

Human Navigational Heuristics in Solving the Euclidean Travelling Salesman Problem

Tamas Makany

University of Southampton, UK

tamas.makany@soton.ac.uk

and

Michael Makowsky

George Mason University, USA

mmakowsk@gmu.edu

Abstract

Humans outperform, within limits, most algorithmic solutions in spatial optimization tasks, such as in the Euclidean Travelling Salesman Problem (TSP). It has been recently suggested that people follow simple, perceptual heuristics – like preferring closed shapes or following the boundaries – rather than applying computationally expensive algorithms. In the present paper, human navigational heuristics were simulated in standard 10-nodes TSPs with different levels of difficulty. Within an agent-based model, an artificial travelling agent evaluates a disutility function at each step with respect to the distance to potential object, memory of previous routes taken, and the distance from home. The number of steps in each tour was recorded and solutions were analyzed. We conclude that human navigational heuristics can lead to near-optimal solutions in tasks like the TSP, and that solutions modelled herein are both cognitively feasible and efficiently robust to a variety of circumstance.

Keywords: agent-based modelling, disutility function, heuristics, human performance, spatial navigation, optimization, Travelling Salesman Problem

Acknowledgement

The authors wish to express their gratitude for the stimulating discussions with the members and lecturers of the Complex Systems Summer School 2006 in Santa Fe, NM at the St. John's College and at the Santa Fe Institute. We are especially grateful for Lukas Kroc and Spandan Tiwari for the long conversations on computational optimization. We also very much appreciated the expert comments of Douglas Hoskins.

Introduction

The Euclidean Travelling Salesman Problem (TSP) is a well-known and studied combinatorial optimization problem in many sciences, including mathematics (Golden, Bodin, Doyle, & Stewart, 1980), geometry (Mitchell, 2004), geography (Norback & Love, 1977), and psychology (Chronicle et al., 1996). In brief, the classical TSP is looking for the shortest spatial route for a salesman, who has to visit a number of cities to sell his product and return to his hometown. In a more formal sense, TSP is looking for the optimal algorithm to find the shortest tour that visits every point of a set S of n points. There is a substantial amount of research on the computational aspects of the problem (and the related special cases) that can provide a relatively clear picture of the NP-hardness of such mathematical problem (Mitchell, 2004).

There are many standard approximation algorithms known for TSPs (Golden et al., 1980; Mitchell, 2004). The most common ones include the *Nearest Neighbour (NN)* and the *Convex Hull with the Cheapest-Insertion (CHCI)* procedures. The *NN algorithm* connects an arbitrary chosen leading node to the closest unconnected node within a set of n nodes (where $n > 2$), and make that node as the new leading one. This sequential linking procedure continues until $n-1$ connections are made. After the last node was made to the leading node it is connected to the starting node to complete the TSP tour. In the worst case, the NN can be arbitrarily factor larger than the optimal tour. The *CHCI method* searches among the smallest convex set containing all nodes not inserted so far, and choose a node whose insertion causes the lowest increase in the length of the tour. There are more sophisticated and complex procedures for solving the TSPs that involves a combination of simpler algorithms (e.g., *1.5-opt*, *2-opt* and *3-opt heuristics*), or the

recently introduced best known factor, the *polynomial-time approximation* (for more details, see Arora, 1998).

Natural selection of human and animal navigation in the course of time has evolved highly efficient non-algorithmic methods of spatial translocation. Avian migration across continents, rodent way finding in long underground ducts, or even human navigation in huge metropolises like London provide ground for observing the various examples of spatial optimizations taking place in nature. Many of them still lack a formal scientific description or algorithmic approximation. To tackle the complex, dynamically changing environmental circumstances, information shortage, time-pressure and most of all the limited cognitive processing capacity, living organisms developed good, quick and relatively effortless ways of dealing with spatial information, called *navigational heuristics*. The main difference between these cognitive heuristics and mathematical approximations (also referred to as ‘heuristics’) is that neither birds nor university professors are able to perform algorithmic calculations with such complexity at the very moment when their navigational decisions take place.

However, cognitive heuristics and relating strategies – like following the boundary of an area or preferring routes with closed global shapes – might provide better solutions than simple computational algorithms in optimization tasks (Chronicle et al., 2006). A study by MacGregor, Ormerod, and Chronicle (2000) attempted to model human performance on the TSP. They proposed a mixture of heuristics: (i) using a convex hull with the cheapest insertion criterion method to link boundary points, (ii) avoiding line crossings, and (iii) avoiding “indentations”. Although their model performed near the optimum on a wide range of solutions, it did not provide an explanation of why and how humans would apply these set of rules. In fact, these authors only collected the most successful strategies rather than finding their underlying controlling factors (for further criticisms on this model see Vickers, Lee, Dry, & Hughes, 2003).

Previous studies in navigation found two main initial *control parameters* that determined subsequent spatial behaviour: (i) energy cost of exploration, and (ii) cognitive cost of memorizing routes (Makany, Redhead, & Dror, 2006). Individual settings of these parameters resulted in navigation heuristics, like ‘minimizing travelling distance’ or

‘preferring novel to well-known routes’. As people can set their own heuristics, the parameter range of optimal performance might vary largely across individuals, tasks and environments (these are called *order parameters*).

Based on the above review of the different approaches of solving TSPs, we propose that humans apply *navigational heuristics* in spatial optimization tasks. In the next part of the paper, we will introduce an agent-based model that utilizes the two main navigational heuristics – energy and memory – described by Makany et al. (2006), and a third one that deals with maintaining relative proximity with the home node.

The Model¹

We implemented an Agent-Based Model of Human Navigational Heuristics (ABM) within the freeware NetLogo (Wilensky, 1999) modelling environment. Three standard, 10-nodes TSP layouts were used with increasing levels of difficulty (Easy, Medium, Hard) based on the number of internal nodes (1, 4, and 5, respectively) that fall within the outermost node boundary (see the TSP coordinates in Appendix II). One artificial agent at a time explores and navigates through each two dimensional lattice layout. The lattice is 280 pixels width and 180 pixels height with the location of origin in the bottom left corner. The task of the agent is to complete a tour of visiting all 10 nodes and return to the arbitrarily set starting node (Object0).

As preparation for each step of the tour, the agent explores her environment and determines a preferred next object to travel to. This individual preference is represented – for modelling purposes only – in the calculation of a *disutility function* (1). However, it should be noted that it is not assumed that any human navigator would perform the same algorithmic calculation. It is only an approximation of incorporates the navigational heuristics and allows for relative weightings.

¹ A java applet of the model is available via the web at

<http://www.soton.ac.uk/~tm304/sfi/tsp/>

$$(1) f(d) = e_{i,j}^{\alpha} + m^{\beta} - e_{0,j}^{\gamma}$$

The disutility function (Equation 1) one serves as the object function of the agent. Before each step the agent evaluates the function with regards to each object, j , not yet visited, and chooses that which presents the minimal disutility relative to her current object, i . The arguments of the function are as follows: distance to the object ($e_{i,j}$), sum of steps already taken on the interceding patches (m), and the distance from the “home” object² ($e_{0,j}$). Each argument is associated with a parameter weighting: α , β , and γ , respectively.

The mechanism that the agent uses to walk around is highly sensitive. She faces the next selected object first, takes a step “directly” towards it, then she rounds off her coordinate to be in the middle of her current step. As such an agent may have a path that is a generic shape (e.g., a convex hull) but at the micro level of the steps is slightly different, with turns that are smoother or rougher. This must be kept in mind when decisions, such as whether to go clockwise or counter-clockwise result in small differences in total steps, despite nearly identical tour shapes.

The agent operating here is intentionally myopic – evaluations of the objective function are made solely with care to the next object. There is no strategic evaluation of the “big picture.” This lack of foresight, combined with heuristic evaluations of distance and repeated steps that are computationally simple, the resulting total construct is one which fits comfortably within the assortment of constraints that fall under the heading of bounded rationality.

Results

First, we compared the best performance of our Model on each difficulty level with the known optimum and with average human path lengths (see Appendix I). In addition, we also included the average solutions for two common algorithms (Nearest

² The “home” object is the arbitrarily chose first object which the agent starts from and seeks to finish at.

Neighbour, NN; and Convex Hull with Cheapest-Insertion Procedure, CHCI) in the comparison. See Figure 1 for a graphical illustration and for the data table. It was found that our Model performed within the range of the other approaches with finding a near-optimal path. In addition, the harder the task became the better our Model performed. Interestingly, in the case of the 5-internal node (Hard) problem we have found a better than optimal result. As we did not intend to improve the known optimum for this special TSP, further analysis is needed to validate this finding.

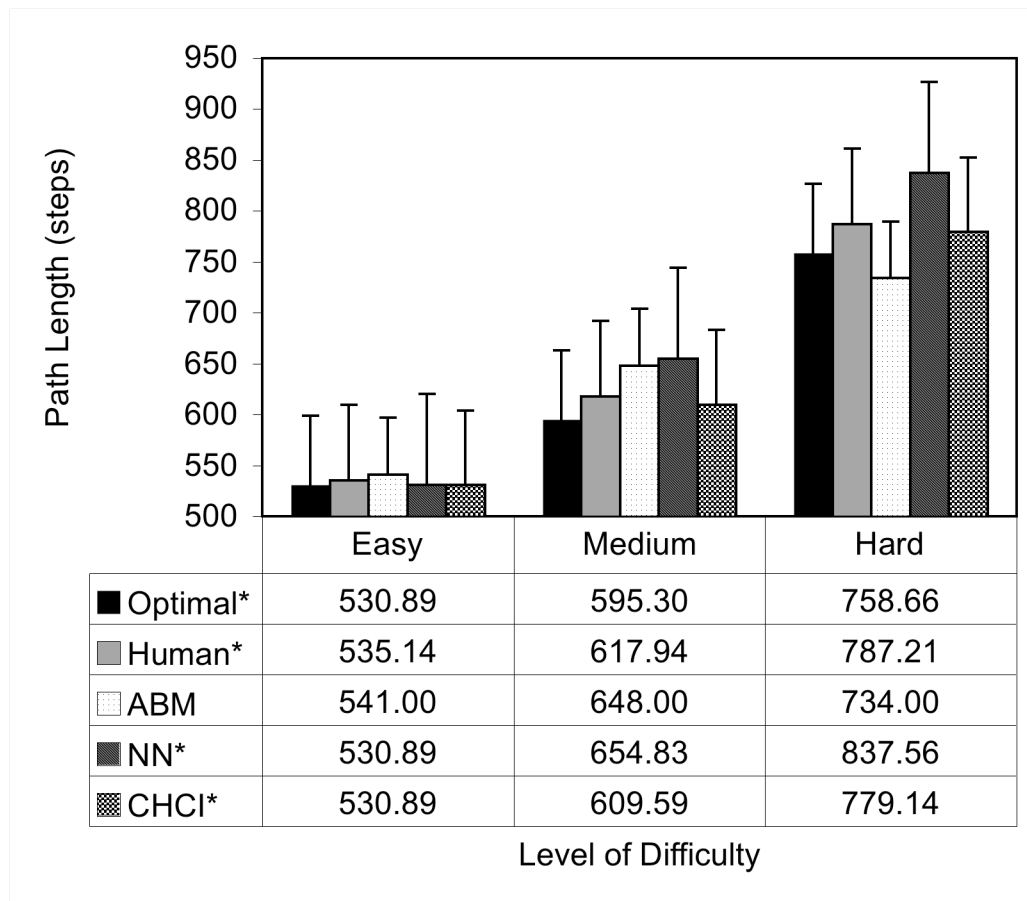


Figure 1. Comparison of TSP solutions with different approaches. Data from our Model (ABM) is compared with previously known optimum path lengths (Optimal), average human performance (Human), the Nearest Neighbour (NN), and the Convex Hull with the Cheapest-Insertion (CHCI) algorithms. Marked (*) datapoints are taken from the study by MacGregor and Ormerod (1996).

In the next step, we fit our model to the human performance on each level of difficulty. Those combinations of the parameters (alpha, beta, gamma) were selected that

scored the closest to the average human performance (Figure 2). The correlation coefficient between the fitted model and the human data was [$r = .991$]. It shows a very good (99%) fit between the two samples and means that our model with these fitted parameters performs equally as the human population. The fitting procedure was performed in order to be able to infer conclusions on how humans might solve different TSPs.

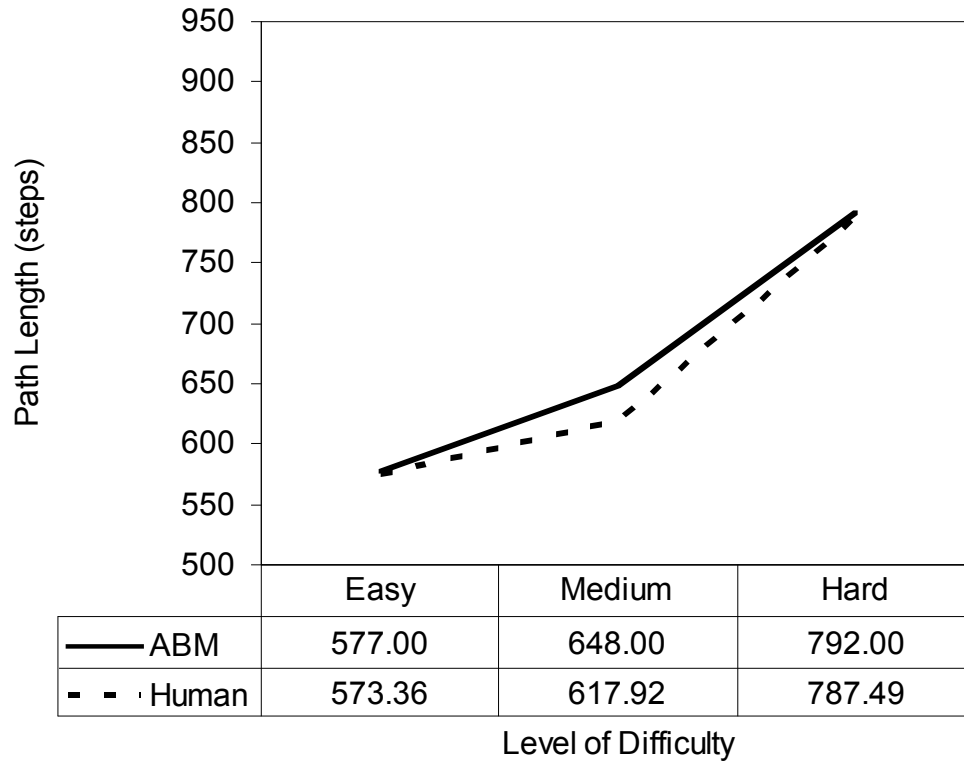


Figure 2. The graph shows the best fit of our Model (solid line) with the average human performance (Human; dashed line). After the fit, the two datasets did not differ from each other statistically.

Finally, we plotted the parameter range of the fitted Model that simulated human performance (Figure 3). The parameters are highly varied across the three TSPs, suggesting that human solutions may not be based on a simple algorithm. An emerging pattern could be observed, when the distribution of these parameters was analysed. A Spearman rank correlation revealed that *alpha* and *gamma* were positively associated with each other in all three cases, meaning that a change in alpha was followed by a

similar change in gamma and vice versa [$r_{\text{easy}} = .973$; $r_{\text{medium}} = .564$; $r_{\text{hard}} = 1.0$; with all $p < .001$]. *Beta*, on the other hand, could take any value from its range [0.1 to 1.0] without influencing the optimality of the solutions.

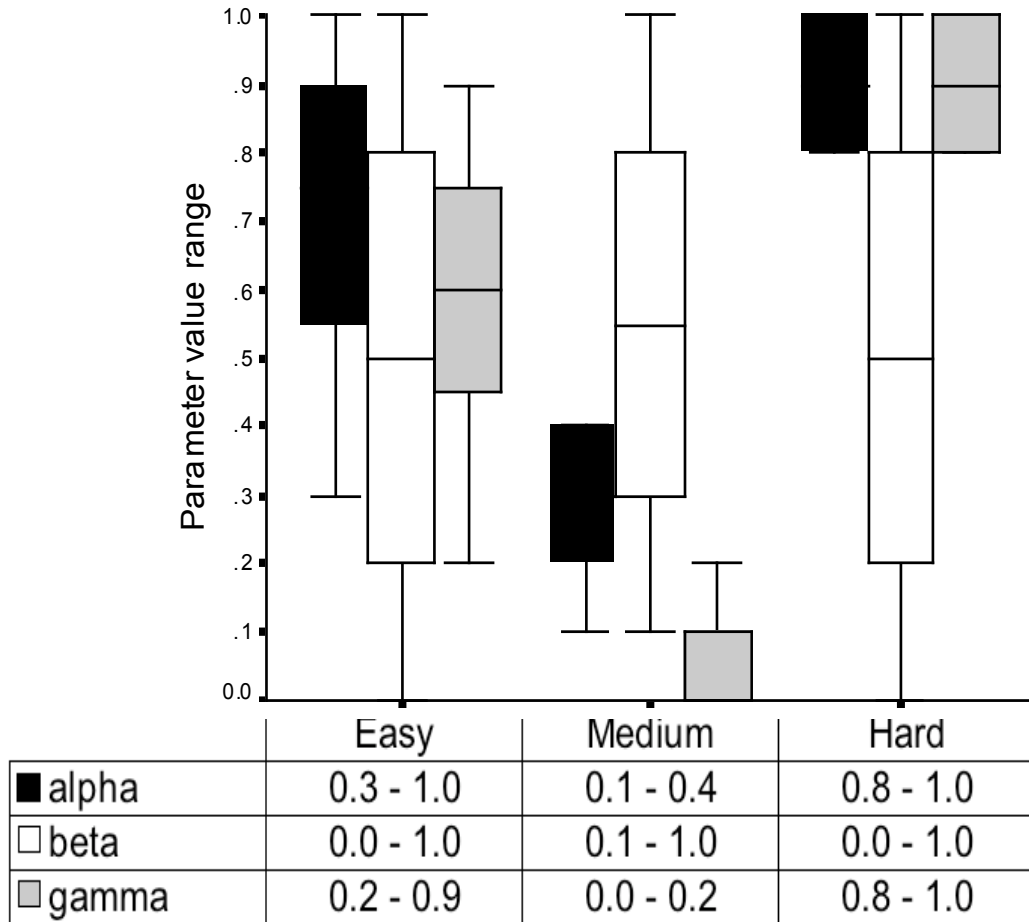


Figure 3. The graph shows the range of the parameters (α , β , γ) that simulated human performance in the three TSP tasks. While α and γ are positively correlated to each other and dependent of the difficulty of the task, β seems to behave as an independent factor.

Discussion

This paper presented an agent-based simulation of human navigation heuristics in three standard 10-nodes Euclidean Travelling Salesman Problem (TSPs). The simple heuristics included minimising the travelled distance to the next object (α), preference to previously taken routes (β), and maintaining proximity towards home

(γ). We found that by applying certain combinations of these heuristics, near-optimal solutions could be found. In fact, heuristic models work efficiently even in difficult conditions, where otherwise considerable computational power is needed for most algorithms to find the optimal tour (Chronicle et al., 2006; MacGregor et al., 2000). Based on our results, we argue that navigational heuristics could lead to optimal solutions in spatial optimization tasks without extensive algorithmic calculations.

In addition, we investigated a specific range of parameter settings that resulted in a good approximation of the average human performance reported by MacGregor and Ormerod (1996). Human performance could be generically reproduced with a broad range of parameter values. This relative insensitivity towards the specificity of settings has two kinds of implications:

First, it shows that that our model is robust against multiple designs, and also powerful enough to perform near the optimum. For example, we could reproduce highly efficient *convex hull* solutions with multiple parameter combinations.

Second, the insensitivity of parameter settings implies that humans might apply various navigation strategies when solving optimization tasks. Task difficulty – the number of internal nodes in an TSP set – seems to influence navigational decisions based on metric distances only (*alpha & gamma*), but it is independent of topological and geometrical route memory requirements (*beta*). This latter finding may also contribute to further understandings of how computationally and cognitively limited humans could – in some cases – outperform sophisticated algorithms with an immense memory capacity.

The sufficient efficiency, while minimising on cognitive effort and data collection, lends the model to inclusion in the broad spectrum of behaviours often classified as satisficing behaviours (Simon, 1957). It satisfies in that it provides near-optimal solutions under a variety of circumstances, avoiding worst-case scenarios familiar to the rules versus discretion debate in political economy and data analysis. This worst-case scenario avoidance bears particular attention when considered in light of the growing literature on loss aversion, and the long standing notion of generic risk-aversion (Kahneman, Slovic, & Tversky, 1982).

Conclusions

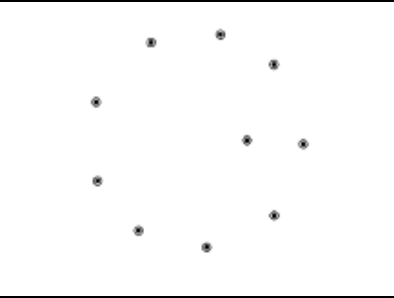
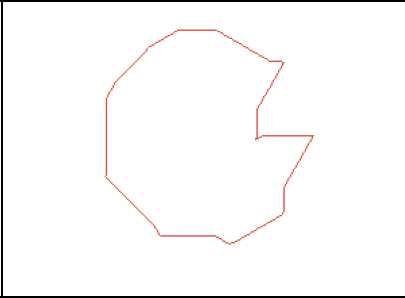
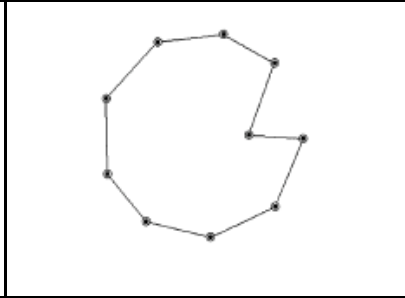
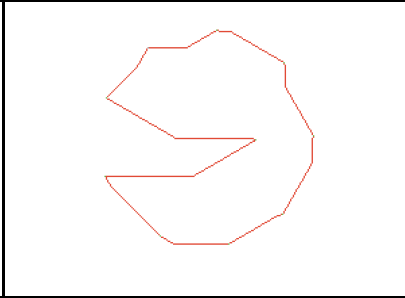
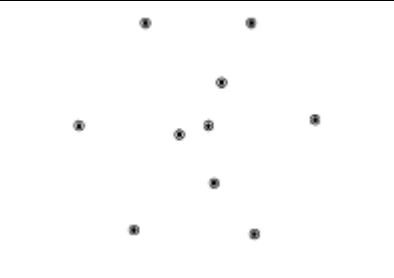
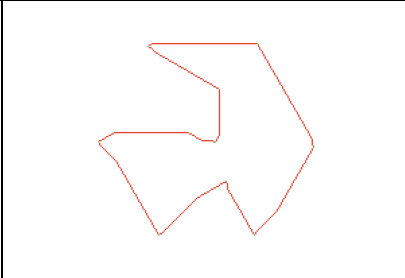
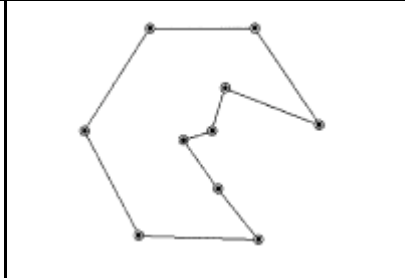
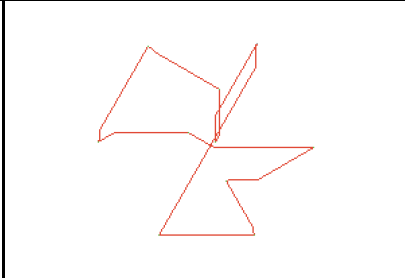
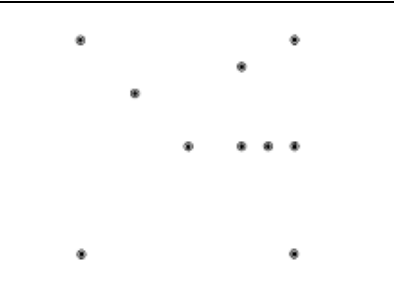
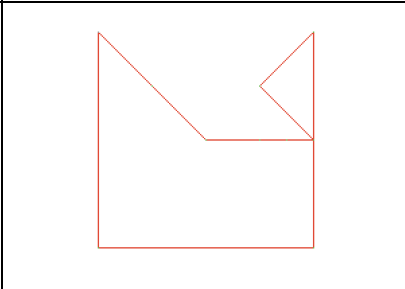
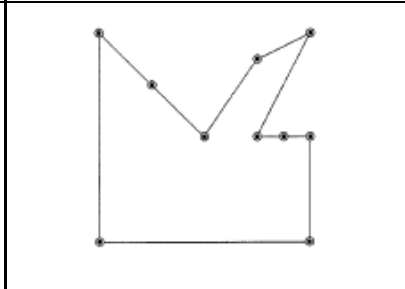
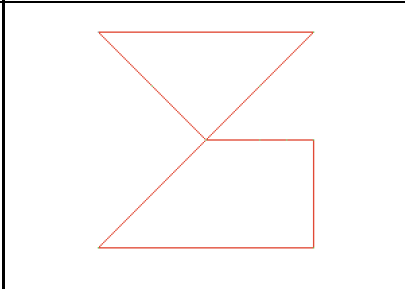
Although this paper represents only a preliminary study it seems apparent that investigating human heuristics and other similar decision-making techniques could be equally fruitful for a variety of interdisciplinary researchers including computational optimization experts, cognitive psychologists, mathematical engineers, urban planners, or behavioural economists. As a next step for this research, a more systematic analysis needs to be performed to explore other, more difficult and complex TSPs. In addition, the exact function of the invariant beta parameter (memory heuristic) requires further investigations. In the future, the human navigation based heuristic approach could be integrated into more generic optimization models (e.g., evolutionary algorithms) to improve their efficiency.

References

- Arora, S. (1998). Polynomial time approximation schemes for Euclidean travelling salesman and other geometric problems. *J. Assoc. Comput. Mach.*, 45, 753-782.
- Chronicle, E.P., MacGregor, J.N., Ormerod, T.C., & Burr, A. (2006). It looks easy! Heuristics for combinatorial optimization problems. *The Quarterly Journal of Experimental Psychology*, 59, 783-800.
- Golden, B., Bodin, L., Doyle, T., & Stewart, W. (1980). Approximate travelling salesman algorithms. *Operations Research*, 28, 694-711.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: University Press.
- 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. (2000). A model of human performance on the travelling salesperson problem. *Memory & Cognition*, 28, 1183-1190.
- Makany, T., Redhead, E.S., & Dror, I.E. (2006). Spatial exploration patterns determine navigation efficiency: Trade-off between memory demands and distance travelled. *Submitted for publication*.

- Mitchell, J.S.B. (2004). Shortest paths and networks. In J.E. Goodman & J. O'Rourke (Eds.), *Handbook of discrete and computational geometry* (pp. 607-641). New York, NY: CRC Press.
- Norback, J.P., & Love, R.F. (1977). Geometric approaches to solving the travelling salesman problem. *Management Science*, *23*, 1208-1223.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*, 99-118.
- 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.
- Wilensky, U. (1999). NetLogo. <http://ccl.northwestern.edu/netlogo/> Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Appendix I. Standard Euclidean Travelling Salesman Problems (TSPs) with three different levels of difficulty (Easy, Medium, Hard) based on the number of internal nodes (1, 4, 5, respectively) that fall within the outermost boundary. Tours of our agent-based model are shown in the ABM column.

	TSP	ABM	Human*	Non-optimal Solution
Easy 1-internal nodes				
Medium 4-internal nodes				
Hard 5-internal nodes				

*Human solutions are taken from Ormerod & Chronicle (1999).

Appendix II. X-, Y- Coordinates for the three TSPs, based on the coordinates by MacGregor and Ormerod (1996).

EASY	X-	Y-
Object0	220	92
Object1	198	146
Object2	148	170
Object3	97	157
Object4	66	120
Object5	65	62
Object6	106	18
Object7	157	12
Object8	197	34
Object9	177	89

MEDIUM	X-	Y-
Object0	220	83
Object1	178	160
Object2	97	158
Object3	60	87
Object4	105	18
Object5	176	18
Object6	155	58
Object7	150	126
Object8	147	87
Object9	127	94

HARD	X-	Y-
Object0	220	170
Object1	60	170
Object2	60	10
Object3	220	10
Object4	220	90
Object5	180	130
Object6	100	130
Object7	140	90
Object8	180	90
Object9	200	90