Perception, 2003, volume 32, pages 871 – 886

DOI:10.1068/p3416

The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean travelling salesperson problems

Douglas Vickers, Pierre Bovet¶, Michael D Lee, Peter Hughes
Department of Psychology, University of Adelaide, South Australia 5005, Australia; ¶ Faculty of Psychology and Educational Sciences, University of Geneva, 40 boulevard du Pont-d'Arve, CH 1211 Geneva, Switzerland.; e-mail: Douglas.Vickers@adelaide.edu.au;
website: http://www.psychology.adelaide.edu.au/members/staff/dougvickers/doug.html
Received 1 July 2002, in revised form 27 February 2003

Abstract. The planar Euclidean version of the travelling salesperson problem (TSP) requires finding a tour of minimal length through a two-dimensional set of nodes. Despite the computational intractability of the TSP, people can produce rapid, near-optimal solutions to visually presented versions of such problems. To explain this, MacGregor et al (1999, *Perception* **28** 1417 – 1428) have suggested that people use a global-to-local process, based on a perceptual tendency to organise stimuli into convex figures. We review the evidence for this idea and propose an alternative, local-to-global hypothesis, based on the detection of least distances between the nodes in an array. We present the results of an experiment in which we examined the relationships between three objective measures and performance measures of optimality and response uncertainty in tasks requiring participants to construct a closed tour or an open path. The data are not well accounted for by a process based on the convex hull. In contrast, results are generally consistent with a locally focused process based initially on the detection of nearest-neighbour clusters. Individual differences are interpreted in terms of a hierarchical process of constructing solutions, and the findings are related to a more general analysis of the role of nearest neighbours in the perception of structure and motion.

1 Introduction

An important class of real-world questions that have been neglected in both experimental and differential psychology have to do with problems of combinatorial optimisation, for which there is no known algorithm that can be guaranteed to produce a definitive solution within a practicable time. Such problems, which are ubiquitous in the fields of mathematics and computer science, can often be stated simply and readily understood. A typical example is the Euclidean version of the travelling salesperson problem (TSP), which may be formulated as follows: Given a set of n interconnected towns, with a specified distance between each pair, select a route that (a) visits each town once and once only; (b) returns to the starting point; and (c) keeps the total distance covered to a minimum. To arrive at a solution by comparing all possible routes involves considering $(n-1)!/2$ itineraries. This is feasible when n is a small number, such as 5. However, when n is even moderately large, the combinatorial explosion of possibilities means that it is no longer possible to arrive at a solution through an exhaustive examination of all the alternatives. For example, when n is 25, the number of possible itineraries is so great that a computer evaluating a million possibilities per second would need almost 10 billion years to evaluate them all (Stein 1989).

More generally, the TSP is a member of the class of so-called NP-complete problems, for which it is believed there is no algorithm that can be guaranteed to arrive at an optimal solution within a feasible, polynomial time. That is, no algorithm can solve the problem in a time proportional to n^c, or better, where n is the number of relevant input variables and c is some constant (Lawler et al 1985; Wilf 1986). Exponential algorithms may exist, capable of solving the problem in a time proportional to c^n. However, these algorithms are infeasible for all but the smallest values of n. Such problems arise in thousands of naturally occurring situations that are well known and of practical interest. Although definitive solutions cannot be guaranteed by strict algorithmic procedures, near-optimal solutions can often be obtained by a variety of heuristic methods (Reinelt 1994).

2 Evidence for the influence of global-to-local perceptual organising processes

Despite the computational intractability of the TSP, there has recently been a growing interest in human performance on such problems (Graham et al 2000; Lee and Vickers 2000; MacGregor and Ormerod 1996; MacGregor et al 1999, 2000; Vickers et al 2001; Vickers et al, in press, b). One striking result, common to all studies since Polivanova (1974), is that human performance on visually presented versions of small-scale TSP tasks is comparable to that of the most successful computational procedures. On this basis, MacGregor and Ormerod (1996) have concluded that the solution processes employed correspond to natural, uniformly present, organising processes of the visual system.

This conclusion is consistent with a variety of other findings concerning the spontaneously perceived or aesthetically constructed organisation of the nodes in TSP arrays, as well as the judged goodness of alternative solutions, viewed as abstract patterns (Polivanova 1974; Pomerantz 1981; Ormerod and Chronicle 1999; Vickers et al 2001). In particular, MacGregor and Ormerod (1996) have proposed that participants use the convex hull bounding an array as a first step in constructing a solution and that this use is linked to a preference for organising stimulus elements into convex figures, as suggested by Gestalt analyses of figure – ground differentiation (Palmer 1999, pages 282 – 283). On this basis, MacGregor et al (2000) have developed a computational model that uses the convex hull to establish a sub-tour, into which the remaining nodes are then serially inserted.

At first sight, there appears to be a significant amount of evidence consistent with such a proposal (MacGregor and Ormerod 1996; MacGregor et al 1999, 2000). However, this interpretation of the evidence has been questioned on a number of grounds. First, human performance on more difficult, randomly constructed arrays does not appear to be achieved by a uniform, species-constant perceptual process. Instead, there are reliable individual differences, that are consistent over different problems, are correlated with performance on other optimisation problems (for which the convex hull has no relevance), and that are correlated with general intelligence (Vickers et al 2001; Vickers et al, in press, a).

Second, as argued by Lee and Vickers (2000), the problem arrays employed by MacGregor and Ormerod (1996) were constructed in a highly constrained way that, as the authors themselves acknowledged, is very likely to have encouraged participants to make use of the convex hull.

Third, although variations in a number of response measures have been cited as evidence in favour of the convex-hull hypothesis (including an information-theoretic measure of path uncertainty and standardised measures of the extent to which participants' solutions exceed known optimal values), different measures are cited on different occasions, and the effects of separate independent variables (such as number of interior nodes and the total number of nodes) have been confounded (MacGregor and Ormerod 2000; Vickers et al, in press, b).

Fourth, evidence that participants connect the nodes of the convex hull in order of adjacency (while passing through interior nodes to do so) can at most count as extremely weak corroboration of the hypothesis (Robert and Pashler 2000), since any procedure that achieves optimal or near-optimal solutions will inevitably do likewise.

Fifth, as first pointed out by Flood (1956), because any solution must avoid intersections to be optimal, and because any tour that fails to connect the nodes of the convex hull in order of adjacency must give rise to an intersection, any tour that follows the rule of avoiding intersections must of necessity visit the nodes of the convex hull in order (Quintas and Supnik 1965).

3 Evidence for the influence of local-to-global perceptual organising processes

The above restrictions are severe in that they nullify much (if not all) of the evidence cited in support of global-to-local processing through the application of a heuristic procedure based on the convex hull. This suggests that it may be useful to consider an opposite, local-to-global, account. For example, the standard nearest-neighbour (SNN) algorithm solves TSPs by starting at a node, selected at random, and sequentially linking the next-nearest neighbour to the last selected node (Golden et al 1980). Despite its simplicity, the SNN algorithm has some attractive features as a model of human TSP performance (MacGregor et al 2000), such as mimicking the serial, link-by-link human production of solutions. Unfortunately, human performance on problems with more than 20 nodes is superior to that of the SNN algorithm, which means that the algorithm is at best incomplete as an account of TSP performance (Graham et al 2000).

In a recent paper, Vickers et al (in press, b) have suggested an alternative, locally focused approach that has much in common with the SNN algorithm. According to this approach, the visual system exploits spontaneous perceptual organisation based on the identification of least-distance links between nearest neighbours. If such least-distance links provide an initial, suggested structure for a TSP solution, then the TSP task can be viewed as a hierarchical nearest-neighbour (HNN) procedure of connecting up the nearest-neighbour clusters in as economical a way as possible.

There are many ways in which such a procedure might be realised. However, irrespective of its precise implementation, such an approach has certain desirable characteristics. First, because the least distances between the nodes in a random array never cross, this hypothesis would account for the general avoidance of intersections in TSP solutions (van Rooij et al 2003).

Second, because nearest-neighbour links may be assumed to be detected by a parallel process, and because the number of nearest-neighbour clusters is a linear function of n, the time taken by an HNN procedure would be expected to be a linear function of n, as found by Graham et al (2000).

Third, there is some evidence that nearest neighbours play an important role in the perception of structure and motion in random-dot patterns generally. For example, Vickers et al (submitted, c) found that nearest neighbours, and minimum spanning trees (MSTs) incorporating them, accounted well for the representation of random-dot patterns (constellations in a desktop planetarium). [An MST is a structure, containing no loops and linking all n nodes of a pattern with $n-1$ edges, for which the sum of the edge lengths is a minimum (Johnsonbaugh 2001)]. Nearest-neighbour statistics also predicted mean links and cluster lengths, as well as the number of links and clusters, detected by participants in random-dot patterns with varying numbers of dots. In addition, these authors found that the extent of motion seen in random-dot kinematograms varied inversely with dot density and was related to the mean distance between nearest neighbours.

Fourth, as pointed out by Graham et al (2000), the applicability of the convex hull is limited to closed TSP tasks and does not extend, for example, to open problems of finding a shortest path between a starting node and a (different) finishing node. In contrast, nearest-neighbour relations have potential applicability both to open and to closed TSP tasks, as well as to a variety of other optimisation problems. For example, Vickers et al (in press, a) examined the performance of a group of participants on TSP, MST, and Steiner tree (ST) tasks. [An ST also connects all n nodes of an array, with a minimal sum of edge lengths, but allows the introduction of additional nodes

(Hwang et al 1992)]. These authors calculated a measure of path complexity for each participant's solution to each problem. For each solution link, numbers 1 to n were assigned, according to whether the nodes of that link were connected to the nearest (1), second nearest (2), or n^{th} nearest node, and the sum of these numbers was divided by the number of links to give a measure of path complexity. The lower the path complexity, the more the participant made use of nearest-neighbour links in arriving at a solution. Vickers et al (in press, a) found that measures of path complexity correlated strongly with performance (percentage above optimal) on each type of problem (Pearson's $r = 0.99$, 0.98, and 0.88 for the TSP, MST, and ST, respectively; $N = 48$, in each case). Thus, the spontaneous perception of structure, based on distances between nearest neighbours, appears to play a significant role in a variety of tasks concerned with finding minimal structures.

In short, the HNN approach is consistent with other arguments, mentioned previously, that the task of solving TSPs coincides with natural organising tendencies of the visual system (eg those of Polivanova 1974; MacGregor and Ormerod 1996; Ormerod and Chronicle 1999; and Vickers et al 2001). However, it implies that the explanation of both the perception of organisation and of the initial stages in the solution of TSP tasks is to be found in local-to-global processing, rather than the converse.

4 Experiment

Although the restrictions summarised earlier do not contradict the convex-hull hypothesis, they do imply that it needs to be subjected to more stringent testing. In addition, the possibility that performance is mediated by an opposite type of process seems worth pursuing. MacGregor and Ormerod (1996) themselves suggest that one way in which their hypothesis might be subjected to more exacting investigation would be to remove the requirement to complete a circuit. A second way would be to employ random configurations. A third way would be to increase the number of nodes in each problem in order to provide more potential for participants to produce sub-optimal tours and to display individual differences. The following experiment was undertaken with these points in mind.

4.1 *Method*

4.1.1 *Participants.* A total of eighty-one students from the University of Adelaide and the wider community served as volunteer participants in the experiment.

4.1.2 *Stimuli.* Fifty 30-node arrays were generated, in which the coordinates were randomly selected according to a uniform distribution. Benchmark optimal or near-optimal solutions were found for each problem, with the use of a simulated annealing algorithm (Press 1992), according to whether the task was framed as an Open (O) or a Closed (C) problem. Of the fifty 30-node arrays, five were identified as having the same optimal solution for the Open and the Closed tasks, except for the longest link in the case of the Open paths. A further five 30-node arrays were identified as having the largest percentage difference in total length between the Open and the Closed optimal solutions (minus the largest link). In addition, none of the optimal Open paths for these arrays followed the points of the convex hull in order of adjacency.

The stimuli for the experiment comprised two subsets of five Same (S) and five Different (D) 30-node arrays (which were the most populated arrays we could find that could be partitioned in this way). Figures 1a and 1b show the shortest Open and Closed solutions, respectively, for one of the arrays from subset S, while figures 1c and 1d show the shortest Open and Closed solutions, respectively, for an array from subset D.

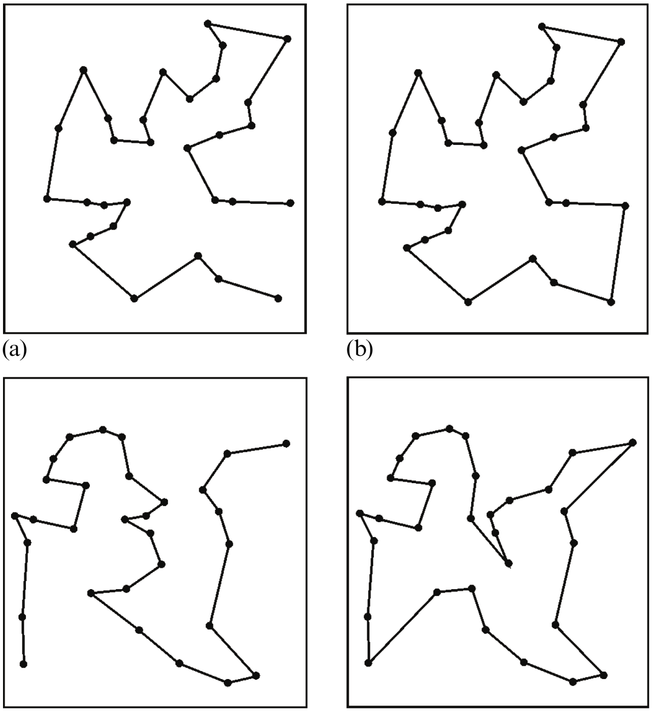

(a) (b) (c) (d)

Figure 1. Examples of problem arrays and optimal solutions: the shortest Open (a) and Closed (b) solutions for an array from the subset of Same problems; the shortest Open (c) and Closed (d) solutions for an array from the subset of Different problems.

4.1.3 *Design.* Participants were randomly assigned to two groups, O – C and C – O. Each group performed the experiment in two consecutive sessions of about 20 min each, separated by a 5-min break. There were ten problems in each session. Group O – C performed the Open path task for all ten problems, followed by the Closed tour task for all ten problems. Group C – O performed the Closed tour tasks first, followed by the Open path tasks. Within each group, each participant performed the ten problems in a different, unique random order in each session. The stimuli for the Closed tour tasks were the same as for the Open path tasks, except that they were reflected about the 45° diagonal. No participant reported having recognised the transformed and untransformed random configurations as otherwise identical.

4.1.4 *Instructions.* Participants were given instructions by way of an introductory information sheet, in the form of verbal instructions by the experimenter, and by means of instructions on the computer screen at the start of each session (which could be re-displayed at any time by clicking on a button). Instructions for Open paths emphasised that the path could begin at any node (town), but had to finish at a different node (town). Instructions for Closed tours emphasised that the tour could begin at any node, but had also to finish at that node. No indication was given that self-intersecting solutions might be inefficient.

4.1.5 *Procedure.* Participants were tested individually or in small groups at well-separated computers. Problems were presented, one at a time, in a 6 inch × 6 inch square in the centre of a standard 16-inch computer display. Participants could begin at any point by left-clicking on a node with the computer mouse. Then, whilst holding down the mouse button, they drew a path by positioning the mouse cursor on 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, participants could undo any links they had drawn. Participants were thus free to connect the nodes in any order, to work alternately from two nodes, or to work in several separated clusters of nodes.

Participants signalled when they had completed a problem by clicking on a screen button labelled 'Proceed to next test'. At this point, if a participant's completed tour was invalid (eg because not all nodes had been connected or because a node had more than two links), a warning message was posted on the screen and the participant was obliged to construct a valid solution before proceeding to the next problem. In an effort to preserve motivation, the optimal solution was then displayed, with the participant's own solution superimposed in light-grey (and slightly displaced), along with his/her 'score' (ie solution length, expressed as a percentage above the length of the optimal solution). When they felt ready, participants clicked on the 'OK' button in the message box, and were presented with the next problem. Each session began with a practice problem, followed by the ten test problems for the corresponding task.

5 Results

As described below, three measures of performance were examined: the relative frequency of responses of a particular type; the extent to which total solution lengths exceeded those of the corresponding optimal solutions; and a group measure of response uncertainty. Frequency measures tallied the number of solutions containing intersections and the number linking the nodes of the convex hull in order of adjacency. To permit comparison across the different problem types, participants' solution lengths were also expressed as a percentage above the benchmark solution length. Thus, a participant's score for a perfect solution would be 0, and any solution that exceeded the optimum would have a positive score. In addition, following MacGregor and Ormerod (1996), an information-theoretic group measure of response uncertainty was calculated for each problem. The frequency with which each point in the array was connected with each other point was tabulated in an $n \times n$ matrix. Probabilities (p_i) were obtained by dividing the frequencies by the number of participants (81). Shannon's (1948) standard information-theory formula [given in equation (1)], was then applied to calculate the total response uncertainty (H), with k being the total number of connections made by the participants.

$$H = -\sum_{i=1}^{k} p_i \log_2 p_i \, . \qquad (1)$$

5.1 *Comparisons between performance on Open and Closed problems*

5.1.1 *Relative frequency measures.* Out of the total of 1620 solutions (20 for each of eighty-one participants), only 0.7% contained an intersecting pathway (5 out of 810 Closed and 6 out of 810 Open pathways).

All Closed solutions for both Same and Different problems adhered to the cyclic order of the convex hull. However, 26% of Open solutions for Same problems and 55% of Open solutions for Different problems did not follow the nodes of the convex hull in order of adjacency.

5.1.2 *Deviations from optimality.* The mean percentage deviations from optimality are listed in table 1. From table 1, it may be seen that performance on Open Same problems was consistently poorer than that on Closed Same problems (averaging 6.2%, as opposed to 4.7%, respectively). Performance on Open Different problems was also poorer than that for Closed Different problems (averaging 8.1%, as opposed to 4.9%, respectively).

Table 1. Means, standard deviations, and minimum and maximum values for percentage deviations from optimality for each Same and Different problem under Closed and Open task instructions.

Task	Same problems					Different problems				
	S1	S2	S3	S4	S5	D1	D2	D3	D4	D5
Closed tours										
mean	3.30	4.81	6.59	2.73	6.01	5.56	5.02	5.28	5.59	3.16
minimum	0	0	0	0	0.16	0.12	0	0.04	1.23	0
SD	4.44	4.57	5.35	3.04	4.26	4.65	3.95	3.59	3.70	3.71
maximum	21.31	27.65	20.26	11.98	26.69	20.39	19.99	19.91	22.94	17.50
Open paths										
mean	5.63	6.63	7.80	3.87	6.82	10.33	8.38	10.37	6.12	5.13
minimum	0	0	0	0	0.09	0	0	0	0	0
SD	4.88	6.16	4.79	4.57	4.79	5.87	5.69	8.50	5.54	5.35
maximum	19.25	28.96	21.34	20.57	26.20	25.85	32.36	64.55	28.39	20.77

As can be seen from figure 2a, in the case of Same problems, performance on Open tasks (as measured by the mean percentage deviation from the optimal length) was a linear function of performance on Closed tours, with the correlation between performance on Open and Closed tasks being given by Pearson's $r = 0.93$. In the case of Different problems, illustrated in figure 2b, the mean percentage deviation from the optimal length for Open problems was also an approximately linear function of performance on Closed tours, with Pearson's r coefficient of 0.64.

5.1.3 *Response uncertainty.* Measures of response uncertainty for each problem are listed in table 2. From table 2, it appears that performance on Open Same problems consistently showed about the same response uncertainty as that on Closed Same problems (averaging 21.8%, as opposed to 21.4%, respectively). Performance on Open Different problems showed slightly less response uncertainty (though not consistently) than for Closed Different problems (averaging 21.4%, as opposed to 23.3%, respectively).

Figure 2c shows response uncertainty on Open Same tasks, as a function of that on Closed Same tasks. The best description of the relation is a linear one, with a correlation of $r = 0.93$. Figure 2d shows the same relationship in the case of Different tasks. Again, the best description is a linear function, with a correlation of $r = 0.91$.

5.1.4 *Individual differences.* As detailed above, overall mean performance on Open problems is related to that on Closed problems. The relation between performance on different problem types is also seen when mean percentage deviations from optimality for individual participants are correlated between different tasks and problem types, as summarised in table 3. Because they are each based on 81 pairs of scores, these correlations are all strongly reliable.

5.2 *Performance across Open and Closed problems as a function of objective measures of solution difficulty and complexity*

To date, the two independent measures of problem characteristics that have been investigated are the total number of nodes and the number of nodes on the convex hull (or, equivalently, the number of interior nodes inside the convex hull). However, as Lee and Vickers (2000) point out, the number of hull-bound points approaches an asymptotic value of around 12 for arrays with more than 50 points. With 30 nodes, the expected variation in the number of hull-bound nodes in a sample of randomly distributed nodes is very restricted (mean = 8.2; SD = 1.1). Although it is possible to select arrays with greater or lesser numbers of hull-bound points, it is not known how such selection might interact with other geometrical properties of the arrays that might be

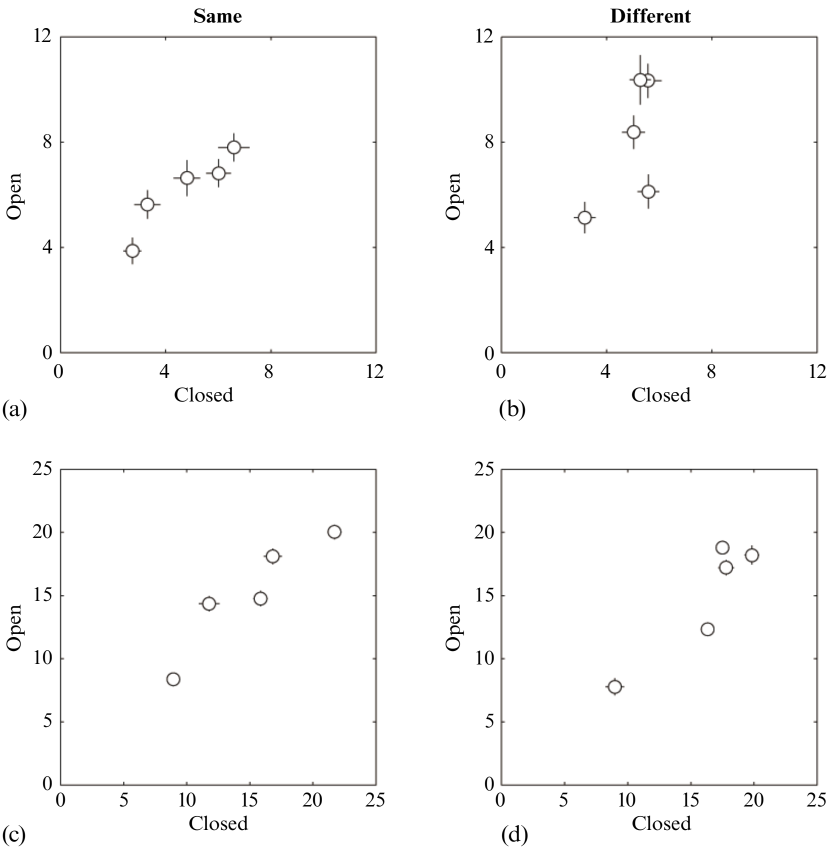

Figure 2. Performance measures on Open tasks as a function of the corresponding measures for Closed tasks: the relation for percentage deviations from optimality for the five Same problems (a) and the five Different problems (b); and the relation for response uncertainty for the five Same problems (c) and for the five Different problems (d).

Table 2. Measures of response uncertainty for each Same and Different problem under Closed and Open task instructions.

Task	Same problems					Different problems				
	S1	S2	S3	S4	S5	D1	D2	D3	D4	D5
Closed tours	16.2	22.8	24.0	12.4	31.4	25.9	25.0	29.4	23.7	12.4
Open paths	20.7	21.3	26.1	12.1	28.9	24.8	27.1	26.3	17.8	11.2

Table 3. Pearson's r correlations between percentage deviations from optimality (averaged over problem instances within each task and problem type) for different pairs of task and problem type. Each correlation is based on 81 pairs of mean scores.

Task	Closed		Open	
	Same	Different	Same	Different
Closed				
Same		0.61	0.48	0.47
Different			0.51	0.55
Open				
Same				0.70
Different				

important for locally focused perceptual processing. Moreover, although MacGregor and Ormerod's (1996) convex-hull hypothesis would suggest that the convex hull might retain its perceptual salience in the case of Open problems, there is no other strong reason to believe that the number of hull-bound points would be relevant to solution difficulty in such problems.

In a first attempt to grapple with such problems, therefore, we devised three objective, independent measures of solution difficulty and complexity. Each of these measures can be applied to both Open and Closed TSP tasks.

5.2.1 *Self-intersections.* The first is a measure of the overall potential of any array to produce self-intersecting paths (SIs), and it is intended to gauge the extent to which a participant's solution behaviour might be guided by the constraint of (deliberately or incidentally) avoiding crossed paths. It consists simply of counting, for each array, the number of segment pairs that potentially intersect.

5.2.2 *Path complexity.* The second measure is an index of path complexity (PC) for a given solution, that is similar to that employed by Vickers et al (in press, a), but that, in this case, is intended to gauge the extent to which a given array might present difficulty in continuing a path for a locally focused nearest-neighbour (NN) algorithm. The measurement procedure uses the benchmark solution for the problem. This identifies a set of links between the nodes, as illustrated for an elementary tour in figure 3a, for which an SNN algorithm would yield the optimal solution. In this simple case, each node will be connected to its nearest neighbour on one side and to its second nearest neighbour on the other side. Thus, each link will have associated with it two orders of proximity—one order contributed by each node. These orders can be represented by positive integers and summed for each link.

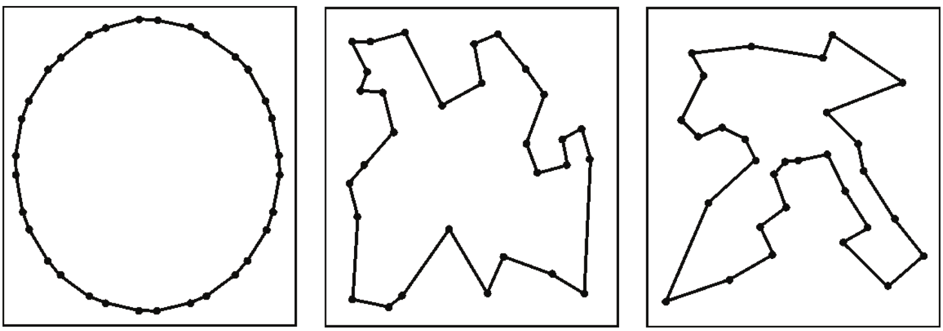

(a) (b) (c)

Figure 3. Illustrations of the application of the path complexity measure: a Closed tour, for which the nearest neighbour algorithm provides the optimal solution (a) [Each node is connected to its nearest neighbour (1) on one side and its second nearest (2) on the other. Thus, the proximity orders for each of the 32 links in this tour are $(1 + 2)$, and the path complexity measure is $96/32 = 3.0$ (while the proportion of potential crossings is a high 32.7%)], and two tours [(b) and (c)] similar to those in the present experiment, but differing in their proportions of potential crossings (25.2% and 21.4%, respectively) and in their path complexity measures (4.17 and 4.87, respectively).

In the case of the tour shown in figure 3a, the total of these summed orders of proximity is 96. When divided by the number of links (32), this yields the number 3.0 as an index of the average proximity order for each node. This is the minimum possible for a Closed tour, in which the optimal tour connects only nearest neighbours. Thus, the amount by which the average proximity order for a Closed tour exceeds 3 provides an index of the difficulty of that configuration for an NN algorithm. In the case of Open paths, the minimum average proximity order is $3 - [1/(n - 1)]$, where n is the number of nodes. Figures 3b and 3c show examples of tours of varying path complexities and with different numbers of potential crossings.

5.2.3 *Proportion of common nearest-neighbour links.* According to Vickers et al (submitted, c), an important source of local information concerning visual structure is provided by the set of (least) distances between the nearest neighbours in a pattern. A sequentially executed procedure, such as the SNN algorithm, starts at one node and consecutively links up the least distance remaining at each step. Thus, for each starting position of the algorithm, there will be a number of least-distance links that this particular SNN solution shares with the optimal solution. The proportion of common nearest-neighbour links (CL) is the mean number of links (averaged over all n starting positions) that the SNN algorithm shares with the optimal solution. When this proportion is high, it is assumed that the problem lends itself to a local processing procedure based on detecting least distances. Conversely, when CL is low, it is assumed that the problem will be more difficult, because the solution is quite different from the naturally seen perceptual structure.

Whereas the path complexity index, PC, is calculated from the optimal solution, the proportion of common links, CL, takes into account all possible SNN solutions. Therefore, although we might expect these measures to be inversely related, there is only a moderate correlation between the two for the problems in this experiment (Pearson's $r = -0.67$ and -0.87; $N = 10$; for Closed and Open solutions, respectively). Although the sequential nature of the SNN algorithm makes intersections possible, the least-distance constraint means that crossed paths will generally be avoided.

5.3 *Prediction of deviations from optimality and response uncertainty*

To examine how the theoretical measures of stimulus arrays—SI, PC, and CL—related to the deviation-from-optimality and response-uncertainty performance measures, we undertook a simple linear-regression analysis. There are eight possible linear models, comprising the three theoretical measures. These range in complexity from one where performance is assumed to be constant, and not affected by the measures, up to a four-term linear combination that includes all three measures plus a constant. The remaining six models include a subset of one or two of the theoretical measures plus a constant, and have intermediate complexity.

We fitted each of the eight models to both the deviation-from-optimality and response-uncertainty data using the method of maximum likelihood, and evaluated which of the models provided the best account. Because the models have different complexities, deciding which is the best model cannot be done solely through comparing their levels of fit or correlation (Roberts and Pashler 2000). Accordingly, we used the Bayesian approach to model selection (eg Kass and Raftery 1995; Myung and Pitt 1997; Pitt et al 2002), which balances the competing demands of fit and complexity. This means that the theoretical measures chosen by Bayesian methods for inclusion in the best linear models can be interpreted as providing the account of the data that is the most accurate and parsimonious.

A succinct summary of the outcomes of evaluation procedures by our Bayesian model is given by the odds ratios known as Bayes factors (Kass and Raftery 1995). For two competing models, M_i and M_j, the Bayes factor is given by $B_{ij} = p(D|M_i)/p(D|M_j)$. This ratio measures how much more likely it is that the data, D, would be observed if model M_i rather than M_j, was the true model.

Computationally, we estimated Bayes factors using an index called the Bayesian information criterion (BIC; Schwarz 1978) that is simple to calculate, given the fit and parametric complexities of our eight models. For a model, M, with K parameters, and having maximum likelihood parameterisation θ^* under the likelihood function $p(\cdot)$ to data D with N samples, the BIC is given by $\text{BIC} = -2\log p(D|M, \theta^*) + K \log N$. Qualitatively, it can be seen that the BIC decreases as the log-likelihood fit of a model to data improves, but increases as additional parameters are included in the model. Accordingly, the model with the minimum BIC value corresponds to the most likely model. The Bayes factors for quantifying how much more likely this model is than the others can be estimated by using the approximation: $\mathrm{BIC}_i - \mathrm{BIC}_j \approx -2 \log B_{ij}$.

We assumed that both the deviations from optimality and response-uncertainty data were normally distributed with common variance, and so used the likelihood function,

$$p(D|M, \theta) = \prod_i \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(d_i - \hat{d}_i)^2}{2\sigma^2}\right],$$

where $\hat{d}_i$ is the prediction made by model M at parameterisation θ for the ith datum. This assumption has the advantage of being sensitive to the precision, σ, of the empirical data. Precision values were estimated by using the variance across participants for deviation from optimality (Lee 2001), and by bootstrapping (Efron and Tibshirani, 1986) for response uncertainty.

Table 4 details the Bayesian analysis for the eight linear models in relation to the deviation from optimality measure of performance. Each row corresponds to one of the possible models, and each column corresponds to a problem type, as defined by whether an Open or Closed tour was required through Same or Different stimulus arrays. Each entry in the table gives the Bayes factor for the model on the problem type, in relation to the best model for that problem type, followed, in brackets, by the correlation of the model at its maximum-likelihood parameterisation. The most likely model for each problem type, with a Bayes factor of 1.0, is highlighted in bold.

Table 4. Evaluation of all possible linear models with the use of the theoretical measures against deviation from optimality, broken down in terms of Closed and Open tours through Same and Different stimulus arrays. The odds against each model being true, in relation to the most likely, are shown first, followed by the correlation of the model in brackets. The most likely model for each problem type is highlighted in bold.

Model	Closed tours		Open tours					
	Same arrays		Different arrays		Same arrays		Different arrays	
Constant	>100	(0.00)	>100	(0.00)	>100	(0.00)	>100	(0.00)
PC	>100	(0.46)	2.59	(0.89)	9.01	(0.90)	>100	(0.78)
NN	29.72	(0.87)	**1.00**	**(0.92)**	**1.00**	**(0.97)**	>100	(0.69)
SI	>100	(0.52)	>100	(0.68)	>100	(0.74)	1.82	(0.90)
PC + NN	65.63	(0.87)	1.77	(0.94)	2.13	(0.97)	>100	(0.78)
PC + SI	>100	(0.60)	5.79	(0.89)	10.58	(0.92)	**1.00**	**(0.96)**
NN + SI	**1.00**	**(0.97)**	2.22	(0.92)	2.11	(0.97)	1.05	(0.94)
PC + NN + SI	1.29	(0.99)	3.67	(0.94)	4.11	(0.98)	1.97	(0.95)

For example, the first column in table 4 shows that the most likely linear model for Closed tours on Same arrays includes the NN and SI theoretical measures, and that this model correlates $r = 0.97$ with human performance. The more-complicated linear model that uses all three theoretical measures is 1.29 times less likely, even though it has a correlation of $r = 0.99$. The Bayes factors also show that linear models that use only the NN theoretical measure, and those that use the PC and NN theoretical measures are, respectively, 29.72 and 65.63 times less likely. All of the remaining models are more than 100 times less likely than the one that uses the NN and SI measures.

Because Bayes factors are odds, they have a natural meaning based on gambling, and so can be interpreted for 'significance' in terms of the standards of scientific evidence required for a problem, rather than by reference to conventional critical values.

Table 5 details the Bayesian analysis for the eight linear models in relation to the response-uncertainty measure of performance, in the same format as table 4.

Table 5. Evaluation of all possible linear models with the use of theoretical measures against response uncertainty, broken down in terms of Closed and Open tours through Same and Different stimulus arrays. The odds against each model being true, in relation to the most likely model, are shown first, followed by the correlation of the model in brackets. The most likely model for each problem type is highlighted in bold.

Model	Closed tours		Open tours					
	Same arrays		Different arrays		Same arrays		Different arrays	
Constant	>100	(0.00)	>100	(0.00)	>100	(−0.00)	>100	(0.00)
PC	>100	(0.86)	60.24	(0.97)	>100	(0.74)	>100	(0.82)
NN	>100	(0.71)	>100	(0.95)	>100	(0.53)	>100	(0.93)
SI	>100	(0.70)	>100	(0.71)	>100	(0.43)	>100	(0.74)
PC + NN	>100	(0.91)	**1.00**	**(0.99)**	>100	(0.88)	>100	(0.94)
PC + SI	4.80	(0.97)	>100	(0.97)	>100	(0.74)	>100	(0.89)
NN + SI	>100	(0.72)	>100	(0.95)	>100	(0.53)	**1.00**	**(0.99)**
PC + NN + SI	**1.00**	**(0.99)**	1.22	(1.00)	**1.00**	**(0.96)**	1.01	(0.99)

6 Discussion

6.1 *Comparisons between performance on Open and Closed problems*

Given that participants in the Closed tasks almost always avoided intersections, and produced near-optimal solutions, it follows logically that their solutions must follow the convex hull in sequence. On the other hand, the contrasting result that 55% of solutions for the Open Different problems failed to follow the convex hull suggests that, if participants are influenced by the convex hull, then it is ineffective in the case of Open problems (although this conclusion must be qualified, since none of the optimal solutions to these problems completely obeyed the convex-hull order). However, since the optimal solutions for both the Open and Closed Same problems were identical, it is remarkable that as many as 26% of solutions for Open Same problems failed to connect the nodes on the convex hull in order of adjacency. If performance on TSP tasks were largely determined by the automatic perception of a global aspect, such as the convex hull, then there is no obvious reason why it should not be equally effective in the case of Open problems, particularly when the optimal solution to these problems does follow the convex hull in cyclic order.

Measured in terms of percentage above the optimal tour length, performance on the Open problems was consistently somewhat poorer than that on the corresponding Closed problems (7.1%, as opposed to 4.8%, respectively). A possible reason for this is that, theoretically, an Open tour problem with n nodes is equivalent to a Closed tour problem with $n + 1$ nodes (Lawler et al 1985: 22; Reinelt 1994: 32).

For the problems used in the present experiment, this implies that the Open problems have some 30 times the number of alternative pathways to consider than do the Closed problems. However, the decrement in performance is far short of this order. Indeed, the increase of 2.3% in percentage above optimal is closer to the order that would be expected if performance were a linear function of the number of nodes. Evidence that participants' deviations from optimality and solution times both appear to be linear functions of the number of nodes has been presented by Graham et al (2000). The present result and the linear relationships found by Graham et al both are to be expected on the local-to-global approach described in section 1.

The finding that deviations from optimality and response uncertainty on Open problems both appear to be approximately linear functions of those on Closed problems suggests that, notwithstanding the irrelevance of the convex hull to Open problems, the two tasks appear to share a significant amount of processing. The result that, for deviations from optimality, the relationship appears to be more reliable in the case of

Same problems, rather than Different ones, can be attributed to the fact that the configurations of the optimal solutions for Open and Closed problems were identical for Same problems, while the configurations of optimal solutions for Open and Closed tasks were quite different in the case of Different problems.

On the local-to-global approach proposed by Vickers et al (in press, b), performance on both Open and Closed TSP tasks is mediated by the perception of structure, as identified by the pattern of least distances in an array. According to this approach, the proportion of common links should be related to performance on TSP tasks. As illustrated in figure 4, for both Same and Different arrays, this objective measure for Open tasks is an approximately linear function of the value of the measure for Closed tasks with the same arrays. The fact that this relation mirrors that found empirically lends support to the view that a major part of the processing shared between Open and Closed tasks is concerned with the identification of structure by means of least distances.

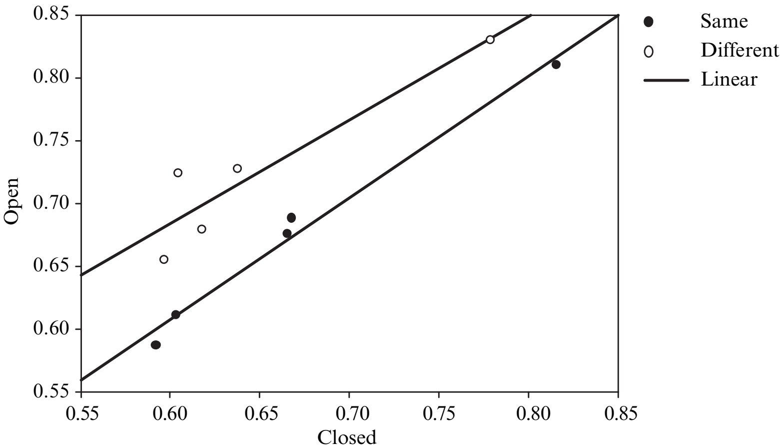

Figure 4. The proportions of common links for the five Same Open problems (filled circles) and the five Different Open problems (empty circles) as a function of the proportions for the five corresponding (Same and Different) Closed problems.

Meanwhile, the finding that there are reliable individual differences in performance on the problems in this experiment is at variance with the absence of such differences reported in the studies of MacGregor and Ormerod (1996), and MacGregor et al (1999, 2000). On the other hand, the evidence for individual differences confirms similar results by Vickers et al (2001), Vickers et al (in press, a), and Vickers et al (in press, b). The simplest explanation for the apparent discrepancy between the two sets of results is that the problems employed by MacGregor and his colleagues were generally much simpler, both in terms of the numbers of nodes and the constraint used in their construction, than the more populated random arrays employed in the latter set of experiments. In consequence, we surmise that performance in the first set of studies was probably limited by a ceiling effect.

The evidence for consistent individual differences across tasks and problem types suggests that, for each participant, there is a substantial overlap in the way these different task and problem types are processed. Such differences are clearly inconsistent with the view that TSP solutions are largely determined by a uniform, species-constant perceptual process. Although this conclusion by itself does not rule out the kind of global processing proposed by MacGregor and Ormerod (1996), it does open the way for more locally focused mechanisms, possibly integrated under top – down control.

6.2 *Prediction of performance across both Open and Closed problems*
More light is thrown on the nature of the shared processing between problem types when we examine the relationships between performance measures and the three objective indices of path difficulty and complexity.

In this respect, the results are reasonably clear. In the case of deviations from optimality, the most accurate and parsimonious predictive model is one based either on common links alone or on a combination of common links and the number of potential self-intersections. These results are consistent with a local-to-global HNN procedure, in which the perceived organisation in an array plays an important role, in the initial stages of the solution process for both Open and Closed problems, through the identification of least-distance links between nearest neighbours. According to this account, the difficulty associated with a problem is determined by the extent to which the naturally seen pattern of least-distance structure is represented in the optimal solution and by the extent to which the later stages of path construction are facilitated by the avoidance of intersections.

In the case of response uncertainty, the most accurate and parsimonious model is one based on common links, in conjunction with either or—most generally—both path complexity and potential self-intersections. Given that response uncertainty measures the variability in the choices made by different participants at each of the nodes, rather than summarising a common tendency, it is not surprising that the best account of the data is based on all three measures. What these results imply is that, although the proportion of common links is a good predictor of the extent to which participants' solutions approach the optimum, this factor does not explain the variability in these solutions. To explain response uncertainty requires taking into account the way in which common links are integrated into a complete path. That is, in integrating structures based on nearest-neighbour clusters, response uncertainty is determined by the extent to which participants are obliged to consider more complex paths, involving higher-order (second-, third-, ...) nearest neighbours, as well as by the extent to which they succeed in reducing the number of alternatives by avoiding intersections.

On the approach proposed here, performance on visually presented TSP tasks is indeed mediated by normal processes of perceptual organisation, as suggested by Polivanova (1974), MacGregor and Ormerod (1996), and by Vickers et al (2001). However, the perceptual processes involved do not appear to be global-to-local ones, based on the convex hull, but are predominantly local-to-global ones, based on the detection of nearest-neighbour clusters. This conclusion is consistent with other recent findings concerning the importance of nearest-neighbour relations for the perception of structure and motion in dot patterns (Vickers et al, submitted, c).

On the other hand, individual differences in performance appear to arise, not from variations in the way in which nearest-neighbour clusters are detected, but from the ways in which different participants connect up these clusters into a complete solution. This hierarchical interpretation corresponds closely to the description, given by participants in the experiment of Graham et al (2000), of the strategies they used. Indeed, our HNN hypothesis is structurally very similar to the hierarchical model of Graham et al, although it assumes a different process for the detection of clusters and leaves open, for the present, the question of how the clusters are put together to make a path.

Meanwhile, our results may also be given an interpretation that differs in another subtle, but important, way from that proposed by MacGregor et al (1999). These authors collectively ascribe their results to the operation of natural, perceptual organising tendencies and suggest that the relative ease with which participants produce high-quality solutions to TSPs may be because the task requirements of the TSP "happen to conform to natural tendencies of the perceptual system" (MacGregor and Ormerod 1996, pages 538 – 539). An alternative—and quite different—perspective is that the perception of some form of minimal structure is a natural tendency of the perceptual system, and Gestalt principles represent different instantiations of this tendency.

7 Conclusions

The results of the present study call into question the notion that performance on TSP tasks is mediated by a global-to-local processing, based on the convex hull. There is no evidence that a hull-based process is employed in Open problems or that it transfers from Closed to Open problems. In contrast, there is good evidence that participants depend on the same kind of primary information in both Open and Closed problems. In general agreement with previously published interpretations, the nature of this information seems to be strongly determined by natural, automatic organising tendencies of the visual system. However, the present results provide strong evidence for the view that performance in both Open and Closed tasks is based on local-to-global processes of detecting structure, based primarily on the detection of nearest-neighbour clusters. On this view, TSP performance is influenced by the same processes of perceptual organisation as are involved in the perception of visual structure and motion generally.

There is also good evidence for reliable individual differences that are consistent with shared processing between different tasks and types of problem. These differences appear to be associated with variations in the way in which nearest-neighbour clusters are integrated into a solution, rather than with differences in the efficiency with which such clusters are detected.

In turn, these conclusions suggest a reversal of perspective, in which the perception of organisation itself may perhaps usefully be thought of as the perception of minimal forms of structure.

Acknowledgments. The research on which this article is based was supported by Australian Research Council grants (DP0211150 and DP0210851) to D Vickers, and to D Vickers and M Lee, respectively. The experimental work reported in this paper was carried out while P Bovet held a Visiting Professorship at the University of Adelaide.

References

Efron B, Tibshirani R, 1986 "Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy" *Statistical Science* **1** 54 – 75

Flood M M, 1956 "The travelling salesman problem" *Operations Research* **4** 61 – 75

Golden B, Bodin L, Doyle T, Stewart W, 1980 "Approximate travelling salesman algorithms" *Operations Research* **28** 694 – 711

Graham S M, Joshi A, Pizlo Z, 2000 "The traveling salesman problem: a hierarchical model" *Memory and Cognition* **28** 1191 – 1204

Hwang F K, Richards D S, Winter P, 1992 *The Steiner Tree Problem* (Amsterdam: Elsevier)

Johnsonbaugh R, 2001 *Discrete Mathematics* (Upper Saddle River, NJ: Prentice-Hall)

Kass R E, Raftery A E, 1995 "Bayes factors" *Journal of the American Statistical Association* **90** 773 – 795

Lawler E L, Lenstra J K, Rinooy K A H G, Schmoys D B, 1985 *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization* (Chichester, UK: John Wiley)

Lee M D, 2001 "Determining the dimensionality of multidimensional scaling models for cognitive modeling" *Journal of Mathematical Psychology* **45** 149 – 166

Lee M D, Vickers D, 2000 "The importance of the convex hull for human performance on the traveling salesman problem: a comment upon MacGregor and Ormerod (1996)" *Perception & Psychophysics* **62** 226 – 228

MacGregor J N, Ormerod T C, 1996 "Human performance on the traveling salesman problem" *Perception & Psychophysics* **58** 527 – 539

MacGregor J N, Ormerod T C, 2000 "Evaluating the importance of the convex hull in solving the Euclidean version of the traveling salesperson problem: Reply to Lee and Vickers (2000)" *Perception & Psychophysics* **62** 1501 – 1503

MacGregor J N, Ormerod T C, Chronicle E P, 1999 "Spatial and contextual factors in human performance on the travelling salesperson problem" *Perception* **28** 1417 – 1428

MacGregor J N, Ormerod T C, Chronicle E P, 2000 "A model of human performance on the traveling salesperson problem" *Memory & Cognition* **28** 1183 – 1190

Myung I J, Pitt M A, 1997 "Applying Occam's razor in modeling cognition" *Psychonomic Bulletin & Review* **4** 79 – 95

Ormerod T C, Chronicle E P, 1999 "Global perceptual processing in problem solving: the case of the traveling salesperson" *Perception & Psychophysics* **61** 1227 – 1238

Palmer S E, 1999 *Vision Science: Photons to Phenomenology* (Cambridge, MA: MIT Press)

Pitt M A, Myung I J, Zhang S, 2002 "Toward a method of selecting among computational models of cognition" *Psychological Review* **109** 472 – 491

Polivanova N I, 1974 "Functional and structural aspects of the visual components of intuition in problem solving" *Voprosy Psikhologii* **4** 41 – 51 (in Russian)

Pomerantz J R, 1981 "Perceptual organization and information processing", in *Perceptual Organisation* Eds M Kubovy, J R Pomerantz (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 141 – 180

Press W H, 1992 *Numerical Recipes in FORTRAN: The Art of Scientific Computing* 2nd edition (Cambridge: Cambridge University Press)

Quintas L V, Supnik F, 1965 "On some properties of shortest Hamiltonian circuits" *American Mathematical Monthly* **72** 977 – 980

Reinelt G, 1994 "The Traveling Salesman: Computational Solutions for TSP Applications" Lecture Notes in Computer Science, 840 (Berlin: Springer)

Roberts S, Pashler H, 2000 "How persuasive is a good fit? A comment on theory testing" *Psychological Review* **107** 358 – 367

Rooij I van, Stege U, Schactman A, 2003 "The convex hull and crossings in the Euclidean Traveling Salesperson problem: Implications for human performance studies" *Memory & Cognition* **31** 215 – 220

Schwartz G, 1978 "Estimating the dimension of a model" *The Annals of Statistics* **6** 461 – 464

Shannon C, 1948 "The mathematical theory of communication" *Bell Systems Technical Journal* **27** 379 – 423, 623 – 656

Stein D L, 1989 "Spin glasses" *Scientific American* **261** 36 – 42

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 (in press, a), "Optimization and intelligence: Individual differences in performance on visually presented optimization problems" *Personality and Individual Differences*

Vickers D, Lee M D, Dry M, Hughes P (in press, b), "The roles of the convex hull and the number of potential crossings upon performance on visually presented traveling salesperson problems" *Memory & Cognition*

Vickers D, Preiss K, Hughes P (submitted, c) "The role of nearest neighbours in the perception of structure and motion in dot patterns"

Wilf H S, 1986 *Algorithms and Complexity* (Englewood Cliffs, NJ: Prentice-Hall)

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 32 2003

www.perceptionweb.com