

Human Performance on Visually Presented Traveling Salesperson Problems with Varying Numbers of Nodes

*Matthew Dry, Michael D. Lee, Douglas Vickers, and Peter Hughes**

Abstract

We investigated the properties of the distribution of human solution times for Traveling Salesperson Problems (TSPs) with increasing numbers of nodes. New experimental data are presented that measure solution times for carefully chosen representative problems with 10, 20, . . . 120 nodes. We compared the solution times predicted by the convex hull procedure proposed by MacGregor and Ormerod (1996), the hierarchical approach of Graham, Joshi, and Pizlo (2000), and by five algorithms drawn from the artificial intelligence and operations research literature. The most likely polynomial model for describing the relationship between mean solution time and the size of a TSP is linear or near-linear over the range of problem sizes tested, supporting the earlier finding of Graham et al. (2000). We argue the properties of the solution time distributions place strong constraints on the development of detailed models of human performance for TSPs, and provide some evaluation of previously proposed models in light of our findings.

Introduction

As humans interact with the world, they are inundated with information. Deciding what information is or is not relevant to a given task is a nontrivial problem: A nearly infinite supply of information is available to a decision-maker, all of which may potentially have *some* influence upon an outcome. In spite of this potential information overload, we are able to negotiate challenging environments with remarkable speed and success. It appears that the way we deal with this wealth of data is to ignore much of it by employing heuristics that focus upon statistical and environmental regularities (e.g., Gigerenzer & Todd, 1999). In order to understand these heuristic processes better, researchers have begun to explore tasks, such as the Traveling Salesperson Problem (TSP), that belong to a class of problems for which it is believed no method exists whereby an optimal solution can be calculated within a practical time (Lawler, Lenstra, Rinooy Kan, & Schmoys, 1985). Because these problems are simple to explain and easily understood, they are well suited

*School of Psychology, University of Adelaide, Australia

Correspondence concerning this article should be addressed to: Matthew Dry, School of Psychology, University of Adelaide, SA, 5005 Australia, Telephone: +61 8 8303 3101, Facsimile: +61 8 8303 3770, Email: matthew.dry@psychology.adelaide.edu.au

to laboratory exploration, yet they have some claim to being representative of the types of real-world problems encountered in everyday life.

Solving a TSP involves finding the shortest pathway through an array of n nodes or "cities," generally returning to the starting node. As the number of nodes in the array is increased, the number of possible solutions increases factorially: Each TSP has $(n - 1)!/2$ possible solutions. This means that a 5-point array has 12 possible solutions, a 10-point array has 181,440 possible solutions, and a 15-point array has 4×10^{10} possible solutions. Despite the apparent intractability of the TSP, research into human performance upon visually presented TSPs has indicated that participants are capable of solving the problems to near optimal accuracy with a minimum of cognitive effort (e.g., MacGregor & Ormerod, 1996; Vickers, Butavicius, Lee, & Medvedev, 2001).

As part of this general program of research, Graham, Joshi, and Pizlo (2000) carried out an experiment in which six participants were tested on instances of Traveling Salesperson Problems with 6, 8, 10, 20, or 30 nodes. An important finding was that participants appeared to spend a roughly constant time per node, implying that the total time they required to arrive at a solution was a linear function of the number of nodes.

If this finding can be corroborated, it is important for two main reasons. First, this experiment is the only one in which human solution times have been measured and compared across instances with several different numbers of nodes. Secondly, there are many algorithms that yield approximate solutions to TSP instances (Lawler et al., 1985). However, no known useful algorithm predicts a simple linear relationship between solution time and the number of nodes, n . If solution time for participants is indeed a linear function of n , then this suggests that the underlying mental processes have a low computational complexity and that participants cannot be searching through the entire (or even a substantial proportion of the) problem space of potential TSP solutions. This, in turn, imposes strong constraints on modeling human performance in arriving at TSP solutions.

Unfortunately, several features of the experiment by Graham et al. (2000) make for considerable uncertainty regarding the finding of a linear relation between solution time and n . First, the range in the number of nodes studied by Graham et al. was quite restricted. Second, no descriptive statistics are given for the problem instances (e.g., the number of convex hull nodes,¹ the number of potential intersections), so that it is not possible to ascertain how representative they were of the population of all problem instances with a given n . Vickers, Lee, Dry, and Hughes (2003) demonstrated that manipulating the number of potential intersections and number of nodes on the convex hull for problems of a set size had a significant effect upon problem difficulty. Furthermore it has been shown that number of potential intersections and number of convex hull nodes for randomly generated problems varies following a near-normal distribution (Lee & Vickers, 2000; Vickers et al., 2003). Given this, it is important to ensure that test stimuli are representative and not drawn from the tails of these distributions. Thirdly, data were collected for six participants only (including the authors).

Fourthly, participants had differing amounts of practice. Fifthly, the experiment was conducted in two stages (the first with 6, 8, or 10 nodes and the second with 10, 20, or 30 nodes), and the two stages involved different amounts of practice. Finally, the three authors responded under a deadline procedure, whereas the three remaining participants were instructed to use “as much time as needed in order to produce a best solution” (Graham et al., 2000, p. 1194). Under these conditions, the averaged times for each problem size can be interpreted as suggestive only.

A more general problem that we also attempt to address is that using total TSP solution time as an index of decision-making processes ignores the time that participants spend physically executing a solution. In order to provide an estimate of the separate cognitive and motor aspects of TSP solution time, we undertake an additional experiment measuring the time taken to trace the paths of representative participant TSP solutions.² We assume that, by removing the need for participants to devise a solution, we could obtain a measure of the pure motor aspects of TSP solution times. Similarly, by subtracting the solution times for the path-following experiment from the total TSP solution times, we obtain an estimate of the decision-making aspects of TSP solution time. In order to differentiate between the two experiments, the main experiment will be referred to as the TSP condition, and the additional experiment will be referred to as the path-following condition.

Overall, therefore, in this paper we carry out a rigorous, systematic investigation of the relation between TSP solution time and n , using representative stimulus arrays with between 10 and 120 nodes. We compare the results against hypotheses concerning the solution processes employed by participants. Using Bayesian model selection we compare the solution times predicted by the convex hull procedure proposed by MacGregor and Ormerod (1996), by the hierarchical approach of Graham et al. (2000), and by five algorithms drawn from the artificial intelligence and operations research literature.

Experiment

Participants

Forty participants (29 females, 11 males) completed the TSP condition. Five participants (2 females, 3 males) completed the path-following condition. The participants' ages ranged from 18 to 52 years, with a mean age of 23. Ethics approval was obtained from the University of Adelaide Psychology Department human ethics subcommittee.

Stimuli


All stimuli were presented on a uniform white background within a 14 cm $\times$ 14 cm square. The nodes were black dots of diameter 1.5 mm, and were constrained to be at least one diameter separate from each other.

For the TSP condition, a total of 10,000 random problems were generated for problems with 10, 20, 30, 40, . . . 120 nodes. Each of these problems was generated by choosing each point independently from a uniform distribution, subject to a constraint that ensured a minimal level of separation between nodes. For these large samples, the modal number of points on the convex hull was observed to be 6, 8, 9, 10, 10, 10, 11, 11, 11, 12, 12, and 12 or 13 for $n = 10, 20, 30, 40, \dots 120$ nodes. Over all the problems with the modal number of convex hull nodes, the average proportion of intersections³ between all nodes (including hull points) was 0.2317. The proportion of intersections for problems with the modal number of hull points in each category varied between 0.2307 and 0.2327. Accordingly, for each value of n , a total of 20 problems were chosen for the experiment that had the modal number of convex hull nodes and were in this range for the proportion of intersections. No restriction was placed on the internode distance mean or standard deviation. An example of a typical 40-node TSP is given in Figure 1a.

For the path-following condition the stimuli were drawn from the pool of participant solutions to problems in the TSP condition. For each value of n , we selected the participant solution that was the closest to the mean of all participant solution lengths for that class. The solution was superimposed on its corresponding TSP. Figure 1b provides an example of the type of stimuli employed in the path-following experiment.

In order to acquaint the participants with the process of solving a TSP, they were given three practice problems (containing 30, 60, and 90 nodes) to complete before commencing the experiment.

Figure 1.
Example 40-node TSP (a) and path-following (b) arrays.



Procedure

The arrays were presented on computer color monitors. The method employed to create pathways was identical for both experiments. Participants could begin a path by left-clicking on a node with the computer mouse. Then, while holding down the mouse button, they drew a path by dragging the mouse cursor to a subsequent node and releasing the button, causing a straight line to be drawn between that node and the previously visited node. By right-clicking on a link to select it and then pressing the delete key on the keyboard, the participants could undo any links they had drawn. The participants were thus free to connect the nodes in any order, to work alternately from two nodes, or to work on several separated clusters of nodes.


In the TSP condition, the participants were instructed to create the shortest continuous pathway that passed through every point in the array. No restrictions were imposed on the length of time that the participants could take to complete a problem. Participants received feedback on how far their solution was from the optimal solution, expressed as a percentage. This feedback was based on estimated optimal solutions that were calculated using a standard simulated annealing algorithm (Press, 1992; Reinelt, 1994). Conversely, for the path-following condition the participants were instructed to trace the circuit displayed on the screen as quickly as possible, and no feedback was provided. In both cases, participants could proceed to the next problem by clicking on a "proceed to next test" button but were not allowed to proceed until a legitimate closed TSP tour was in place.

Each participant completed one problem for each of the 12 numbers of points. The assignment of particular arrays to participants in the TSP condition ensured that each of the 20 arrays of each type was completed by exactly two participants.

Analysis

Using the estimated optimal solution as a benchmark, it was possible to measure the mean participant deviation from optimality for each problem in the TSP condition, expressed as the proportional length by which a participant solution exceeded the optimal solution. [Figure 2a](#) compares the mean participant solution lengths with the estimated optimal solution length, and [Figure 2b](#) shows the mean participant deviation from optimality. Participant solution lengths closely approximate the estimated optimal solution lengths, with deviation asymptoting at around 0.11 when problem size exceeds 70 nodes. Importantly, [Figure 2b](#) demonstrates that participant performance falls within the range of deviation from optimality scores reported in Graham et al. (2000) and MacGregor, Ormerod, and Chronicle (1999).

Figure 2.
The distribution of mean participant solution lengths (a), and mean participant deviation from optimal (b) for TSP arrays with 10, 20, . . . 120 nodes. The error bars represent one standard error. The data points from Graham et al. (2000) represent mean values aggregated across the separate scores reported for Authors and Naive Subjects.




Mean empirical solution times for problems in the TSP and path-following conditions were obtained by averaging across all participants for each of the 12 problem sizes. As stated previously, by subtracting each point in the distribution of path-following solution times from its corresponding point in the distribution of TSP solution times, we were able to obtain an estimate of the decision-making component of TSP solution times. Figure 3 shows the distributions of total solution time and decision time for the TSP condition, as well as the distribution of solution times for the path-following condition. As can be seen, all three data sets appear to follow a linear distribution.

Graham et al. concluded that the linear distribution of TSP solution times ruled out the majority of artificial intelligence and operations research algorithms as plausible models of human performance on the TSP task. However, they did not provide any statistical evidence to support this conclusion. Given the empirical TSP solution time data and the computational complexity of an algorithm (as indexed by the number of operations necessary to provide a solution to a given problem), it is possible to estimate algorithm solution time as a function of number of problem nodes. The linear appearance of the two empirical TSP solution time distributions suggests that participants employ heuristics with a complexity of $O(n)$. Graham et al. (2000) reviewed a number of artificial

intelligence and operations research algorithms capable of providing good approximate solutions to the TSP: nearest neighbor (NN), cheapest insertion (CI), convex hull with cheapest insertion (CCI), convex hull with cheapest insertion and three segment swap (CCI3), and elastic net (EN), with computational complexities of $O(n^3)$, $O(n^3 \ln n)$, $O(n^2 \ln n)$, $O(n^3)$, and $O(n^2)$ respectively. Furthermore, they introduced the hierarchical pyramid algorithm (HP) with complexity $O(n^2)$ and suggested that the HP algorithm could be simplified to a complexity of $O(n \ln n)$. The computational complexities of the NN, CI, CCI, and CCI3 algorithms are based on the estimates presented in Golden, Bodin, Doyle, and Stewart (1980). As Golden et al. indicate, the complexity of the NN and CI algorithms increases by an order of magnitude equal to n [from $O(n^2)$ to $O(n^3)$, and $O(n^2 \ln n)$ to $O(n^3 \ln n)$ for NN and CI respectively] when every node in an array is considered as a potential starting point for the algorithm.

Figure 3.
The distribution of mean participant solution times for TSP and path-following arrays with 10, 20, ... 120 nodes. The error bars represent one standard error.



We employed Bayesian statistical inference to determine which of the six functions predicted by computational complexity could provide the best account of the empirical TSP total solution time and decision time data. Using this approach, the six functions become competing models that make different assumptions about the relationship

between mean solution time and number of problem nodes. The six models, their computational complexity, functional form, slope, and intercept are summarized in Table 1.

Table 1
Summary of Models, Including the Slope, Intercept, Weighted Sum Squared Error Data Fit, BIC and Bayes Factors for the TSP Total Solution Time, and Decision Time Data.

Model Summary			Total Solution Time					Decision Time				
No	Complexity	Functional Form	slope	int.	WSSE	BIC	Bayes Factors	slope	int.	WSSE	BIC	Bayes Factors
1	$O(n)$	$y = ax + b$	1.712	6.29	6.40	11.37	1	1.07	2.18	8.35	13.32	2.42
2	$O(n \ln n)$	$y = ax \ln x + b$	0.357	17.41	6.66	11.63	1.13	0.226	9.16	6.60	11.57	1
3	$O(n^2)$	$y = ax^2 + b$	0.015	31.46	48.44	53.41	1.34×10^9	0.009	18.25	22.20	27.17	2.43×10^3
4	$O(n^3)$	$y = ax^3 + b$	0.000	38.78	124.40	129.37	4.20×10^{25}	0.000	23.13	56.60	61.57	7.18×10^{10}
5	$O(n^2 \ln n)$	$y = ax^2 \ln x + b$	0.003	33.98	65.76	70.73	7.75×10^{12}	0.002	19.92	29.76	34.73	1.06×10^5
6	$O(n^3 \ln n)$	$y = ax^3 \ln x + b$	0.000	39.75	139.90	144.87	9.74×10^{28}	0.000	23.78	63.76	68.73	2.58×10^{12}

Note: Bayes factors are taken in relation to the most likely model, which in this case is Model 1 for the total solution time data and Model 2 for the decision time data.

Figure 4 shows the maximum likelihood fits of the six models to the total TSP solution time data (a) and TSP decision time data (b). In both cases the maximum likelihood fits are based upon the assumption of a Gaussian likelihood function. Given these data fits, it is possible to calculate the Bayesian information criterion for each model (Schwarz, 1978) using:


$$BIC = WSSE + P\ln(N),$$

where WSSE is the weighted sum-squared error, P is equal to the number of model parameters, and N represents the number of data points to which the model is being fitted (in this case, 12). The relative likelihood of each model can then be determined by calculating Bayes factors (Kass & Raftery, 1995). In this case, the six models had equal parametric complexity ($P = 2$); therefore, $P\ln(N)$ was essentially an additive constant of 4.96. It is recognized that log-likelihood ratios based on the maximum likelihood fits will predict the same order of results as Bayes factors. However, given that future studies may wish to evaluate the fits of models with varying parametric complexity, we have presented our results in a form that will facilitate this comparison. Table 1 summarizes the results of these analyses, showing the maximum likelihood fit of the function predicted by each of the six models and the associated BIC and Bayes factors for both the total solution time and decision time.

As can be seen in Figure 4, Models 1 and 2 appear to provide the best qualitative fit to both the total solution time and decision time data, with the more computationally

complex models (Models 3–6) failing to match the basic pattern of results. This finding is supported statistically by the weighted sum squared error fits and Bayesian analysis. The Bayes factors indicate that, for the total solution time data, Model 1 provides the most likely fit to the empirical data, with Model 2 being only 1.13 times less likely. For the decision time data, the order is reversed, with Model 2 being 2.4 times more likely than Model 1. According to Jeffreys’ (1961) guidelines for interpreting Bayes factors, it is difficult to make a meaningful distinction between the likelihood of these two models’ fitting the data (a difference of less than 3.2 is “not worth more than a bare mention”). However, it should be noted that the model that provides the next best fit to the data (Model 3 in both cases) is 1.34×10^9 or 2.42×10^3 times less likely than the best-fitting model for the total solution time and decision time data respectively. On the basis of these analyses, we conclude that human solution times either are linear or closely approximate linearity over the range of problem sizes tested in this experiment.

Figure 4. Comparison between total solution time (a) and decision time (b), and time plotted as a function of complexity. The error bars represent one standard error.



Comparison of the model predicted intercept values provides further evidence in favor of Models 1 and 2. Models 3–6 all predict intercepts that exceed the mean participant solution time for 10 node problems. In contrast to this, Models 1 and 2 both have intercepts that can be plausibly accepted as the amount of preparation time that must necessarily precede the physical execution of a solution.

The same pattern of results was found in a subsequent experiment that employed a subset of the TSP condition stimuli and 70 naive participants. Three problems were chosen from each size category, and each problem was completed a minimum of 22 times (max = 26). The distribution of average participant solution times closely replicated the findings of the initial experiment for both total solution time and decision time data, with Model 2 providing the most likely description of the relationship between problem size and time taken and Model 1 the next most likely description.

Discussion

Our finding—that the most likely polynomial model for describing the relationship between mean solution time and the size of a TSP problem is linear or near-linear over the range of problem sizes tested in this paper—supports the earlier finding of Graham et al. (2000). Our analyses provide a quantified demonstration of the inadequacy of standard artificial intelligence and operations research algorithms to provide a plausible account of human performance. Although previous studies have shown that some of these algorithms are able to perform at a level equal to or better than human performance in regard to solution optimality (Graham et al., 2000) and solution path length (MacGregor, Ormerod, & Chronicle, 2000), in each case the price of good performance appears to be high computational complexity.

More generally, the solution time data suggest that human participants are only searching through a subset of potential solution pathways. This imposes constraints upon the types of heuristics that participants might be employing to solve the TSP task. Recently Vickers, Lee et al. (2003) suggested that the visual system may be solving the TSP task by exploiting spontaneous perceptual organization based on least-distance links between neighboring nodes. The clusters formed by this nearest-neighbor linking can be thought of as providing the initial solution for a TSP, and the TSP task as a hierarchical nearest-neighbor (HNN) procedure of finding the most economical pathway connecting the clusters. A growing body of evidence suggests that such an approach might be able to provide a plausible account of human performance upon the TSP task. First, Vickers, Lee et al. (2003) demonstrated that problem difficulty is related to number of potential intersections. Because nearest-neighbor links never intersect, reliance upon nearest neighbors would place a strong restriction on the number of plausible solution pathways within a problem. Furthermore, this would account for the general avoidance of crossings in TSP solutions (van Rooij, Stege, & Schactman, 2003).

Second, Vickers, Bovet, Lee, & Hughes (2003) found that the solution length of both open and closed versions of the TSP was reliably correlated with the number of nearest-neighbor links that a standard NN algorithm shared with the benchmark solution, and Vickers, Mayo, Heitmann, Lee, and Hughes (2004) demonstrated that solution lengths for TSPs and two closely related visual optimization tasks, the Minimum Spanning Tree and Steiner Tree, were all highly correlated with a measure of path complexity indexing the 1st to k th order neighbor links within participant solutions. Finally, and most importantly in light of the results of the current experiment, the number of clusters formed by nearest-neighbor links increases as a linear function of n (Vickers, Bovet et al., 2003). Given this, we could expect the amount of time taken by an HNN procedure to produce TSP solutions to similarly increase as a linear function of n .

A final point needs to be made in relation to modeling human performance upon visual optimization problems such as the TSP. There have been a number of attempts to elucidate the psychological processes underlying TSP solution. A common feature that each of these approaches shares is that the low-level perceptual structure from which solutions are assumed to be derived (e.g., neighbor relations [Vickers, Lee et al., 2003], clustering derived from Gaussian filtering [Graham et al., 2000], and the convex hull [MacGregor et al., 2000]), have all been shown to be features of numerous other perceptual tasks. For example, both neighbor relations and Gaussian filtering have been used to model the perception of Glass patterns (Caelli, 1981; Dakin, 1997; Dry, Vickers, Lee, & Hughes, submitted), and it could be argued that perception of the convex hull is a form of figure-ground segregation. Importantly, it has been shown that the human visual system is capable of detecting perceptual features such as these with presentation durations of less than a second, and it is assumed that awareness of these features is the result of spontaneous parallel processes acting in tandem across the face of a stimulus.

In this context, it is worth contrasting these “bottom-up” *perceptual* processes with the “top-down” *cognitive* process of tour planning and execution. It may be that we need to make a distinction between the spontaneous parallel processes involved in the initial perception of the structure in TSP arrays and the serial process of linking individual nodes or clusters of nodes. It could be argued that the computational complexity of a *psychological* model should not be based upon the process or processes involved in deriving the initial representation of a problem array but rather upon the process or processes involved in moving from the initial representation to the solution end-state. This is not to say, however, that these perceptual processes are not worthy of attention, for we can assume that the representations derived from these distinctive perceptual processes will result in very different predictions about the computational complexity of the solution process. In any case, both the top-down and bottom-up processes should be of interest to optimization researchers.

Notes

1. The convex hull is a boundary so that no line joining any two nodes in the array can fall outside it.
2. We are indebted to our reviewers for suggesting that we obtain a separate measure of the motor component of TSP solution time.
3. Within each array, there are $n(n-1)/2$ unique node pairings. The number of intersections within an array is given by the number of instances in which the vectors joining node pairs cross. The proportion of intersections can then be calculated as the number of intersections observed within a given array of size n , divided by the total number of potential intersections for all size n arrays, where the number of potential intersections is given by $[n(n-1)/2][(n-2)(n-3)/2]/2$.

References

- Caelli, T. M. (1981). Some psychophysical determinants of discrete moire patterns. *Biological Cybernetics*, 39, 97–103.
- Dakin, S. C. (1997). The detection of structure in Glass patterns: Psychophysics and computational models. *Vision Research*, 37(16), 2227–2246.
- Dry, M., Vickers, D., Lee, M., D., & Hughes, P. (submitted). The role of nearest neighbors in the perception of Glass patterns.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Golden, B., Bodin, L., Doyle, T., & Stewart, J. (1980). Approximate Traveling Salesman algorithms. *Operations Research*, 28(3), 694–711.
- Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The Traveling Salesman Problem: A hierarchical model. *Memory & Cognition*, 28(7), 1191–1204.
- Jeffreys, H. (1961). *Theory of probability*. Oxford, UK: Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–796.
- Lawler, E. L., Lenstra, J. K., Rinooy Kan, A. H. G., & Schmoys, D. B. (1985). *The Traveling Salesman Problem: A guided tour of combinatorial optimisation*. Chichester, UK: Wiley.
- Lee, M. D., & Vickers, D. (2000). The importance of the convex hull for human performance on the Traveling Salesman Problem: A comment on MacGregor and Ormerod (1996). *Perception & Psychophysics*, 62(1), 226–228.
- 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. (1999). Spatial and contextual factors in human performance on the Travelling Salesperson Problem. *Perception*, 28(11), 1417–1427.
- MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2000). A model of human performance on the Traveling Salesperson Problem. *Memory & Cognition*, 28(7), 1183–1190.
- Press, W. H. (1992). *Numerical recipes in FORTRAN: The art of scientific computing* (2nd ed.). Cambridge: Cambridge University Press.

- Reinelt, G. (1994). The travelling salesman: Computational solutions for TSP applications. In *Lecture notes in computer science* (Vol. 840). Berlin: Springer-Verlag.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461–464.
- 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., 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(7), 871–886.
- Vickers, D., Butavicius, M., Lee, M., & Medvedev, A. (2001). Human performance on visually presented Traveling Salesman Problems. *Psychological Research-Psychologische Forschung*, 65(1), 34–45.
- 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.
- 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.