

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

Amy L. Walwyn¹ and Daniel J. Navarro¹

Abstract:

An experiment is reported comparing human performance on two kinds of visually presented traveling salesperson problems (TSPs), those reliant on Euclidean geometry and those reliant on city block geometry. Across multiple array sizes, human performance was near-optimal in both geometries, but was slightly better in the Euclidean format. Even so, human solutions varied systematically across geometries, with the differences reflecting sensible adaptations to the different distance structures that these geometries imply. These results suggest that human near-optimality for visually presented TSPs may apply to a more general class of problems than the standard Euclidean case.

Keywords

optimization, traveling salesperson problem, visual perception

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 and extends research from ALW's honors thesis. DJN was supported by an Australian Research Fellowship (ARC grant DP0773794). We thank Matt Dry, Michael Lee, Jim MacGregor, and Zyg Pizlo for helpful comments and discussions.

¹University of Adelaide

Introduction

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

Not surprisingly, human performance on the TSP is tied to the 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 a verbal 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, it seems likely that the critical predictor of human performance would be the degree to which the relevant representational system is matched to the format 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 operate in a universe that is (locally) Euclidean, so it is natural to expect the visual system to be adapted to this structure (e.g., Shepard, 1984, 1994), via representations that are innately specified. In contrast, it would be highly surprising if people were optimal with respect to problems formulated in an arbitrary Riemannian geometry: nothing in our evolutionary or everyday experience prepares us for such problems. However, between these two extremes lies 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 TSPs 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 human performance to be more robust. The only work of which we are aware that discusses such problems is by Saalweachter and Pizlo (2008), who considered TSPs involving obstacles or mazes, and found that people are 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 ages 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 12 10-point random dot arrays and 6 40-point arrays, presented in pencil-and-paper form.² Each point was sampled from a uniform distribution over a 158 mm x 158 mm square, subject to the constraint that the points be visually distinct: all distances were calculated relative to the unit square. The use of random arrays (rather than arrays chosen to maximize the difference between the Euclidean and city block minimal tours) was a deliberate choice. Given that this is the first investigation of this class of problems, we felt it would make sense to consider the "typical" case. Moreover, we were somewhat concerned that selecting atypical arrays to maximize effect size might alter the array structure in undesirable ways. Minimal tours for the 10-point problems were found by exhaustively checking all possible solutions; for the 40-point problems these were estimated in both city block and Euclidean geometries using standard numerical optimization methods (stochastic hill-climbing with multiple restarts; e.g., Nocedal & Wright, 1999).

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 counter-balanced. 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.

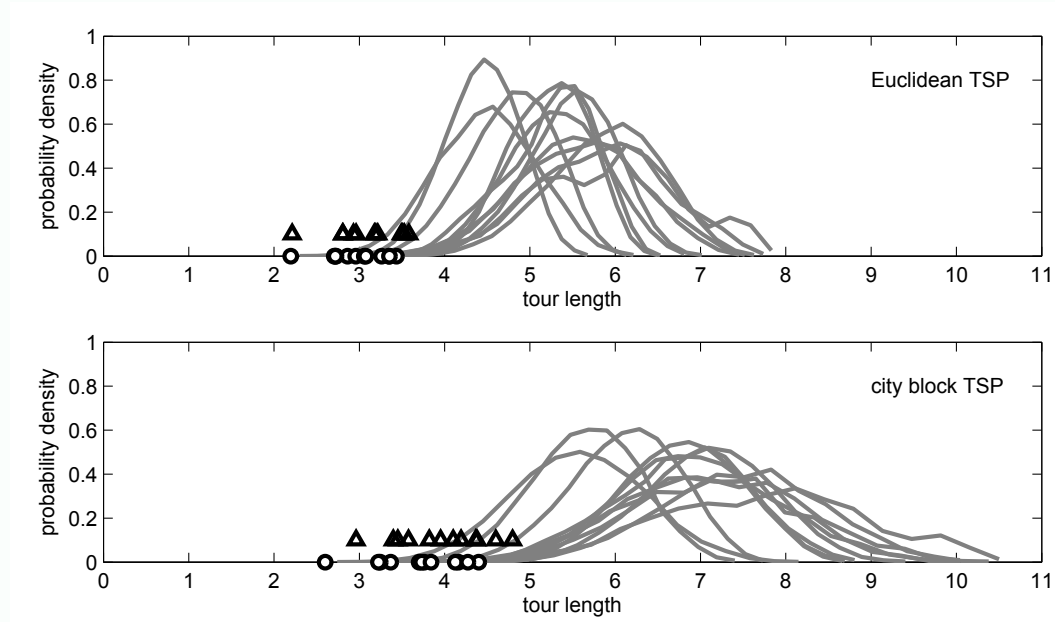
Exclusions. Four participants produced extremely long paths for at least one of the problems, relative to the other 36 participants. This happened for both Euclidean and city block problems. Since we could not be certain that these four participants were following the instructions properly, all data from those four people were excluded. Additionally, some 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 Statistical Character of the Problems. The 10-point TSPs have only $9!/2 = 181440$ possible solutions, so it is feasible to compute the lengths of all tours. The distributions over tour lengths for all 12 problems are shown in Figure 1, along with the length of the minimal tour. The descriptive statistics for these distributions are shown in Table 1, along with counts of the number of tours that achieve the minimum length and the number of tours that lie within 2.5% of the minimum (i.e., are not more than 2.5 “percent above optimal,” or PAO). Notice that for the city block TSPs, the optimal solution is not necessarily unique. Moreover, regardless of metric the number of tours that are less than 2.5% above optimal is extremely variable as a function of array, suggesting that although PAO is often used as a measure of human performance it should be interpreted with considerable caution.

For the 40-point problems, an exhaustive search of the solution space is infeasible, since there are $39!/2 = 6.8 \times 10^{45}$ tours. Nevertheless, we can approximate the relevant distributions in the following manner. First, we approximate the minimal tour by applying standard combinatorial optimization routines. Second, we estimate the main body of the distribution by constructing 100,000 random tours for each problem, and plotting the implied distribution over solution lengths. The results, shown in Figure 2, exhibit the same pattern as Figure 1: the city block tours are longer and more variable. Moreover, the fact that the distributions are approximately Gaussian suggests that for each array there is a natural scale upon which to measure the length of minimal solutions: namely, the z-score, the number of standard deviations below the mean tour length. For the city block problems the minimal tour lies (on average) 11.2 standard deviations below the mean. For the Euclidean problems the minimum is slightly further away, at 11.8 standard deviations.

Figure 1. Exact distribution over solution lengths for all 10-point TSPs, whether in Euclidean form (top) or in city block form (bottom). Circular markers show the minimum for each problem, while triangular markers show the mean human solution lengths for the various problems. Notice that the minimal tour length and the mean tour length tend to be longer for the city block problems, and also that the solutions for the city block problem are more variable.



The Near-Optimality of Human Solutions. The lengths of all solutions produced by participants are shown in the left panel of Figure 3, plotted as 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, suggesting that human solutions are near-optimal in both cases. In order to formalize this intuition for the 10-point problems, it is possible simply to count the number of tours that are better than the human tour. Across all problems, 12% of the Euclidean solutions were optimal, and 21% of the city block solutions were optimal. Furthermore, 73% of the Euclidean solutions were in the “top 20,” compared to only 58% for the city block case. Noting that the top 20 corresponds to the best 0.003% of possible solutions, we can reasonably conclude that people are near-optimal for these problems in both metrics. As to whether the solutions are better for one metric or the other, the evidence is ambiguous: the median rank of the Euclidean solutions was 6 (interquartile range: 2-23), and for the city block solutions it was 14 (interquartile range: 3-52.25). A Kruskal-Wallis test comparing these medians is non-significant but borderline ($z = 1.75, p = 0.08$).

Table 1. Statistics for the 10-point TSPs, showing the length of the minimal tour, the mean tour length, the standard deviation of tour lengths, the number of unique tours that achieve the minimal length, and the number of tours that are not longer than the minimum by more than 2.5%.

TSP	city block					Euclidean				
	min	mean	std dev	#min	#close	min	mean	std dev	#min	#close
1	3.70	6.12	0.64	2	6	2.86	4.79	0.51	1	7
2	4.16	7.49	0.94	4	12	3.35	5.89	0.66	1	8
3	3.76	6.80	0.76	1	2	2.96	5.30	0.57	1	1
4	3.37	7.03	0.97	1	3	2.71	5.56	0.70	1	3
5	3.73	7.23	1.03	1	11	3.06	5.84	0.77	1	2
6	4.40	7.10	0.75	2	12	3.43	5.50	0.53	1	10
7	3.24	7.01	0.99	1	1	2.73	5.58	0.74	1	4
8	3.22	5.65	0.62	1	2	2.72	4.44	0.43	1	4
9	4.13	6.75	0.78	1	31	3.26	5.23	0.55	1	27
10	3.84	7.71	1.22	2	14	3.08	5.90	0.88	1	7
11	2.60	5.57	0.77	1	1	2.20	4.49	0.58	1	1
12	4.27	6.80	0.69	1	2	3.35	5.23	0.47	1	1

For the 40-point problems this approach is infeasible. In view of the concerns with PAO, we associate each human solution with its corresponding z-score. On average, the human solutions to the Euclidean problems fell 11.2 standard deviations below the mean tour length; for the city block problems they fell 10.5 standard deviations below. The median deviation from optimality (in terms of number of standard deviations) for the Euclidean problems was 0.46, whereas for the city block problems this difference was 0.63, again suggestive of near-optimality. Moreover, a Kruskal-Wallis test comparing these deviations from optimality suggests the Euclidean solutions are somewhat superior ($z = 3.37, p < 0.001$).

Differences in Human Tours as a Function of Geometry. The previous section leads us to conclude that human performance is near-optimal for both metrics, but that it may be slightly better for the Euclidean problems. This leads one to ask whether people really treat the city block TSP any differently than the Euclidean TSP. That is, it may be that people construct city block tours using the exact same method that they use for Euclidean problems, and it happens to be the case that this does not impair performance by very much. Since the 10-point problems do not elicit a lot of variability (e.g., only one problem

Figure 2. Approximate distribution over solution lengths for all 40-point TSPs, whether in Euclidean form (top) or in city block form (bottom). Circular markers show the minimum for each problem, while triangular markers show the mean human solution lengths for the various problems. Notice that the minimal tour length and the mean tour length tend to be longer for the city block problems, and also that the solutions for the city block problem are more variable.

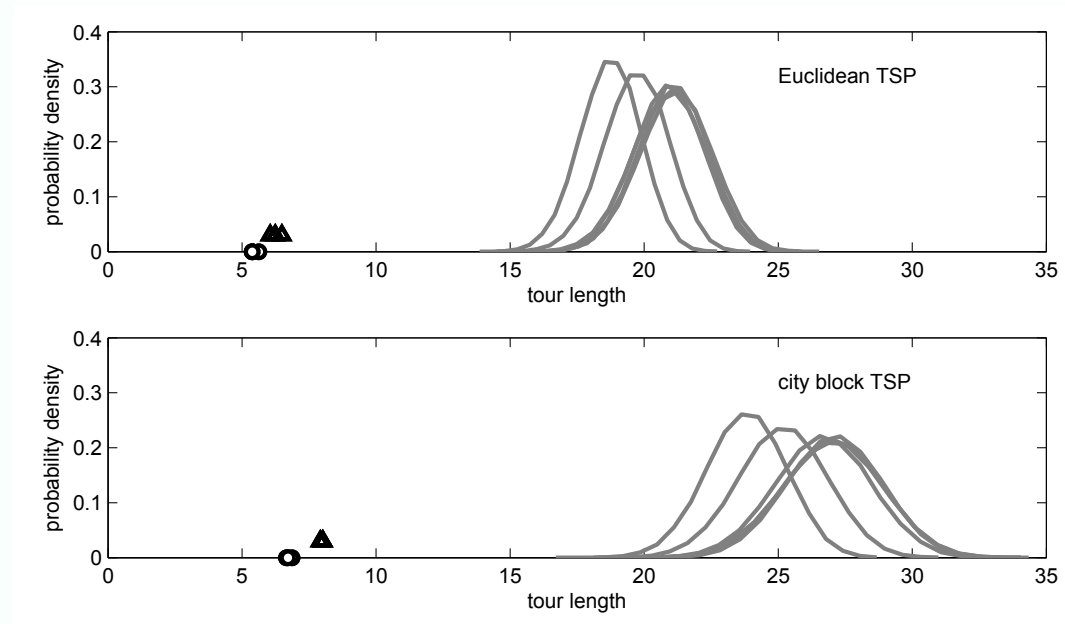
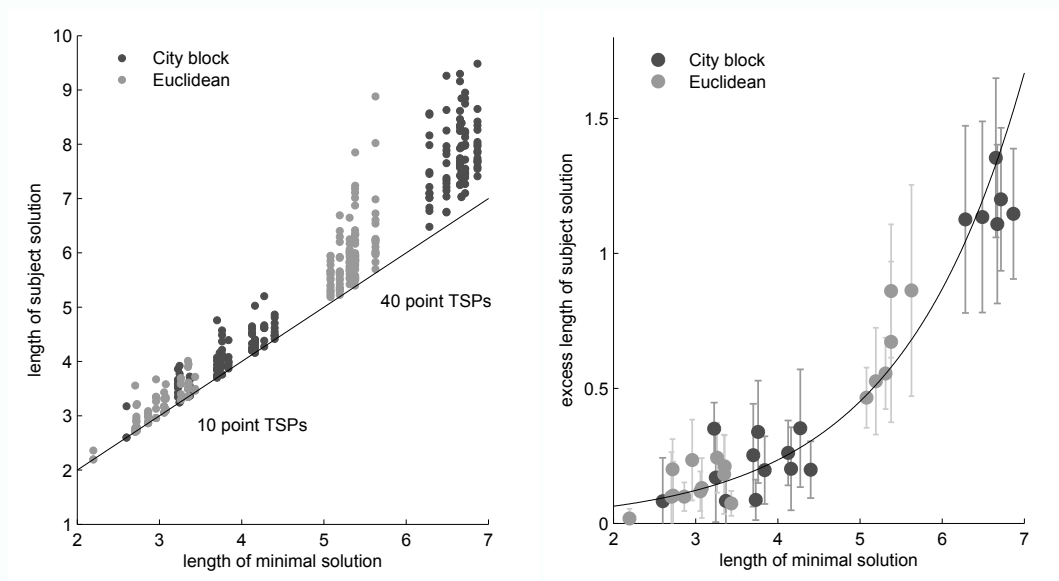


Table 2. Approximate statistics for the 40-point TSPs, showing the length of the minimal tour, the mean tour length, and the standard deviation of tour lengths. On the basis of these, we also compute the implied z-score for the minimal tour.

TSP	city block				Euclidean			
	min	mean	std dev	z-min	min	mean	std dev	z-min
1	6.87	27.07	1.77	-11.43	5.63	21.06	1.29	-11.95
2	6.65	27.00	1.86	-10.93	5.38	20.90	1.33	-11.65
3	6.72	26.66	1.79	-11.16	5.38	20.90	1.29	-11.98
4	6.87	27.02	1.85	-10.92	5.63	21.17	1.34	-11.56
5	6.65	23.79	1.48	-11.55	5.38	18.72	1.12	-11.92
6	6.72	25.16	1.66	-11.12	5.38	19.70	1.20	-11.96

Figure 3. 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.



has a unique city block minimal solution that differs from the corresponding Euclidean optimum), we restrict the discussion to the 40-point problems. For these problems, the estimated minimal tour was never the same for the two metrics: the number of disagreements ranged from 5 to 15. An example is shown in Figure 4.

The fact that the optimal solutions differed across metrics means that there are two different benchmarks against which we can compare any given human tour, one appropriate to the city block TSP, the other appropriate to the Euclidean problem. With this in mind, for every human 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 for the same array. For an array in which the best solutions differ by k edges, this measure varies from $-k$ to k , and hence weights the more diagnostic TSPs (those that induce larger differences in the optimal solutions) more heavily. The mean relative Euclidean-ness for human solutions to each 40-point TSP is shown in Figure 5. A 2×6 ANOVA revealed a significant effect of metric ($F_{1,192} = 4.25, p = 0.041$), with the human solutions to the city block TSPs scoring lower in relative Euclidean-ness than the solutions to the corresponding array in Euclidean form. (There is also a main effect of array [$F_{5,192} = 3.40, p = 0.006$], but this is to be expected given the difference in diagnosticity of the arrays and not of any particular interest.)

Figure 4. The best solutions found 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 coordinate (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.

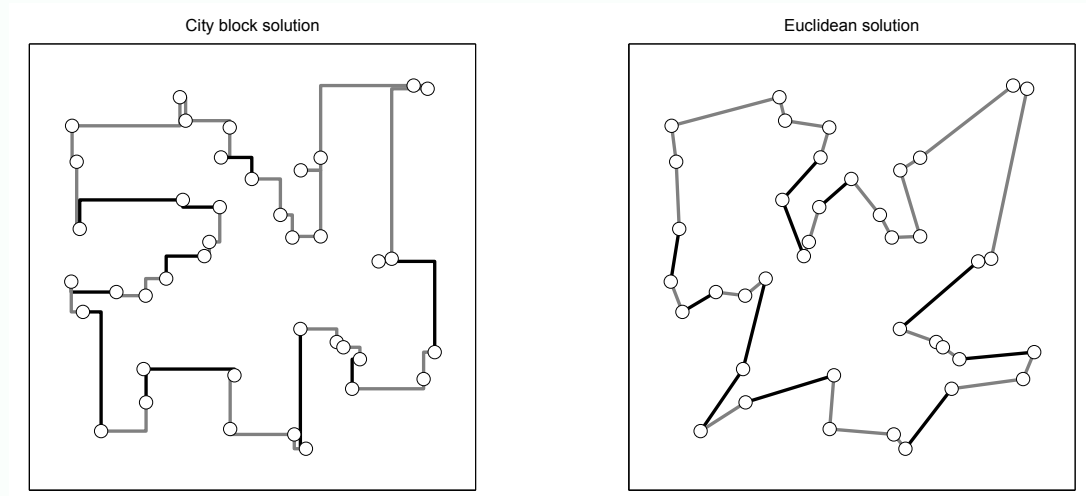
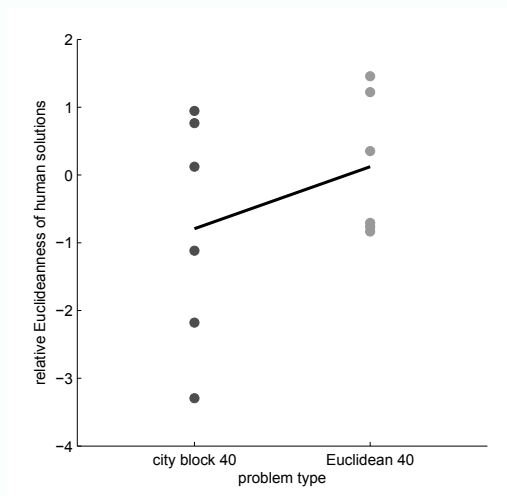
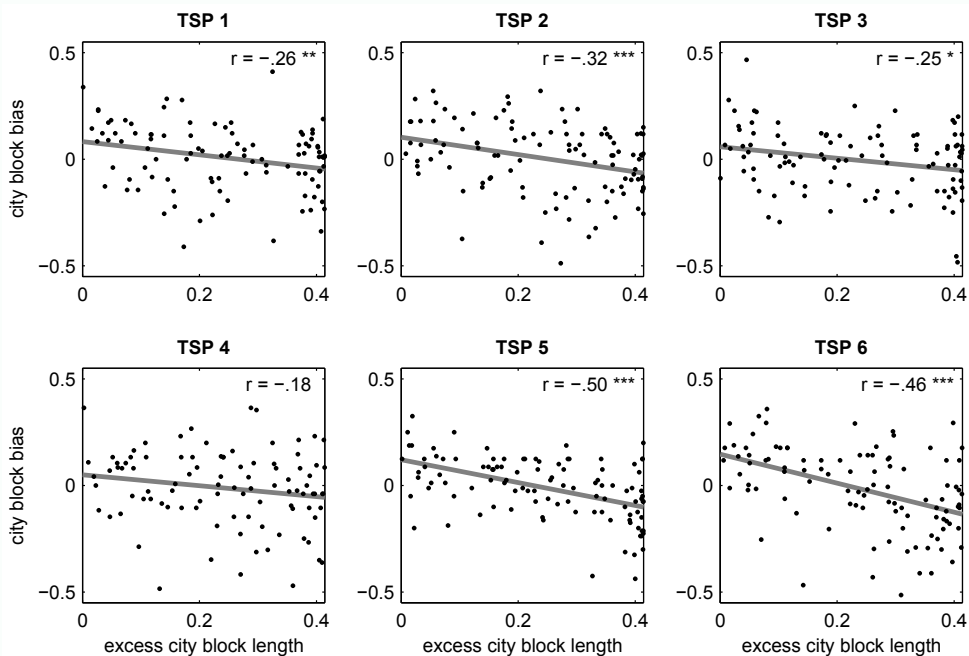


Figure 5. Differences between conditions in terms of the extent to which the human solutions to 40-point TSPs more closely resemble the optimal Euclidean solution or the optimal city block solution. Overall, there is evidence that the human solutions alter as a function of format in the appropriate manner.



The good thing about the previous analysis is that it allows us to speak directly about the relationship between human performance and the two different types of optimal solution. As such, it makes comparisons between one kind of good TSP solution (the human tours) and another kind of good TSP solution (the minimal tours). The disadvantage, however, is that minimal tours do not represent typical human performance, and so may not provide quite the right standard against which to compare human tours. An alternative approach is to take advantage of the fundamental differences between the two metrics. Specifically, for any two points, the city block distance between them varies from being identical to the Euclidean distance (when the points lie on the same horizontal or vertical line), to being two times as long (when the horizontal displacement is the same as the vertical displacement). If people do adapt their solutions to suit the metric, then those links that yield much longer city block distances than Euclidean distances should be somewhat less likely to appear in people's city block solutions, relative to those links

Figure 6. An analysis of the relative likelihood of different links appearing in human solutions. Each dot corresponds to a link that was included by at least two participants in at least one metric. The horizontal location denotes the extent to which the length of that link is longer in the city block metric, which necessarily ranges from 0 to $\sqrt{2} - 1$. The vertical location plots the probability that the link is included in a city block solution, minus the probability that it appears in a Euclidean solution. Regression lines are shown in gray, and the significant negative correlations indicate that people systematically tended to use links in a metric-appropriate fashion. (Notation: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).



that are the same length in both metrics.⁴ As is illustrated in Figure 6, this is precisely what happens. Each panel plots the probability that a given link appears in a city block solution minus the probability that it appears in a Euclidean solution, as a function of the extent to which the city block length for that link is longer (ranging from 0 to $\sqrt{2} - 1$). Only those links that appear in at least two human solutions are shown. As expected, all correlations are negative, and five of the six are significant.

Conclusion

In terms of solution length, we found evidence that human tours were nearly optimal for both Euclidean and city block metrics, but slightly superior for the Euclidean case. Nevertheless, we also found evidence consistent with the proposition that the city block solutions were not solely reliant on Euclidean representations, since the human tours differed systematically between metrics, in a manner that reflects sensible adaptation to the differences between the two types of problem.

Discussion

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 a perceptual system built for solving such problems. Unlike the Euclidean case, in which human TSP performance might well be optimal solely due to natively endowed visual representations, 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 TSPs) 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 have learned to use high-level reasoning to adapt a purely Euclidean visual representation to other problems in a very successful manner. Intuitively, this seems quite plausible: many navigation problems in evolutionarily relevant contexts involve distortions to an otherwise regular metric structure (e.g., navigation through hilly terrain). In view of this, robustness in the face of such distortions would be a very useful capability. 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. Future work could consider the extent

to which theoretical principles such as crossing avoidance and convex hull use (e.g., MacGregor & Ormerod, 1996; Vickers et al., 2003) generalize from the Euclidean TSP to the city block version, and whether models such as Pizlo et al.'s (2006) pyramid search theory are similarly robust. However, a detailed discussion of these ideas is beyond the scope of this paper. In the meantime, 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.

Endnotes

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).
3. The fact that the missing data fall disproportionately in the city block cases does suggest that this is 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 (more so 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. In a sense, this is the same regularity exploited by the relative Euclidean-ness measure, since the optimal city block tour is more likely to avoid edges that are particularly long under the city block metric. For relative Euclidean-ness, only those edges that belong to one of the two minimal solutions are considered, whereas for the analysis in Figure 6 all edges are considered. As such, the two analyses are complementary.

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.
- Nocedal, J., & Wright, S. J. (1999). *Numerical optimization*. New York: Springer Verlag.
- 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*, 83–101.
- Polivanova, N. I. (1974). [On some functional and structural features of the visual intuitive components of problem-solving processes]. *Voprosy Psikhologii, 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* (pp. 339–358). New York: Springer.
- Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review, 91*, 417–447.
- 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. D., & Medvedev, A. (2001). Human performance on visually presented Traveling Salesman Problems. *Psychological Research, 65*, 34–45.
- Vickers, D., Mayo, T., Heitmann, M., Lee, M. D., & Hughes, P. (2004). Intelligence and individual differences in performance on three types of visually presented optimisation problems. *Personality and Individual Differences, 36*, 1059–1071.
- Vickers, D., Lee, M. D., Dry, M. J., & Hughes, P. (2003). The role of the convex hull and the number of potential intersections in performance on visually presented traveling salesperson problems. *Memory & Cognition, 31*, 1094–1104.
- Walwyn, A. L. (2006). *Investigating the roles of visual perception and intelligence in relation to the travelling salesperson problem*. Unpublished honors thesis, University of Adelaide.
- Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Cambridge, MA: Addison Wesley.