

Effects of Cluster Location on Human Performance on the Traveling Salesperson Problem

*James N. MacGregor*¹

Abstract

Most models of human performance on the traveling salesperson problem involve clustering of nodes, but few empirical studies have examined effects of clustering in the stimulus array. A recent exception varied degree of clustering and concluded that the more clustered a stimulus array, the easier a TSP is to solve (Dry, Preiss, & Wagemans, 2012). However, a limitation to this conclusion arises because degree of clustering may have been partially confounded with cluster location. An experiment was conducted to test the effects of cluster location while holding degree of clustering constant. Stimuli with a cluster near a boundary were solved more quickly and accurately than stimuli with the same cluster located more centrally. The results support and extend the previous findings of MacGregor, Ormerod, & Chronicle (1999). They also qualify the results of Dry et al. (2012). To the extent that degree of clustering may have been confounded with the location of clusters in their stimuli, it is unclear to what extent each factor may have affected performance.

Keywords

problem solving, optimization, clustering

¹University of Victoria. Please direct correspondence to jmacgreg@uvic.ca.

Introduction

Over the past decade, interest in human performance on the travelling salesman problem (TSP) and related tasks has been steadily increasing, and a substantial body of empirical and theoretical results has been established. The TSP requires finding the shortest path through a set of locations, returning to the starting location. As well as representing a fundamental class of issues in computer science and operations research, the TSP underlies many important applications in science and technology.

A recent review of human performance in the task cited some 50 research papers, many reporting empirical findings (MacGregor & Chu, 2011). The review also summarized five different theoretical approaches, and identified several factors which differentiated among them. One factor was the extent to which models assume that nodes are clustered during the solution process. Strikingly, while almost all theoretical accounts assume that clustering plays a fundamental role, almost no empirical research has systematically examined the effects of clustering.

One of the few exceptions was a study by Hirtle and Gärling (1992) using “open” versions of the TSP (where start and end points differ). The experiment used two levels of node distribution (clustered versus homogenous) and found evidence that participants attended to clusters in constructing paths. Although no effects of level of clustering were reported, MacGregor and Chu (2011) conducted a descriptive reanalysis of the results which suggested that performance was generally superior for the clustered than the homogenous point distributions.

MacGregor, Ormerod, and Chronicle (1999) reported an experiment using six 12-node problems in which the location of interior nodes was varied (central, intermediate, and peripheral). The hypothesis was that locating interior nodes closer to the boundary would make the task simpler and result in superior performance. Although the experiment did not address clustering explicitly, the interior nodes formed a single salient cluster in the central and peripheral conditions. The results indicated superior performance when the cluster was located near the boundary than centrally. Specifically, two-thirds of solutions were optimal when the internal cluster was close to the boundary, compared with less than one-third of solutions when the cluster was located centrally.

The most extensive and systematic study of clustering in TSPs to date has been reported recently by Dry, Preiss, and Wagemans (2012). In addition to reporting the results of an experiment, the paper applied the Clark-Evans Index of Aggregation, R , as a measure of clustering. Clark and Evans (1954) introduced the statistic as a means of measuring the distribution of a species within a defined area, based on the observed proximity of individuals to the proximity expected by chance. Observed proximity, r_o , is calculated from the distance of each member to its nearest neighbour, averaged across members. That is, where the distance between nodes i and j is given by d_{ij} ,

$$r_0 = \frac{1}{n} \sum_{i \neq j}^n \min(d_{ij})$$

The value expected by chance, r_E within a defined area, A , is given by the Poisson distribution,

$$r_E = .5 \sqrt{A/n}$$

The R statistic is the ratio of the observed value to the expected value. By definition, the value of R will be 1 if the observed distribution is random. The theoretical minimum value of R is 0, when members occupy the same locations. Thus, as R values decrease from 1 to 0, the average spacing of members decreases and clusters form. As values increase above 1, average spacing increases to a maximum value of 2.1491, when a regular hexagonal pattern is formed, with each member having 6 equidistant nearest neighbours (Clark and Evans, 1954).

Dry et al. (2012) proposed that, if the perception of clusters plays a role in the cognitive representation of a TSP, then the presence in stimuli of salient clusters should facilitate the formation of the problem representation and result in faster and better solutions than with stimuli where nodes are more regularly spaced. To test this, they created fifteen 49-node TSPs, 5 each at three different levels of R corresponding to highly clustered, random, and highly regular distributions. Twenty-five participants each completed the 15 TSPs, and performance was measured in terms of solution path length (expressed as a percentage above the optimal, or PAO) and solution times. The results partially supported the research hypothesis, in that PAO increased significantly with increasing R values. Solution times increased in the same direction, but differences between conditions were non-significant. The authors interpreted the findings as providing empirical support for the proposal that human solvers employ a clustering heuristic when solving TSPs. However, because of the way stimuli were constructed an additional, potentially alternative, explanation of the results suggests itself. This relates to the earlier hypothesis proposed by MacGregor et al. (1999), concerning the effects of the locations of interior nodes.

To produce the stimuli for the low R (highly clustered) condition, Dry et al. (2012) first randomly determined the sizes of clusters, ranging from 2 to 5 nodes, which in turn established the number of clusters. Next, locations were randomly sampled, one for each of the clusters, and initially each member of a cluster was superposed at that location. In other words, nodes in the same cluster were initially at zero distance from each other, corresponding to an initial R value of 0. Finally, random "jitter" was applied to each node in a cluster, displacing its position until the R value increased to 0.62 +/- 0.001.

Because of the method of generating them, stimuli in the low R condition may have differed from random stimuli not only in the degree of clustering but also in the locations of nodes and clusters. The reason is that by initially placing all members of a cluster at the same location, the number of nodes falling on (or near) the convex hull may have been increased artificially. For example, if clusters range in size from 2 to 5, the expected

cluster size will be 3.5 and the expected number of clusters will be 14 (49 nodes /3.5). In sampling 14 locations randomly, we would expect approximately 7 to fall on the convex hull, based on Monte Carlo simulations reported by Philip (2004). However, because of the initial multiple occupancy of an average of 3.5 nodes at the same locations, the effect would be that an expected 24.5 nodes (7 x 3.5) would initially fall on the convex hull. In contrast, if 49 nodes were randomly located, the expected number on the convex hull would be approximately 10, based on Philip's (2004) results.

Although Dry et al. (2012) selected stimuli to have approximately the same number of nodes on the convex hull, this would not control for the number of nodes close to the hull. While the effects of random jitter would be to displace from the hull some (or all) node members that started on it, the fact that nodes remained in sufficient proximity to retain integrity as a cluster implies that they would have remained in similar proximity to the convex hull. MacGregor et al. (1999) reported that performance is superior when interior points are located nearer the boundary than centrally. This raises the possibility that the location of nodes, as well as the degree of clustering, may have influenced performance in the Dry et al. experiment.

As a first step towards investigating these possibilities, an experiment was conducted to examine the effects of cluster location while holding constant the degree of clustering (as measured by R). The primary purpose was to test whether performance was superior when clusters were located near the boundary than located centrally. Also, the experimental design allowed a test of a more basic question, of whether adding a cluster of points to an array increases the difficulty of a TSP. On one hand, it seems that it would, since the effect is to increase the number of nodes. On the other hand, if processing of clusters is automatic, then adding a cluster may be virtually cost-free.

Experiment

The main purpose of the experiment was to examine whether the location of a cluster of nodes—central or peripheral—influences performance.

Method

Participants

Participants were 20 volunteers recruited from the campus community at the University of Victoria.

Stimuli

The stimuli were fifteen point sets. Each point set was displayed within an on-screen area of 125x134mm, using nodes of radius 3mm. There were five randomly-generated 20-node sets, five 30-node sets each with a central cluster of 10 nodes, and five 30-node sets with a cluster

of 10 nodes near the boundary. The stimuli were produced in the following way. First, five clusters of 10 nodes were randomly-generated so that a cluster fit within a circle of radius 17.5 mm. Second, five 20-node point sets were generated randomly, subject to the constraints that there was sufficient unoccupied space to place a circular area of 17.5mm radius (a) within a central rectangle of 62.5x67mm, and (b) between two adjacent boundary nodes. Third, each cluster was randomly assigned to one of the five point sets and used to create central cluster and boundary cluster stimuli. In the former case, the cluster was randomly placed within the central area, subject to the constraint that no point in a cluster was closer than 6mm to an existing point. In the latter case several steps were required: (i) a preliminary direction was randomly selected from the 4 cardinal and 4 ordinal points of the compass; (ii) boundary arcs were identified that lay in the selected direction from the centre, and the longest of these was identified (for example, the longest arc on the north-west boundary of the point-set); (iii) the cluster was placed as close as possible to midway between the two end nodes of the selected arc, such that no point in the cluster was closer than 6mm to an existing point or 3mm to the boundary arc. In this way, each of the five original 20-node point-sets was used as the basis for a group of three stimuli, one being the original 20-node stimulus, the second being the original stimulus plus a 10-node cluster located in the centre, the third being the original stimulus plus the same 10-node cluster located near the periphery.

The stimuli are shown in Figure 1. The average R-value for the centre cluster stimuli was 1.07 and for the boundary cluster stimuli, 1.06. Dry at al. (2012) reported that, as R increases, the number of potential intersections also increases. Because increasing potential intersections leads to better performance (Vickers, Lee, Dry, & Hughes, 2003), Dry at al. (2012) controlled for number of potential intersections in generating stimuli. However, in the present case, because the R-values of the critical stimuli were held constant, it was not necessary to take this precaution.


Procedure

Testing was conducted individually and each participant produced a TSP tour for all fifteen stimulus point-sets. The order of presentation of stimuli was randomized, except that no member of the same group of three was presented immediately after another. Also, the second member of each group to be presented was rotated through 180°, to reduce the chances that a participant might notice similarities between related stimuli. Stimuli were presented on screen and participants responded by pointing-and-clicking. At the end of the procedure participants were asked if they noticed whether or not some problems contained a cluster of points and, if so, where clusters were located.

Results and Discussion

On being asked, all participants reported that they noticed that some problems contained a cluster. However, only six participants noticed that clusters were sometimes centrally located and at other times, closer to a boundary.

Figure 1. TSP instances used in the experiment. Each row shows a random 20-node instance (first column), and the same instance with a 10-node cluster inserted near the centre (second column) and near the boundary (third column).



The length of the optimal tour was found for each of the 15 stimulus sets, and participants' tours were expressed in terms of percentage length above the optimal (PAO). One participant appeared not to have understood, or engaged in, the task, producing tours more than 4 standard deviations above the average for both the 20-node and 30-node problems, and these data were excluded from further analysis. For the remaining 19 participants, tours had mean lengths (and standard deviations) of 4.4% (3.8%), 5.5% (4.7%) and 6.3% (2.8%) above optimal, for the no cluster, boundary cluster, and central cluster conditions, respectively.

Results were analyzed using repeated measures ANOVA, which indicated an overall significant difference among conditions, $F(2, 36) = 6.40$, $MSe = .0004$, $p < 0.01$. Three contrasts were of interest, one each between the 20-node condition and each of the 30-node conditions, and the third between the two 30-node conditions. The first two addressed the question of whether adding a cluster to a given display increased problem difficulty, and the third addressed whether the location of a cluster had an effect.

Adding a cluster of 10 nodes to a 20-node problem increased solution length from 4.4% above optimal to 5.5% above when the cluster was near the boundary, and to 6.3% above optimal, when the cluster was near the centre. The former contrast was not significant, $t(18) = 1.58$, $p = 0.13$, while the latter was, $t(18) = 3.05$, $p < 0.01$. The contrast between the two 30-node conditions was also significant, $t(18) = 2.21$, $p < 0.05$.

For solution times, means (and standard deviations) were 55.66 (32.42), 85.60 (64.84), and 98.49 (73.27) seconds, for the no cluster, boundary cluster, and central cluster conditions, respectively. The overall difference was significant, $F(2, 36) = 15.02$, $MSe = 610590$, $p < 0.01$. In this case, all three contrasts were significant, $t(18) > 3.00$, $p < 0.01$.

The results indicated that adding clusters of 10-nodes to randomly-generated 20-node TSPs is not cost-free, but that the impact depends on cluster location, with clusters near the centre having a greater negative effect than clusters near the boundary, on both solution quality and completion time. The results appeared to be independent of differences in the degree of clustering as measured by the Clark-Evans Index, which was virtually identical in the two cluster conditions, at 1.07 for the centre and 1.06 and for the boundary cluster stimuli.

General Discussion

Although most models of human TSP performance propose clustering of nodes as part of the solution process, few empirical studies have directly examined the effects on performance of clustering. A recent exception reported the effects of degree of clustering on both solution length and response time, concluding that the more clustered a TSP stimulus array, the easier it is to solve (Dry et al. 2012). The authors interpreted the findings as providing empirical support that people employ a clustering heuristic in producing TSP tours. However, a caveat to this conclusion arises because the method of generating stimuli may have resulted in confounding degree of clustering with cluster location. Specifically, as clustering increased, clusters may have been more likely to fall closer to the boundaries than to the centre of a stimulus array. Previous research has found that TSPs are easier to solve when internal nodes are located near a boundary rather than centrally (MacGregor et al., 1999), and so proximity to a boundary provides an additional and, potentially alternative, explanation to that of clustering.

At the same time, the Dry et al. findings may provide an alternative explanation for the MacGregor et al. (1999) result. That is, in the latter, degree of clustering may have differed

as well as the location of nodes, and caused the difference in performance. The present experiment was conducted to test whether varying cluster location affects performance while holding degree of clustering constant. The results of the experiment indicated that this was the case. When stimuli with equivalent R values were compared, those with a cluster near a boundary were solved more accurately and more quickly than those where the same cluster was located centrally. The results support and extend the previous findings of MacGregor et al. They also qualify the results of Dry et al. To the extent that degree of clustering may have been confounded with the location of clusters in their stimuli, it is unclear to what extent each factor may have affected performance.

The present study was conducted as a preliminary investigation, and the results have provided evidence that the degree of clustering alone, as measured by R, does not fully explain performance. What remains to be investigated is whether degree of clustering, as measured by R, explains performance at all, or whether the apparent effects of R can be completely accounted for in terms of cluster location. To answer this, future research will vary both degree and location of clustering independently.

Acknowledgments

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

References

- Clark, P. J., & Evans, F. C. (1954). Distance to nearest neighbor as a measure of spatial relationships in populations. *Ecology*, *35*(4), 445–53.
- Dry, M., Preiss, K., & Wagemans, J. (2012). Clustering, randomness, and regularity: Spatial distributions and human performance on the traveling salesperson problem and minimum spanning tree problem. *Journal of Problem Solving*, *4*(1), 1–17. Retrieved from <http://docs.lib.purdue.edu/jps/vol4/iss1/2>
- Hirtle, S. C., & Gärling, T. (1992). Heuristic rules for sequential spatial decisions. *Geoforum*, *23*(2), 227–38. [http://dx.doi.org/10.1016/0016-7185\(92\)90019-Z](http://dx.doi.org/10.1016/0016-7185(92)90019-Z)
- MacGregor, J. N., & Chu, Y. (2011). Human performance on the traveling salesman and related problems: A review. *Journal of Problem Solving*, *3*(2), 1–29. Retrieved from <http://docs.lib.purdue.edu/jps/vol3/iss2/2>
- MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (1999). Spatial and contextual factors in human performance on the travelling salesman problem. *Perception*, *28*(11), 1417–27.
- Philip, J. (2004). The area of a random convex polygon. TRITA-MAT-04-MA-07, June. Retrieved from <http://www.math.kth.se/~johanph/area12.pdf>

Vickers, D., Lee, M. D., Dry, M. J., & Hughes, P. (2003). The roles of the convex hull and the number of potential intersections in performance on visually presented traveling salesperson problems. *Memory & Cognition*, *31*(7), 1094–104. <http://dx.doi.org/10.3758/BF03196130>