Minimal Paths in the City Block: Human Performance on Euclidean and Non-Euclidean Travelling Salesperson Problems

Amy L. Walwyn

School of Psychology

University of Adelaide

Daniel J. Navarro

School of Psychology

University of Adelaide

Abstract

An experiment is reported comparing human performance on two kinds of visually-presented travelling salesperson problems (TSPs), those reliant on a Euclidean geometry and those reliant on a city-block geometry. Across multiple array sizes, human performance was near-optimal in both geometries, with the magnitude of deviations from optimality being predicted accurately by the length of the shortest solution. Moreover, human solutions appeared to be adapted to the geometry of the problem: for arrays to which the minimal tour differs depending the geometry, human solutions also differ in the appropriate manner. These results suggest that human near-optimality for visual-format TSPs may be tied to the familiarity of the geometry, rather than being inherently constrained to an innately Euclidean perceptual representation.

Minimal Paths in the City Block: Human Performance on Euclidean and Non-Euclidean Travelling Salesperson Problems

Introduction

The idea that the human mind embodies a simplicity or minimality principle has been proposed by many authors in many domains (e.g., Köhler 1929, Zipf 1949, Chater & Vitányi 2003, Feldman 2003). One area in which this proposition has been examined closely is with respect to visually-defined combinatorial optimization problems (e.g., Vickers, Mayo, Heitmann, Lee & Hughes 2004), with particular reference to the *travelling salesperson problem* (e.g. Polivanova, 1974, MacGregor & Ormerod 1996, Vickers, Butavicius, Lee & Medvedev, 2001, Vickers, Bovet, Lee & Hughes 2003, Dry, Lee, Vickers & Hughes, 2006, Pizlo et al. 2006, Chronicle, MacGregor, Lee, Ormerod & Hughes, 2008). In the standard closed-tour Euclidean form of the problem, participants are shown an array of n points, and asked to construct a path that passes through all n points exactly once and returns to the starting point. One of the main reasons for studying the TSP is that it ties directly to the minimality principle, and poses a very complex computational problem. An n point problem has $(n-1)!/2$ possible solutions, and finding the globally minimal solution is known to be a nondeterministic polynomial (NP) time problem, making it extremely difficult. Despite this, the general pattern of results across many studies is that human solutions to visually-presented closely approximates the globally optimal solution.

Not surprisingly, the impressive performance on the TSP is tied to the visual representation of the problem. When presented in visual form, people solve the TSP near-optimally, but since the inception of work on this topic (Polivanova 1974) it has been clear that people perform much worse on equivalent *verbal* representations of the minimization problem. Though intuitively obvious, this result poses a theoretical puzzle. If minimality is a fundamental organizing principle for higher-order cognition (e.g., Feldman 2003) as much as for visual perception, why should linguistic/symbolic representation of the problem impair performance so badly? Presumably, the answer is that the visual system is attuned to the two-dimensional Euclidean geometric structure of the visual TSP, whereas no such match exists between symbolic conceptual representations and the verbal TSP. Accordingly, the critical aspect to human performance is the degree to which the relevant representational system is matched to the specifics of the problem.

Viewed from this perspective, while it seems clear that the visual system exploits the *geometric* structure of the task, how heavily do we rely on the fact that the problem is *Euclidean*? A case can be made that Euclidean geometry should be privileged. In our daily experience, we interact with a (locally) Euclidean geometry, so it is natural to expect the visual system to be adapted to this structure (e.g., Shepard 1994), via representations that are innately specified. By way of contrast, it would be highly surprising if people were optimal with respect to problems formulated in some arbitrary Riemannian geometry: nothing in our evolutionary or everyday experiences prepares us for such problems. However, in between these two extremes lie a range of interesting intermediate cases. For instance, efficient spatial navigation when driving a car requires people to construct minimal paths within a geometric structure (i.e., the road network) that can be highly non-Euclidean. In the simplest case, many cities have streets laid out on a grid, implying that the distances between two points are best described by a city-block metric rather than a Euclidean one.[1] Given the familiarity of the city block metric, and its amenability to simple verbal description (i.e., participants are told that paths must be constructed using horizontal and vertical lines only), it is a natural candidate for exploring human performance on TSP problems specified in non-Euclidean geometries. If human optimality on such problems is closely tied to innately specified perceptual representations, then one would expect a sharp decrement in performance even for familiar non-Euclidean problems such as the city block TSP. However, to the extent that prior learning allows for the construction and use of novel, richer mental representations (or alternatively, allows people to adapt existing representations to solve different problems), we might expect that human performance is more robust. The only work of which we are aware that discusses such problems is by Saalweachter and Pizlo (2008), who considered TSP problems involving obstacles or mazes, an found that people are very close to optimal if the obstacles do not overly distort the Euclidean structure, but deviate from optimality in more complex cases. Along similar lines, this paper reports the results of an experiment comparing Euclidean TSP performance to city block TSP performance.

Experiment

Method

Participants. Forty people aged 17-52 participated in the study (16 males, 24 females), recruited from the general community and from the University of Adelaide undergraduate research participation pool. Students received course credit for participating, and other participants received a $10 gift voucher.

Materials. Stimuli consisted of twelve 10-point random dot arrays, and six 40-point arrays, presented in pencil-and-paper form. Each point was sampled from a uniform distribution over a 158mm $\times$ 158mm square, subject to the constraint that the points be visually distinct.[2] Minimal tours for all problems were estimated in both city block and Euclidean geometry using standard numerical optimization methods.

Procedure. Each participant completed six visual TSPs in Euclidean form and another six in city block form. In both cases, three of the six items were 10 point problems and the other three were 40 point problems. Stimuli were grouped by number of points and by geometric format, but in all other respects presentation order was counterbalanced. For any given problem, participants were asked to create a path that visited all points and returned to the start point, in as short a distance as possible. They were free to begin the tour at a point of their choice. In the Euclidean condition people were asked to connect points using direct lines; in the city block format, however, they were told that they could only use horizontal and vertical lines. Participants were shown a sample problem before the real ones.

Exclusions. Four participants produced extremely long paths to at least one of the problems, relative to the other 36 participants. This happened for both Euclidean and city block problems. For the sake of conservativeness, all data from those four people were excluded. Additionally, a number of solutions were lost due to transcription errors or incomplete tours. In total, this left 215 cases (107 city block, 108 Euclidean) for the 10-point problems and 199 cases for the 40-point problems (93 city block, 106 Euclidean).[3]

Results

The lengths of all 414 solutions produced by participants are shown in the left panel of Figure 1, plotted a function of the length of the minimal solution to the corresponding problem. The right hand side plots the average deviation from the optimum for each TSP, which varies smoothly as a function of the length of the best solution. The Euclidean and city block solutions sit on the same curve, which suggests that human solutions are near-optimal in both cases. However, it might be that people select tours relying solely on the Euclidean metric, and it just so happens that the optimal city block solutions are sufficiently similar that this still leads to good solutions. This is not implausible on its face: optimal solution lengths in the two metrics are highly correlated in both the 10-point and 40-point problems, at $r = 0.99$ ($p = 2.3 \times 10^{-9}$) and $r = 0.96$ ($p = .003$) respectively. Moreover, the optimal tours do tend to agree with one another across metrics. In the

10-point problems in fully 7 of 12 cases the tours were identical, so in those cases people could safely disregard the cover story and produce optimal city block solutions by trying to solve the corresponding Euclidean problem. In the other five cases, the solutions shared 4-8 of the 10 links. For the 40-point problems this effect is much weaker. In no such case did the an array have identical minimal tours across metrics, and the solutions agreed on 25-35 edges. This is illustrated in Figure 2.

The fact that the optimal solutions differed in some cases suggests a natural way to determine if people did in fact tailor their solution to the appropriate metric. For any given solution, we construct a measure of its "relative Euclidean-ness", namely the number of links it shares with the optimal Euclidean solution, minus the number of links it shares with the optimal city block solution to the same array. For an array in which the best solutions differ on k edges, this measure varies from $-k$ to k, and hence weights more diagnostic TSPs (those that induce larger differences in the optimal solutions) more heavily.

The mean relative Euclidean-ness of the human solutions to the various TSPs are shown in Figure 3 (the seven 10-point problems with no differences in minimal solution are omitted). Using this measure, there are no observable differences in the 10-point problems, even if we restrict the analysis to the five arrays in which there are some differences in the optimal solution across metrics. That is, a 2×5 ANOVA revealed no significant effect of the TSP metric on the relative Euclidean-ness of the solutions ($F_{1,85} = 0.05, p = 0.83$). However, when we turn to the 40-point problems, the expected effect emerges ($F_{1,192} = 4.25, p = 0.041$), with the human solutions to the city block TSPs scoring lower in relative Euclidean-ness to than the solutions to the corresponding array in Euclidean form.[4]

Discussion

Consistent with previous research (e.g., Dry et al 2006) people deviate further from optimality as a function of the number of nodes, and the length of the shortest solutions. No evidence was found that people's solutions were especially good for the Euclidean relative to the city block problems. Moreover, in the 40-point problems, since the best Euclidean tours differed from the best city block tours for the same arrays, we were able to test the proposition that people were just using Euclidean-derived solutions for the city block problems. Contrary to this proposition, we observe that people tended to produce solutions more appropriate to the geometry in which the problem is embedded. To the extent that these findings generalize and replicate, it appears that the near-optimality of human performance on visual TSPs is not restricted to the Euclidean case.

Conclusion

The fact that people are able to construct TSP solutions that are more appropriate to a city block problem than a Euclidean one is instructive. Arguably, the natural world does not present people with grid-navigation problems, so it is unlikely that evolution would have endowed us with any perceptual apparatus built for solving such problems. Unlike the Euclidean case, in which human TSP performance might well be optimal solely due to natively endowed visual representation, city block optimality is presumably more reliant on reasoning and learned representations. In that sense, the fact that people are near-optimal on two familiar minimal path problems (Euclidean and city block TSPs) but highly suboptimal on an unfamiliar but logically-equivalent one (verbal TSP) is reminiscent of Griffiths and Tenenbaum's (2006) proposition that optimality may attach not only to evolutionarily relevant problems, but to everyday cognition generally. Of course, it is not clear what mental representations support the strong human performance in the city block problems. It may be that through everyday experience people are able to build novel representations more suited to navigation on grids. However, it seems more likely that people may have learned to use high level reasoning to adapt a purely Euclidean visual representation to other problems in a very successful manner. Indeed, this seems consistent with Saalweachter and Pizlo's (2008) finding that people performed well on simple occluded-Euclidean problems, but poorly when the geometry of the problem became more alien. In any case, the high level of performance observed for the city block TSP suggests that the claim that "the perception of optimal structure might be a natural, automatic tendency of the human visual system, as opposed to a specific, task-determined and capacity-limited achievement" (Vickers et al. 2001, p. 36) may hold more generally than previously thought.

Acknowledgements

Correspondence concerning this article should be addressed to Daniel Navarro, School of Psychology, University of Adelaide, SA 5005 Australia (+61 8 8303 5265, daniel.navarro@adelaide.edu.au). This article summarizes research from ALW's honours thesis. DJN was supported by an Australian Research Fellowship (ARC grant DP0773794).

References

Chater, N. & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences, 7*, 19-22.

Chronicle, E. P., MacGregor, J. N., Lee, M. D., Ormerod, T. C. & Hughes, P. (2008). Individual differences in optimization problem solving: Reconciling conflicting results. *Journal of Problem Solving, 2*, 41-49.

Dry, M. J., Lee, M. D., Vickers, D. & Hughes, P. (2006). Human performance on visually presented traveling salesperson problems with varying numbers of nodes. *Journal of*

Problem Solving, 1, 20-32

Feldman, J. (2003). The simplicity principle in human concept learning. *Current Directions in Psychological Science, 12*, 227-232.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science, 17*, 767-773.

Köhler, W. (1929). *Gestalt psychology.* New York: Liveright.

MacGregor, J. N. & Ormerod, T. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics, 58*, 527-539.

Pizlo, Z., Stefanov, E., Saalweachter, J., Li, Z., Haxhimusa, Y. & Kropatsch, W. G. (2006). Traveling salesman problem: A foveating pyramid model. *Journal of Problem Solving, 1*

Polivanova, N. I. (1974). [On some functional and structural features of the visual intuitive components of problem solving processes]. *Voprosy Psikhologi, 4*, 41-51

Saalweachter, J. & Pizlo, Z. (2008) Non-Euclidean traveling salesman problem. In: T. Kugler, J.C. Smith, Y-J. Sun, T. Connolly (Eds), *Decision Modeling and Behavior in Complex and Uncertain Environments.* New York, NY: Springer (pp. 339-358).

Shepard, R. N. (1994) Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review, 1*, 2-28.

Vickers, D., Bovet, P., Lee, M. D., & Hughes, P. (2003). The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean Travelling Salesperson Problems. *Perception, 32*, 871-886.

Vickers, D., Butavicius, M., Lee, M., & Medvedev, A. (2001). Human performance on visually presented Traveling Salesman Problems. *Psychological Research 65*, 34-45.

Vickers, D., Mayo, T., Heitmann, M., Lee, M., & Hughes, P. (2004). Intelligence and individual differences in performance on three types of visually presented optimisation problems. *Personality and Individual Differences, 36*, 1059-1071.

Walwyn, A. L. (2006). *Investigating the roles of visual perception and intelligence in relation to the travelling salesperson problem.* Unpublished honours thesis, University of Adelaide.

Zipf, G. K. (1949). *Human behavior and the principle of least effort.* Cambridge, MA: Addison Wesley

Footnotes

[1]Euclidean and city block spaces are two special cases of a Minkowski space, in which the distance d_{ij} between two points i and j is $d_{ij} = (\sum_{k=1}^{v} |x_{ik} - x_{jk}|^r)^{1/r}$ where v is the dimensionality of the space, x_{ik} is the co-ordinate value of point i along dimension k and r is the underlying metric. Setting $r = 1$ produces the city block metric, and setting $r = 2$ gives the Euclidean metric.

[2]The reason for using more 10-point arrays was that the full experiment also used verbally-specified problems as a control condition, omitted here since performance was predictably poor (see Walwyn, 2006). Additionally, the use of random arrays (rather than arrays chosen to maximize the difference between the Euclidean and city block minimal tours) was that we are genuinely interested in the typical case, and moreover were not convinced that selecting atypical arrays to maximize effect size would not alter the array structure in undesirable ways.

[3]The fact that the missing data fall disproportionately in the 40 point city-block cases, does suggest a data missing not-at-random problem: it appears to be the case that in the pencil-and-paper format the 40-point problem can get a little cluttered (moreso than the Euclidean version), making it easier for both the participant (and the experimenter) to miss a point when generating (or transcribing) the solution. The missing solutions do not seem to be atypical in any other respect however.

[4]There is a main effect of array for both the 10-point problems ($F_{4,85} = 3.81, p = 0.007$) and the 40-point problems ($F_{5,192} = 3.40, p = 0.006$), but this is to be expected given the difference in diagnosticity of the arrays and not of any particular interest.

Figure Captions

Figure 1. Overall character of the solution lengths. In general, human performance is near-optimal (left panel), and the deviations from optimality vary smoothly as a function of the length of the shortest solution (right panel). In both panels, lighter dots are solutions to the Euclidean problems, and darker dots refer to city block problems. The shorter lengths all correspond to the 10-point problems, and the longer ones refer to the 40-point problems.

Figure 2. The best-found solutions for one of the 40 point arrays, in both the city block metric (left) and Euclidean metric (right). Connections that exist in one solution but not the other are shown with darker lines (11 such in both arrays). To see why the solutions differ, note that Euclidean distances are the same as the city block distances for two points that differ only in one co-ordinate (vertical or horizontal), but can differ by a factor as large as $\sqrt{2}$ if the points differ along both dimensions. This systematically alters the structure of the distances, producing subtly different solutions.

Figure 3. Differences between conditions in terms of the extent to which the human solutions more closely resemble the optimal Euclidean solution or the optimal city block solution. For the five 10-point problems in which the optimal tours are distinguishable, there are no clear differences between participant solutions, presumably because people are so close to optimal when the number of points is so small (Figure 1). For the 40-point problems, however, the human solutions alter as a function of format in the appropriate manner.

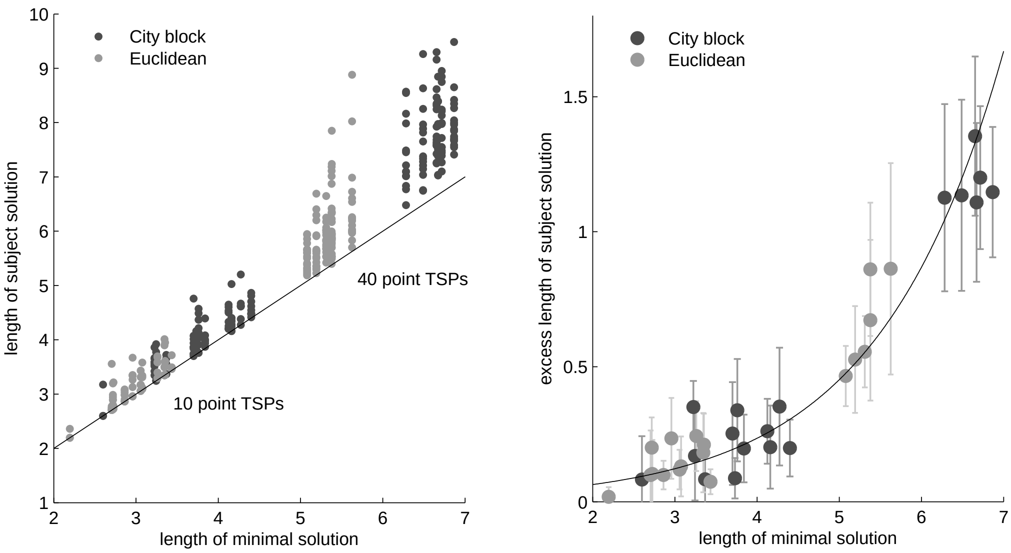

City block
Euclidean
length of subject solution
length of minimal solution
40 point TSPs
10 point TSPs
City block
Euclidean
excess length of subject solution
length of minimal solution

City block solution

Euclidean solution

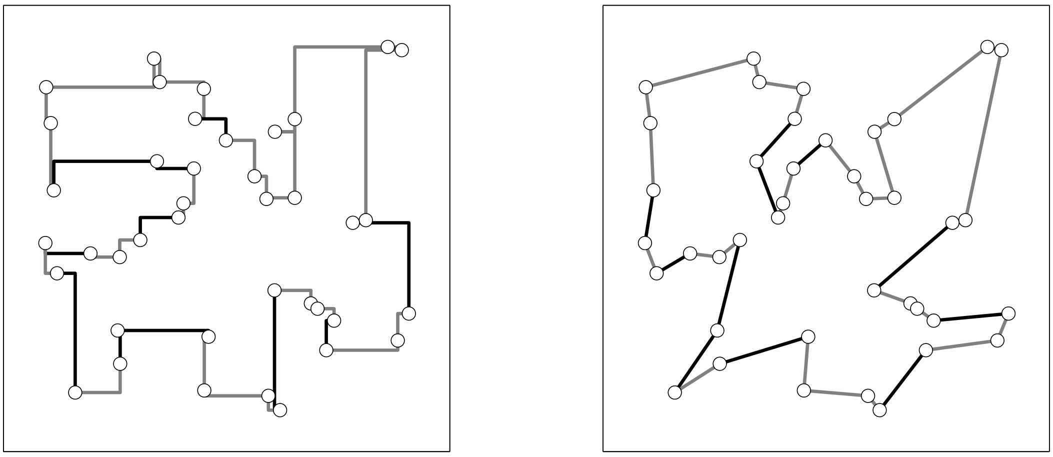

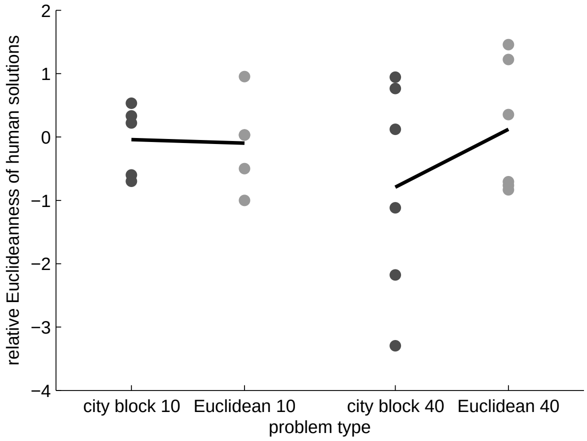

relative Euclideanness of human solutions
2
1
0
−1
−2
−3
−4
city block 10
Euclidean 10
city block 40
Euclidean 40
problem type