

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The role of clustering in the efficient solution of small Traveling Salesperson Problems

Permalink

<https://escholarship.org/uc/item/5q76r06x>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Marupudi, Vijay
Harsch, Rina
Rao, V. N. Vimal
et al.

Publication Date

2021

Peer reviewed

The role of clustering in the efficient solution of small Traveling Salesperson Problems

Vijay Marupudi (vijaymarupudi@gatech.edu)

School of Interactive Computing
Atlanta, GA 30308 USA

Rina Miyata Harsch (harsch@umn.edu)

Department of Educational Psychology,
Minneapolis, MN 55455 USA

V. N. Vimal Rao (rao00013@umn.edu)

Department of Educational Psychology,
Minneapolis, MN 55455 USA

Jeffrey Kramer Bye (jbye@umn.edu)

Department of Educational Psychology,
Minneapolis, MN 55455 USA

Jimin Park (park1481@umn.edu)

Department of Educational Psychology,
Minneapolis, MN 55455 USA

Sashank Varma (varma@gatech.edu)

School of Interactive Computing
Atlanta, GA 30308 USA

Abstract

Human solutions to the Traveling Salesperson Problem (TSP) are surprisingly close to optimal and unexpectedly efficient. We posit that humans solve instances of the TSP by first clustering the points into smaller regions and then solving each cluster as a simpler TSP. Prior research has shown that participants cluster visual stimuli reliably. That is, their clustering and re-clustering of the same stimulus are similar, especially when the stimulus is relatively more clustered. In this study, participants solved the same TSP instances twice. On the second presentation, half of the instances were flipped about the horizontal and vertical axes. Participants solved the TSP reliably, with their two tours of the same instance sharing 77 percent of the same edges on average. In addition, within-participant reliability was higher for more clustered versus more dispersed instances. Our findings are consistent with the proposal that people use clustering strategies to solve the TSP.

Keywords: traveling salesperson problem; problem solving; computational complexity; computational thinking; mathematical cognition

The traveling salesperson problem (TSP) entails finding the shortest path connecting a set of points in a plane, starting at one point, visiting all other points exactly once, and returning to the starting point. Such a solution is called a *tour*. There is no known algorithm for solving the TSP in polynomial time, i.e., time which is a polynomial function of the number of points. The TSP is at least as hard as the NP-complete problems: nondeterministic problems that can be verified in polynomial time. However, it is currently not possible to verify that a solution to the TSP is optimal without solving the problem itself, making the TSP an NP-hard problem. Algorithms for finding the exact solution to an NP-hard

problem have time complexity $O(2^n)$. This makes the search computationally intractable for problems of non-trivial size. For this reason, for such problems, the TSP is typically solved using approximation algorithms that utilize heuristics and constraint relaxation to avoid the exponential growth associated with finding optimal solutions. As a result, they find solutions that are “good enough”, and their performance is evaluated by how close their solutions are to optimal one i.e. how much excess length they have above the short possible tour (Perron & Furnon, 2019). It is therefore surprising that humans are approximately optimal on the TSP. When solving problem instances with between 10 and 120 points, humans are able to find tours that are $\pm 12\%$ of the length of the optimal tour, and are sometimes considerably closer (Dry, Lee, Vickers, & Hughes, 2006). Equally surprising, human solution times appear to scale linearly or nearly linearly with problem size, with a time complexity of either $O(n)$ or $O(n \log(n))$ (Dry et al., 2006). This capability generalizes developmentally: it is present even in 7 year olds (van Rooij, Schactman, Kadlec, & Stege, 2006). It also extends to variants of the TSP in non-Euclidean city-block space (Walwyn & Navarro, 2010). Finally, this ability is not just generative; it is also receptive: When provided a set of points and tours of them, humans judge more optimal tours as having “goodness of figure” (Ormerod & Chronicle, 1999). This is a hint that the surprising efficiency of humans on this computationally hard problem arises in part because of the properties of the visual system.

In reporting the first study of human TSP problem solving,

Krolak, Felts, and Marble (1970) marveled at the efficiency of human performance. However, it was MacGregor and Ormerod (1996) who launched the current study of the mental representations and strategies that humans use to achieve their surprising performance. These researchers noticed that participants were sensitive to the points on the boundary of a problem instance that formed the *convex hull*. The convex hull is the smallest set of points in a figure which define a convex space that contains all other points inside of them. For instances with relatively few points on the convex hull and relatively many points on the interior, participants provided poorer solutions. This led MacGregor and Ormerod to posit the convex hull hypothesis, which states that participants aim to connect the points of the convex hull in order, and that they pick up close-by interior points along way from one convex hull point to another. Instances with larger proportions of interior points introduce more uncertainty into this strategy, resulting in less-optimal performance. The authors presented a computational model of the hypothesis and other aspects of human behavior, such as the sequentiality of TSP problem solving (Macgregor, Ormerod, & Chronicle, 2000). The model fit their data well.

A limitation of these studies is that the instances were designed to conform to the convex hull hypothesis. Carefully varying the number of points on the convex hull against the number of interior points resulted in instances that were roughly circular. In studies with more random instances, the convex hull hypothesis has fared less well. Two later studies (Vickers, Lee, Dry, & Hughes, 2003; Dry & Fontaine, 2014) used random instances and found that participants judged those that contained relatively more interior points *easier* to solve than those with relatively more points on the convex hull. Vickers et al. (2003) also found that participants performed better at TSP problems with relatively more interior points, further challenging the convex hull hypothesis.

Alternative hypotheses have been proposed about the strategies that people use to solve the TSP. Many start with the observation that humans appear sensitive to *hierarchical structure* in spatial memory (McNamara, Hardy, & Hirtle, 1989). This sensitivity may allow them to cluster the points and use a divide-and-conquer approach to simplify the complexity of problem solving, and in this way generate approximately optimal solutions. For example, they might partition the 30 points of a problem instance into 5 groups of roughly 6 points each, solve each cluster as a separate TSP instance, and then connect these local solutions into a solution for the original, overall instance. Graham, Joshi, and Pizlo (2000) proposed such a model inspired by the structure and function of the human visual system. It hierarchically segments an instance into regions and then interactively solves the sub-regions in a top-down fashion. This model was relatively successful in fitting human performance. Other models have used *clustering* techniques to explain human performance, such as Best and Simon (2000) and Kong and Schunn (2007). Clustering approaches can be understood as imposing a relatively shal-

low hierarchy over the points of a problem instance. They have shown greater success in fitting human data than pure convex-hull-based models (Kong & Schunn, 2007).

Some of the aforementioned models use clustering algorithms, such as K-means, and posit that humans might be doing the same. However, relatively little evidence has been offered to show that humans are sensitive to the cluster structure of problem instances and that this structure impacts the solutions of TSP problem instances they generate. An important study in this regard is Dry, Preiss, and Wagemans (2012), who presented participants with the points were relatively more clustered or relatively more dispersed. They found that participants produced more optimal solutions for instances that were more clustered.

Here, we seek more direct evidence for whether participants use a clustering approach to solving the TSP. We built on previous work (Marupudi et al., in preparation) in our lab showing that people have a stable clustering ability. Specifically, in an earlier study, participants were remarkably reliable when generating “clusterings” of the same dot stimulus on two occasions, even when the stimulus is flipped horizontally and vertically on its second presentation, and even when a distractor task intervenes between the two sets of presentations/clusterings. Moreover, the stability of their clusterings was reduced for relatively dispersed instances versus relatively clustered instances (where “dispersed” and “clustered” are defined neutrally, by statistical indices of spatial randomness). If we assume that people adopt a clustering strategy when solving TSP instances, then their TSP performance should show the same patterns as the clustering task. Specifically, we hypothesize that when participants perform the TSP task on the same instance on two different occasions, their solution paths will be more stable for more clustered instances than for more dispersed instances. Alternatively, if participants are *not* clustering, then we would not expect the reliability of TSP solutions to follow the cluster structure of the instance.

Methods

Participants Forty-six undergraduate students at a large public university in the Midwest completed the study. They took approximately 45 minutes (median time) to complete the experiment online and were compensated with a \$15 gift card. The protocol for the study was approved by the local IRB.

Design This study followed a $4 \times 2 \times 2$ within-subjects design. The factors were Number Of Points (10, 15, 20, 25), Cluster Structure (clustered, dispersed), and Orientation Congruency on the second presentation (same, flipped). All factors were varied orthogonally.

Materials We selected instances of different levels of Cluster Structure by randomly generating TSP instances with the desired number of points and filtering out those that did not meet the appropriate standard of “clusteriness” or “dispersion”. We operationally defined these levels using the Z score measure, adapting the variance and edge effects estimates of

Donnelly (1978). The Z score measure is defined as

$$Z = \frac{\bar{d} - E(d_i)}{\sqrt{\text{Var}(\bar{d})}}$$

where \bar{d} is the nearest neighbor distance,

$$\bar{d} = \frac{\sum_{i=1}^N d_i}{N}$$

$E(d_i)$ is the expected value of the nearest neighbor distance for random patterns where A is the area and B is the perimeter,

$$E(d_i) = 0.5\sqrt{\frac{A}{N}} + \left(0.0514 + \frac{0.041}{\sqrt{N}}\right) \frac{B}{N}$$

and $\text{Var}(\bar{d})$ is the variance

$$\text{Var}(\bar{d}) = 0.070 \frac{A}{N^2} + 0.037B \sqrt{\frac{A}{N^5}}$$

We corrected the Z score for instances where the points might appear as a single cluster, which might be biased towards convex hull strategies. This happens when the random stimulus generation process places points near the center of the space and leaves the regions near the border unfilled. Specifically, we calculated the Z score for the bounding box described by the outermost points for calculating measures A and B , instead of the entire visible space (800×500). As a result, the dots of all instances spanned most of the space visible to participants, and were not oddly clustered in a single region.

We then randomly selected from instances with corrected Z score values in the range 1 ± 0.05 , our operationalization for 'clustered' stimuli, and instances in the range -2 ± 0.05 for our operationalization for 'dispersed' stimuli. We chose these specific values because we judged the instances they selected to be *qualitatively* more clustered and more dispersed, respectively, while at the same time keeping the goal of the experiment hidden to participants. This subtlety can be seen in the examples shown in Figure 1.

Number Of Points ranged between 10 - 25 points in increments of 5 points, for a total of 4 levels. For each Number Of Points, we generated 4 clustered and 4 dispersed instances. This resulted in $4 \times (4 + 4) = 32$ instances. Pilot testing confirmed that participants could complete trials without experiencing excessive mental or physical fatigue.

Procedure All participants solved the same 32 TSP instances on the first presentation in a randomized order, then completed an unrelated distractor task lasting approximately 5 minutes, and finally solved the same 32 instances in a different randomized order a second time. On the second presentation, 16 of the 32 instances were flipped horizontally and vertically as specified by the Orientation Congruency factor which varied orthogonally with Number of Points and Cluster Structure. (The distractor task involved estimating the value of various factorial expressions, e.g., $8 \cdot 7 \cdot 6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1$.) Thus, each participant completed 64 TSP trials across the experiment. They did so using an online interface on a custom plugin using the




Figure 1: Examples of clustered (Z score: -2.003) and dispersed (Z score: 1.009) instances.

jsPsych library (De Leeuw, 2015). Participants constructed their tour by clicking on the points in sequence. Their choices were reflected in realtime with a change in the color of the points and the extension of the line indicating their tour to that point. The interface did not permit backtracking; this was to maintain a correspondence with real world movement which often cannot be undone without spending additional resources.

We collected the sequence of points in each tour and the response time for each click. Of the 32 instances in the first presentation, half (16) were statistically clustered and half (16) were statistically dispersed. On the second presentation, half of the instances within each Cluster Structure type were flipped (i.e., 8 clustered, 8 dispersed); the other half were presented in the same orientation.

We defined the reliability between two solutions of a TSP instance as the proportion of edges (connections between points) shared between the two tours. This measure ranges between 0 and 1, with 0 indicating no edges shared between the two tours and 1 indicating all edges shared, i.e., identical tours.

Results

TSP Reliability The primary research questions asked (1) whether participants are stable in their TSP performance, or more precisely, whether they produce similar tours when solving the same problem instance on different occasions, and (2) whether humans first cluster the points in an instance as a precursor to generating a tour. Support for (1) would indicate a stable mechanism underlying TSP performance (i.e. high reliability), and support for (2) would indicate that this mechanism is clustering (i.e. higher reliability for clustered instances). To address these research questions, we calculated the reliability of the two solutions participants produced for each instance. We then fit a series of linear mixed effects models predicting the TSP reliability using Number Of Points, Cluster Structure,

and Orientation Congruency as predictors. The model with all three factors and all interactions as fixed effects, and with a random intercept for each participant, was adopted because it had the lowest AICc value.

There was an effect of Cluster Structure ($\chi^2(1) = 8.3, p < 0.005, \beta = 0.11$). This indicates that clustered instances were solved 0.11 more reliably than dispersed instances. Specifically, there was 11% greater overlap in the tours generated on the two occasions for clustered instances than for dispersed instances. There was also an interaction between Cluster Structure and Number Of Points ($\chi^2(1) = 20.15, p < 0.001, \beta = -0.01$). Each additional point in an instance reduced TSP reliability by 0.01 (i.e., resulted in 1% less overlap in the tours), but only for dispersed instances; for clustered instances, there was no significant effect of Number Of Points on TSP reliability (Figure 2). This indicates that participants might be clustering in order to solve the TSP, and this trend follows a similar performance profile to those we found when participants clustered stimuli, addressing research question (2) that one mechanism for humans solving the TSP is clustering

The model did not include a significant effect of Number Of Points or of Orientation, nor were any interactions involving Orientation significant. Thus, with respect to the main research question (1) concerning the reliability of TSP performance, it was high overall.

Prior studies have investigated the general characteristics of human TSP performance. We attempted to replicate these findings in a series of follow-up analyses that further informs the role of clustering. We focused on findings regarding the near-optimality of human performance. These analyses looked at participants' initial solutions of the 32 TSP problems. They did not include their solutions to these problems on the second occasion to avoid potential confounds of practice effects.

Percent above optimal We first evaluated whether we replicated the most striking finding in the human TSP literature: that people produce tours that are within a few percentage points of optimal, even as the Number Of Points increases. This trend was observed by Graham et al. (2000) and Dry et al. (2006), among other studies. We calculated the optimal solution for each of our 32 problem instances using the Concorde TSP solver, which uses linear programming techniques (Applegate, Bixby, Chvatal, & Cook, 2006). We then used the following formula to calculate the percent above optimal (PAO) for participants' solutions.

$$PAO = 100 \times \frac{\text{participant tour length} - \text{optimal tour length}}{\text{optimal tour length}}$$

Finally, we fit multiple linear mixed effects models predicting the percent above optimal using Number Of Points and Cluster Structure as independent variables. The best-fitting model, measured by the lowest AICc value, included both fixed effects and random effects of Number Of Points and Cluster Structure along with a random intercept term per participant (Relative likelihood = 0.89). The model included a sig-




Figure 2: The effect of Cluster Structure on TSP reliability of instances of varying Number Of Points. Increasing the number of points has a minimal effect on reliability for clustered instances but a deleterious effect for dispersed instances.

nificant effect of Number Of Points ($\chi^2(1) = 58.21, p < 0.001, \beta = 0.33$) and Cluster Structure ($\chi^2(1) = 14.77, p < 0.001, \beta = 1.87$). Each point adds 0.33% extra tour length from optimality. Additionally, dispersed instances were 1.87% less optimal than clustered instances (Figure 3).

Thus, we replicated the surprising finding of approximately optimal human performance on the TSP, a problem that is computationally intractable. Although the deviation from optimality increases with the Number Of Points, it is less than 10% for all trial types except for the dispersed instances with the maximum number (25) of points. A model with an interaction term between Cluster Structure and Number Of Points was eliminated in model selection for having a low AICc value, suggesting that clustered instances provide a constant advantage over dispersed instances to participants across all Number Of Points.

Time taken Another striking finding in the human TSP literature is the linearity of people's solution times with the Number Of Points. We therefore investigated whether the participants in our study replicated the $O(n)$ algorithmic time complexity that previous studies (e.g., Graham et al. (2000)) have shown. We also assessed how Cluster Structure impacts participants' solution times. This question is critical given our hypothesis that people solve TSP instances efficiently by capitalizing on clustering. Model selection using AICc on a series of linear mixed effects models led to the adoption of a model with fixed and random effects of Number Of Points and Cluster Structure, along with the fixed (but not random) effect for the interaction between number of points and Cluster Structure (Figure 4). The model included an effect of Number Of Points ($\chi^2(1) = 300.67, p < 0.001, \beta = 1093$). Each point




Figure 3: Percent above optimal (PAO) as a function of Cluster Structure and Number Of Points. Each point leads to an increase of 0.33% in PAO. Dispersed stimuli were 1.87% less optimal than clustered stimuli.




Figure 4: Time taken to complete the TSP for instances with different Cluster Structures and Numbers Of Points. Each additional point lead to a 1093 ms increase in time taken. There was no effect of Cluster Structure.

was associated with an increase of 1093 milliseconds of problem solving time. The linear trend in the data explained large portion of the variance ($R^2 = 0.33$), replicating previous findings in the literature (e.g., Graham et al. (2000)). Contrary to our expectations, we did not find an effect of Cluster Structure on the time taken to complete the TSP.

To summarize, the current study demonstrates, once again, that people are able to generate approximately optimal solutions to TSP problems in time linear in the Number Of Points, at least for small problems. However, it also demonstrates that people are flexible in achieving this surprising performance, perhaps using alternative strategies when clustering is less available.

Discussion

Prior studies of human solution of the traveling salesperson problem (TSP) have focused on calculating the percent above optimal for people's solutions and the time they take to gen-

erate them. The consistent finding is that people are approximately optimal in the quality of their tours and approximately linear in the time they require to generate them, even though the TSP is a computationally hard problem. These studies have found mixed evidence for potential strategies involved in this surprising performance: the convex hull, crossing avoidance (Van Rooij, Stege, & Schactman, 2003), and clustering hypotheses. Here, we focused on the clustering hypothesis and provided a stronger test than in past studies.

In prior work, we have found that participants are more reliable in their clusterings of a problem instance that is, statistically speaking, more clustered than more dispersed (Marupudi et al., in preparation). In this study, we found that participants produced more reliable TSP solutions for more clustered instances than for more dispersed instances. This similarity in the performance profiles of participants on the clustering and TSP tasks is consistent with the proposal that participants first cluster the points of a TSP instance and then exploit this local structure to generate a tour.

Increasing the Number Of Points did *not* affect the reliability of TSP solutions for clustered instances, but did do so for dispersed instances. This follows directly from our clustering hypothesis: Dispersed instances have less evident cluster structure. Thus, people probably generate less consistent (i.e., more variable) clusterings of these stimuli, and thus less consistent (i.e., more variable) tours. It is important to note that the Clustering Structure of instances did *not* affect the time participants required to solve TSP instances. This replicates the null effect of cluster structure previously documented by Dry et al. (2012). This implies that clustering is *not* a strategy that people use to save time. Rather, they appear to take advantage of the Cluster Structure of an instance when it is evident, and in this way improve the quality of their solution. This leaves room for alternative strategies (i.e., convex hull, crossing avoidance, nearest neighbour) that might be used when the cluster structure of an instance is not evident.

One failed prediction concerned the percent above optimality (PAO) of TSP solutions. We expected to find an interaction effect between Number Of Points and Cluster Structure on this variable. Specifically, we expected that with increasing Number of Points, participants would deviate less from optimality for the clustered instances than for the dispersed instances. This is because the cluster structure of the former would guide them through the larger problems, while the lack of such structure in the latter would provide minimal such affordances. This prediction was not supported, and it is unclear why this was the case. It is possible that with a greater and greater Numbers of Points, the benefit of clustering begins to diminish: as the clusters increase in size, they become difficult TSP problems in their own right. We are currently running a replication of this portion of the study to see if this surprising finding persists. If it does, then this might explain why human performance eventually deteriorates when the Number Of Points grows sufficiently large: the clusters themselves become problems that are too large to solve efficiently.

This study not only successfully replicated effects from previous studies suggesting clustering as a strategy on the TSP, but additionally showed that human solutions are sensitive to the cluster structure of instances, and this can in turn contribute to the reliability of their problem solving. We are currently following up this finding in a new study that asks participants to first cluster a TSP instance and then solve it as a TSP. The prediction is that the TSP tours they generate will follow the structure of their clusterings.

It should be acknowledged that the difference in the optimality with which clustered and dispersed instances are solved, while present, was low (1.87% percentage above optimal difference). The advantage might be larger for “more clustered” instances than the ones used in the present study. We tried to ensure that the clustered and dispersed instances were not noticeably different from each other so that participants would not guess the hypothesis of the experiment. The subtle difference in the clustering versus dispersion of our stimuli was enough to find evidence for the various predictions made here. The use of more clustered instances in a future study might reveal greater evidence for the role of clustering in TSP problem solving than was found here.

A limitation of most models of human TSP performance is that they do not attempt to predict the *exact* tours that participants produce on the TSP. Only one prior study has evaluated how well the tours their model produces correlate with those that humans produce (Kong & Schunn, 2007). Future research looking at the fine-grain structure of tours might yield additional information about whether people apply clustering to TSP instances in order to produce efficient solutions. It might also find evidence for the additional strategies that people employ.

Determining the strategies humans use to solve NP-hard problems like the TSP is important for our broader understanding of the limits of human problem solving, and may in turn inform efforts in artificial intelligence (AI). Approximation algorithms that provide “good enough” solutions to computationally intractable problems may benefit from adopting the strategies, clustering and otherwise, that humans use to efficiently solve (small instances of) these problems. This might result in AI systems that improve upon the current state of the art.

Acknowledgements

We thank **Kori Terkelsen** for coordinating data collection on this study and for helping the transition online when the COVID-19 pandemic arrived.

References

Applegate, D., Bixby, R., Chvatal, V., & Cook, W. (2006). *Concorde TSP solver*.
 Best, B. J., & Simon, H. A. (2000). Simulating human performance on the traveling salesman problem. In *Proceedings of the third international conference on cognitive modeling* (pp. 42–49).

De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods*, 47(1), 1–12.
 Donnelly, K. (1978). Simulation to determine the variance and edge-effect of total nearest neighbour distance. *Simulation methods in archeology*.
 Dry, M., & Fontaine, E. L. (2014). Fast and Efficient Discrimination of Traveling Salesperson Problem Stimulus Difficulty. *The Journal of Problem Solving*, 7(1), 9.
 Dry, M., Lee, M., Vickers, D., & Hughes, P. (2006). Human Performance on Visually Presented Traveling Salesperson Problems with Varying Numbers of Nodes. *The Journal of Problem Solving*, 1(1). Retrieved from <https://docs.lib.purdue.edu/jps/vol1/iss1/4> doi: 10.7771/1932-6246.1004
 Dry, M., Preiss, K., & Wagemans, J. (2012). Clustering, randomness, and regularity: Spatial distributions and human performance on the traveling salesman problem and minimum spanning tree problem. *The Journal of Problem Solving*, 4(1), 2.
 Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & cognition*, 28(7), 1191–1204.
 Kong, X., & Schunn, C. D. (2007). Global vs. local information processing in visual/spatial problem solving: The case of traveling salesman problem. *Cognitive Systems Research*, 8(3), 192–207.
 Krolak, P., Felts, W., & Marble, G. (1970). A man-machine approach toward solving the traveling salesman problem. In *Proceedings of the 7th Design Automation Workshop* (pp. 250–264). Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/800160.805133> doi: 10.1145/800160.805133
 MacGregor, J. N., & Ormerod, T. (1996). Human performance on the traveling salesman problem. *Perception & psychophysics*, 58(4), 527–539.
 Macgregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2000). A model of human performance on the traveling salesman problem. *Memory & Cognition*, 28(7), 1183–1190. Retrieved from <https://doi.org/10.3758/BF03211819> doi: 10.3758/BF03211819
 McNamara, T. P., Hardy, J. K., & Hirtle, S. C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(2), 211–227. doi: 10.1037/0278-7393.15.2.211
 Ormerod, T. C., & Chronicle, E. P. (1999). Global perceptual processing in problem solving: The case of the traveling salesman. *Perception & Psychophysics*, 61(6), 1227–1238.
 Perron, L., & Furnon, V. (2019). *OR-Tools*. Retrieved from <https://developers.google.com/optimization/>
 van Rooij, I., Schactman, A., Kadlec, H., & Stege, U. (2006). Perceptual or Analytical processing? Evidence from children’s and adult’s performance on the Euclidean traveling salesman problem. *The Journal of Problem Solving*, 1(1),

6.

- Van Rooij, I., Stege, U., & Schactman, A. (2003). Convex hull and tour crossings in the Euclidean traveling salesperson problem: Implications for human performance studies. *Memory & cognition*, *31*(2), 215–220.
- Vickers, D., Lee, M. D., Dry, M., & 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), 1094–1104. Retrieved from <https://doi.org/10.3758/BF03196130>
doi: 10.3758/BF03196130
- Walwyn, A. L., & Navarro, D. J. (2010). Minimal paths in the city block: Human performance on Euclidean and non-Euclidean traveling salesperson problems. *The Journal of Problem Solving*, *3*(1), 5.