. Paper presented at the Human Problem Solving Difficult Optimization Problems Workshop, Purdue University, June 2005.
- van Rooij, I., Stege, U., & Schactman, A. (2003). Convex hull and tour crossings in the Euclidean Travelling Salesperson Problem: Implications for human performance studies. *Memory and 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., Lee, M. D., Dry, M., & Hughes, P. (2003). The roles of the convex hull and number of intersections in performance on visually presented Travelling Salesperson Problems. *Memory & Cognition*, *31*, 1094–1104.
- Vickers, D., Mayo, T., Heitmann, M., Lee, M. D., & Hughes, P. (2004). Intelligence and individual differences in performance on three types of visually presented optimization problems. *Personality and Individual Differences*, *36*, 1059–1071.

Acknowledgment

The research was partially supported by a research grant to MacGregor from the National Science and Engineering Research Council of Canada