PERGAMON

PERSONALITY AND INDIVIDUAL DIFFERENCES

Personality and Individual Differences 36 (2004) 1059–1071

www.elsevier.com/locate/paid

Intelligence and individual differences in performance on three types of visually presented optimisation problems

Douglas Vickers*, Therese Mayo, Megan Heitmann, Michael D. Lee, Peter Hughes

Department of Psychology, University of Adelaide, Adelaide, South Australia 5005, Australia

Received 31 July 2002; received in revised form 14 March 2003; accepted 18 April 2003

Abstract

Although problem solving is an essential expression of intelligence, both experimental and differential psychology have neglected an important class of problems, for which it is difficult or impossible for systematic procedures to provide a definitive solution. Two experiments are described, in which participants' solutions to three computationally difficult problems (a Travelling Salesperson, a Minimal Spanning Tree, and a Generalised Steiner Tree problem) all showed consistent individual differences that intercorrelated reliably and correlated moderately with scores on Raven's Advanced Progressive Matrices. The results are interpreted in terms of a theory of visual perception based on the efficient use of information about the relative position of stimulus elements.

Keywords: Optimisation; Intelligence; Problem solving; Cognition; Visual perception; Structure

1. Introduction

In the everyday world, the ability to solve practical problems is generally regarded as an expression of intelligence. Similarly, in the psychological laboratory, the cognitive aspects of problem solving have long been considered as essential to any well conceived notion of intelligence (Resnick & Glaser, 1976; Sternberg, 1982).

* Corresponding author. Tel.: +61-8-8303-5662; fax: +61-8-8303-3770.
E-mail address: douglas.vickers@adelaide.edu.au,
http://www.psychology.adelaide.edu.au/members/staff/dougvickers/doug.html (D. Vickers).

doi:10.1016/S0191-8869(03)00200-9

Research into problem-solving processes and abilities has undergone considerable change (Lovett, 2002). Much early work was carried out by researchers following the Gestalt tradition and required making a single, insightful step that led to a solution through a restructuring of the representation of the problem in question—a process with obvious similarities to perceptual reorganisation (Duncan, 1959). Later work continued to focus on 'knowledge-lean' problems that required no background knowledge but placed greater emphasis on memory and on sustained cognitive strategy through the study of multi-step problems (Newell & Simon, 1972). In research over the last two decades, attention has shifted to studying 'knowledge-rich' domains, such as chess and medical diagnosis, and has focussed in particular on the nature and acquisition of expertise (e.g. Chi, Glaser, & Farr, 1988).

The wide spectrum of cognitive activities covered by this evolution presents a dilemma for attempts to understand individual differences in intellectual ability in terms of specific cognitive processes. On the one hand, as argued by Vernon and Strudensky (1988), the problems that have received most attention in knowledge-lean laboratory studies (such as the Missionaries and Cannibals and Tower of Hanoi problems) do not have strong claims to being ecologically representative. On the other hand, research that has focussed on various real-world content domains has emphasised the specificity of expertise, the limited transfer to a different domain, and the differences in problem-solving performances across different domains (Ericsson & Lehmann, 1996).

One feature common to most laboratory problem-solving tasks is the possession of well defined initial and end states that must be linked by a sequence of operations. The problems in such tasks typically have a unique answer that can be arrived at by an algorithmic procedure (i.e. a systematic, finite series of steps). Indeed, in a comprehensive review of early work, Ray (1955) specified solvability in terms of a unique solution as a desirable criterion for any problem-solving task. In contrast, working within the context of domain-specific, practical problem solving, Smith (1991: 14) has argued for a different perspective, in which a problem is defined as "a task that ... cannot be solved by recall, recognition, reproduction, or *application of an algorithm alone*" [italics added].

This definition draws attention to an important class of real-world problems that have been neglected in both experimental and differential psychology. These are problems for which there is no algorithm that can be guaranteed to produce a definitive solution within a time that is practicable. Because these problems can often be stated simply and readily understood, they lend themselves to laboratory investigation. At the same time, however, such problems are representative of many naturally occurring, real-world situations that are of practical interest, as illustrated by the fields of mathematics and computer science dealing with combinatorial optimisation.

A classic example of a computationally intractable optimisation problem is the so-called Travelling Salesperson Problem (Lawler, Lenstra, Rinooy Kan, & Shmoys, 1985), historically referred to as the Travelling Salesman Problem (TSP). The planar Euclidean version of the TSP can be formulated as follows: given a set of n interconnected towns, represented by nodes on a graph, devise an itinerary that visits each town (node) exactly once, returns to the starting node, and ensures that the total distance travelled is as short as possible. To arrive at a definitive solution to this problem entails an exhaustive consideration of $(n-1)!/2$ pathways. This is feasible when n is a small number, such as 5. However, when n is just 25, the number of pathways becomes so immense that a computer evaluating a million possibilities a second would take almost 10 billion years to evaluate them all (Stein, 1989).

Despite these computational difficulties, human performance on visually presented TSP tasks is impressively close to optimal, and, in some cases, exceeds that of computational heuristics, at least when n is less than 100 (Graham, Joshi, & Pizlo, 2000). For example, MacGregor and Ormerod (1996) found that participants' solutions were so close to the best known solutions that there were no individual differences and there was zero correlation across different problem instances. On these grounds, MacGregor and Ormerod concluded that the solution processes employed correspond to spontaneous, organising tendencies of the visual system that are uniformly present and species-constant.

This conclusion has been questioned by Vickers, Butavicius, Lee and Medvedev (2001). In their first experiment, employing two 10-, 25- and 40-node problem instances, there was a high rank order correlation (Spearman $\rho = 0.81$; $N = 18$) in normalised measures of solution length, averaged over all 15 pairs of instances. In their second experiment, employing a single 40-node problem instance, there was also a low, but significant, correlation (Spearman $\rho = -0.36$; $N = 40$) between normalised solution lengths and scores on Raven's Advanced Progressive Matrices (APM; Raven, Court, & Raven, 1988). As pointed out by Lee and Vickers (2000), the problem instances employed by MacGregor and Ormerod (1996) were restricted to 10 or 20 nodes only and were constructed in a highly constrained way that was conducive to near-optimal (ceiling) performance by most participants.

In view of the apparent conflict in previous results, an experiment was carried out to investigate individual differences in performance with several instances of an even more difficult, visually presented TSP task and to examine their relationship with scores on the APM. A second experiment was carried out to explore the possibility that intelligence might be associated with a more general ability to arrive at near-optimal solutions, other than the TSP, and to investigate the question of whether such an ability is predominantly determined by automatic perceptual tendencies or by an interplay between perceptual organisation and cognitive control.

2. Experiment 1

2.1. Method

2.1.1. Participants

There were 69 volunteer participants, comprising University of Adelaide students and opportunistically selected members of the public, comparable in age, educational background and socio-economic status. Ages ranged from 17 to 53 years, with an approximately equal number of males and females.

2.1.2. Stimuli and materials

Participants were presented with a booklet containing five 50-node arrays, representing TSP instances, and printed on separate sheets of A4 paper. The coordinates for each node were generated by selecting randomly from a uniform distribution. The only constraint applied was that each problem instance should have an equal number of nodes on the convex hull (a boundary enclosing the nodes, such that no line segment joining any two nodes can cross the boundary.) This constraint was applied because MacGregor and Ormerod (1996) hypothesised that participants employ the convex hull to establish an initial sub-tour, before inserting interior nodes in sequence.

In addition, participants completed Set 1 (practice) and Set 2 (test) of the APM.

2.1.3. Procedure

The APM was administered singly or in small groups of <8 participants. Set 1, containing 12 items, was first completed as practice. Immediately afterward, Set 2, containing 36 items, was timed to be completed in 40 min.

Participants tackled the TSP instances either 15 min after collection of the APM material, or in a separate session. In order to balance possible order effects, the five instances were presented to each participant, according to a Latin square design (Fisher & Yates, 1963). Participants were instructed to draw the shortest possible pathway that passed through each node once and once only and returned to the starting node. Participants were allowed to begin at any node and to take as long as they needed. However, no participant took longer than 30 min.

2.2. Results

2.2.1. Overall performance

Benchmark (best) solutions were found for each of the five problem instances using a simulated annealing algorithm (Press, 1992). Participants' solution lengths for each problem were expressed as a proportion above the benchmark solution length (PAB) for that problem. Thus, a PAB score for an optimal solution would be 0 and any solution that exceeded the benchmark would have a positive score. The means of the PAB scores for instances 1–5 were 0.09, 0.07, 0.07, 0.09 and 0.08 (SDs = 0.06, 0.04, 0.07, 0.07 and 0.07), respectively.

2.2.2. Performance across problem instances

Table 1 shows a matrix of Pearson correlations between the PAB scores for each of the 10 possible pairs of instances. Although modest, each of these 10 correlations is positive and

Table 1
Pearson intercorrelations between PAB scores for the 5 TSP instances and APM scores in Experiment 1

TSP problem	TSP problem					APM
	1	2	3	4	5	
1		0.55	0.70	0.48	0.41	−0.34
2			0.58	0.45	0.56	−0.37
3				0.62	0.45	−0.28
4					0.43	−0.34
5						−0.53

statistically significant ($P<0.01$; $N=69$).[1] There was an acceptably high degree of internal consistency among the five instances, with a value of 0.84 for Cronbach's α. Table 1 also shows the Pearson correlations between PAB and APM scores for each of the five problem instances. Although low to moderate, each of these correlations is negative and significant ($P<0.05$; $N=69$). The Pearson correlation between APM scores and the overall average PAB scores was $r=-0.48$, which is clearly significant statistically ($P\leqslant 0.01$).

2.2.3. *Other correlates with TSP performance and APM scores*

Vickers, Preiss, and Hughes (submitted for publication) have presented evidence that an important source of information concerning visual structure is provided by the set of (least) distances between the nearest neighbours of any pattern that consists of an array of nodes. This hypothesised basis for the perception of organisation is closely related to the sequential procedure for arriving at near-optimal TSP solutions executed by the standard 'Nearest Neighbour' algorithm, which starts at one node and consecutively links up the next *remaining* least distance at each step (Golden, Bodin, Doyle, & Stewart, 1980).

To examine the extent to which participants made use of nearest neighbours in arriving at a solution, measures of *path complexity* were calculated for each participant's solution to each problem. For each link of a participant's solution, numbers (1 to n) were assigned, according to whether the nodes of that link were connected to the nearest (1), second nearest (2), or nth nearest node. 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 succeeded in connecting each node to its nearest neighbours (i.e. the more the participant made use of least distances in arriving at a solution).

There was a strong correlation between participants' measures of path complexity, averaged over the five instances, and their PAB measures, also averaged over the five instances (Pearson $r=0.98$; $P\leqslant 0.01$; $N=69$). There was also a moderately good correlation between APM scores and participants' path complexity ($r=-0.48$; $P\leqslant 0.01$; $N=69$).

[1] *Statistical significance of the correlations*. Under the Null-Hypothesis Significance Testing (NHST) approach to statistical inference, the correlations between performances on the same test, between average performances on different tests, between test performances and APM scores, and between path complexity measures and performance scores are all highly significant, with $P<0.01$, except in one case, where $p<0.05$. We are sensitive, however, to criticisms of NHST (e.g. Cohen, 1994; Edwards, Lindman, & Savage, 1963; Hunter, 1997). In particular, we acknowledge that NHST violates the likelihood principle, and so, for the reasons explained by Lindley (1972), does not satisfy a basic requirement for rational, consistent and coherent statistical decision making. Accordingly, we also undertook Bayesian analyses of the data (e.g. Lindley, 1972; Sivia, 1996). In particular, we assessed the statistical significance of correlations using Bayes Factors (Kass & Raftery, 1995), which measure how much more (or less) likely one model is than another, on the basis of the evidence provided by data. For correlational analyses, the two models we compared were a 'null model', which assumed no relationship between the variables, and so assigned data points at random, and an alternative model that assumed the two variables were linearly related. The Bayes Factor for these two models is straightforward to estimate using standard Monte-Carlo techniques (Kass & Raftery, 1995, p. 779), involving the sampling of a large number of specific data distributions predicted by the null and alternative models, and calculating their average level of fit to the observed data. For all of the correlations reported earlier that are significant at the $P<0.01$ level under NHST, we found that the estimated Bayes Factors showed the alternative model was always more than 1000 times more likely than the null model. On this basis, it seems safe to conclude that the reported correlation coefficients indicate that a linear relationship between the variables is a better account than one based on random variation.

2.3. Discussion

Results of this experiment show consistent and reliable individual differences in PAB scores across different problems. This replicates similar findings by Vickers, Bovet, Lee, and Hughes (in press), Vickers et al. (2001), and Vickers, Lee, Dry, and Hughes (in press). Provided visually presented TSP problem instances have a sufficiently large number of randomly distributed nodes, then reliable individual differences emerge.

The relationship found between PAB scores and those on Raven's APM confirms the earlier result of Vickers et al. (2001). Meanwhile, the moderate to strong correlations between the number of closest links and both APM and PAB scores, respectively, suggests two conclusions. First, it indicates that performance on visual TSP tasks is mediated by a locally focussed process of detecting least distances between the nodes in an array. Second, it suggests that a significant factor in intelligence may be the ability to discern perceptual structure (itself a non-trivial combinatorial problem).

3. Experiment 2

To test the generality of the association between intelligence and the detection of minimal structures, and to investigate the possible dependence of such an association on automatic perceptual processes, three optimising tasks were selected. The first was a visually presented TSP task, with 50 randomly distributed nodes, illustrated in Fig. 1a. The second was a Minimum Spanning Tree Problem (MSTP), also with 50 nodes, illustrated in Fig. 1b. A minimum spanning tree is the shortest path that directly links all the nodes in an array. However, it has fewer constraints, because the path does not have to be continuous and closed, and a node can be connected by more than two links. Unlike the TSP, the MSTP is computationally tractable (Ahuja, Magnanti, & Orlin, 1993). The optimal solution is an open, branching path system, which directly links all the nodes, and in which branches occur at the nodes.

The third problem was a 15-node Generalised Steiner Tree Problem (GSTP) (Hwang, Richards, & Winter, 1992). In the elementary, three-node Steiner Tree problem, the shortest path connecting three nodes in a plane can be found by determining the so-called Fermat point, and creating a path of links that branch at that point (as illustrated in Fig. 1c). If all the angles in a triangle, ABC, are less than 120°, then the Fermat point, P, is the point within the triangle, from where the angles APC, CPB, and BPA are all exactly 120°. However, if one of the angles is at least 120° (say, the angle ACB), then P must coincide with C.

GSTPs contain more than three nodes, and the solutions to these problems look like combinations of solutions found in three-node Steiner Tree Problems, with additional nodes (P_1, P_2, ... P_k) as branch points to create the minimum connections between the original nodes, as illustrated in Fig. 1d.

These three problems were selected because they appeared a priori to make differing use of automatic perceptual processes. For example, contrary to the convex hull hypothesis of MacGregor and Ormerod (1996), it appears that an important perceptual component in the TSP solution process is the detection of least distances between each node and its nearest neighbours. At the same time, the optimal solution may not include all the least distances. Thus, the TSP makes possible a tension between a spontaneous perceptual tendency and the cognitive demand of constructing a circuit that is shortest overall. On the other hand, the MSTP does include all the least distances, although it may also require additional links to connect up substructures of least distances. Viewed in this way, the solution procedure for the MSTP is arguably more compatible with a locally focussed, spontaneous perceptual process of detecting least distances. In contrast, the GSTP task is arguably the least compatible with natural perceptual tendencies, because it does not simply link the visible nodes directly but requires the imposition of additional nodes and links that are well separated from any structure formed by direct links and least distances between existing nodes.

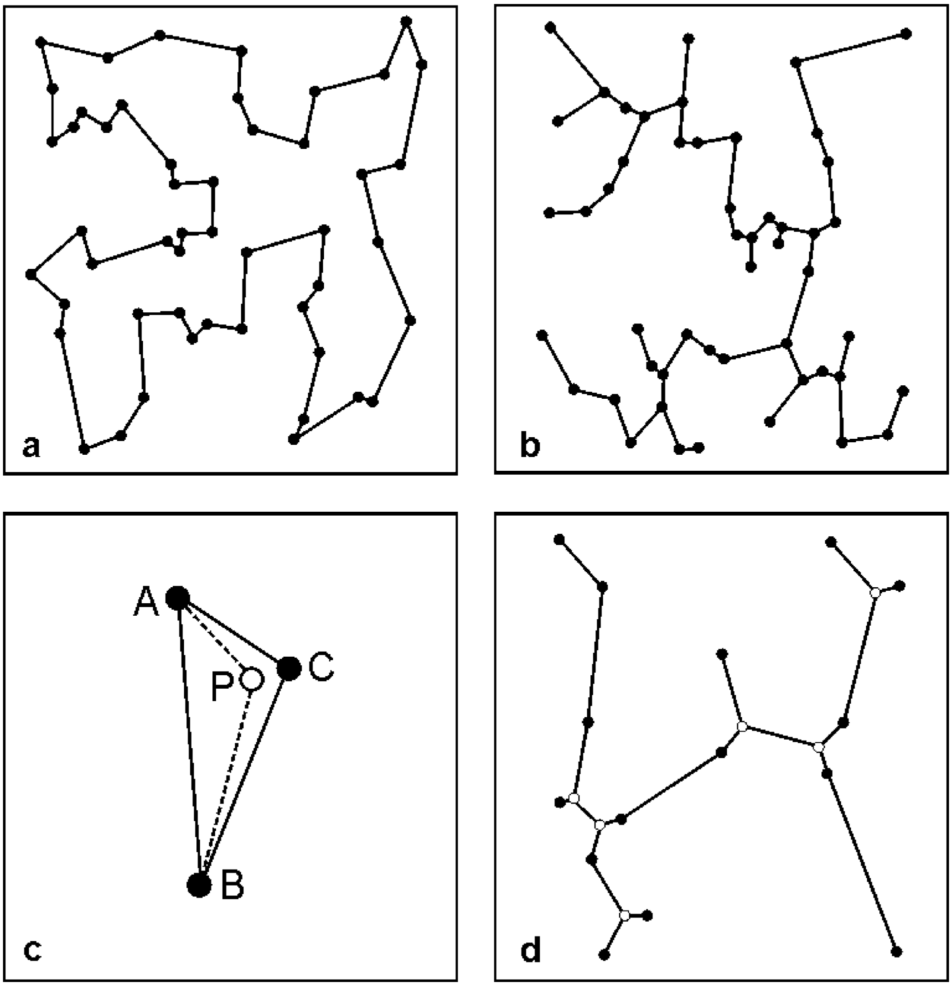

Fig. 1. (a) Shows the 50-node TSP array and (b) the 50-node MSTP array, used in Experiment 2, with the optimal solution in both cases as solid connecting edges. (c) Illustrates the Fermat point (P) for a simple Steiner Tree problem as an open circle, with the optimal solution shown by broken lines. (d) Shows the 15-node GSTP array used in Experiment 2, together with the interpolated points (open circles) and the optimal tree solution (broken lines) found by a computational algorithm.

3.1. Method

3.1.1. Participants

A total of 48 volunteer undergraduate students from the University of Adelaide, aged between 17 and 50 years ($M = 26.2$; SD = 9.2), participated in the experiment.

3.1.2. Problems

Participants provided solutions to each of two visually presented TSP, MSTP and GSTP problem instances. The second of each pair of instances was identical to the first, but was rotated through 90°. The TSP and MSTP instances are shown in Fig. 1a and b, and each consisted of 50 nodes, the coordinates of which were randomly selected from a uniform (rectangular) distribution. The GSTP instance is shown in Fig. 1d, and consisted of 15 nodes, the coordinates of which were also selected randomly from a uniform distribution. For each type of problem, there was also a practice instance. In the case of the TSP and the MSTP, this consisted of 15 randomly distributed nodes, and, for the GSTP, this consisted of seven randomly distributed nodes.

3.1.3. Design

Participants performed the three types of problem in a single session and in an order that was varied according to a Latin square design (Fisher & Yates, 1963), so that equal numbers of the 48 participants solved each type of problem in each of the six different possible orders. After solving a practice problem, participants solved the untransformed and transformed versions of that type of problem, before being presented with the practice instance for the next type of problem. In addition, participants completed Set 1 (practice) and Set 2 (test) of Raven's APM at a separate, prior session.

3.1.4. 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). There were no time constraints on performance, but participants were asked to complete the tests as quickly and as accurately as they could.

For the TSP, participants were asked to connect all the nodes in a continuous tour that visited each node once only, returned to the starting node, and was as short as possible. In the case of the MSTP, participants were asked to connect all the nodes by making a system of links (or paths) between them, using as many links as necessary, but only between existing nodes, and creating a path system with the shortest possible overall length. With the GSTP, participants were asked to connect all the nodes by making a system of links (or paths) with the shortest overall length. In this case, however, participants were told that the existing nodes did not have to be linked directly. If they wished, participants could create additional nodes, but they were required to create at least one additional node.

3.1.5. Procedure

Participants were tested individually or in small groups at well separated computers. Problems were presented, one at a time, in a 6×6-inch square in the centre of a standard 16-inch computer display. Participants could begin at any point by left-clicking on a node (or on a location where they wished to establish a new node) with the computer mouse. 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 (or created node) to select it, and then pressing the "Delete" key on the keyboard, participants could undo any links or nodes they had drawn. 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.

If a participant's completed solution was invalid (e.g. because not all TSP nodes had been connected), a warning message was posted on the screen and the participant was obliged to construct a valid solution before proceeding to the next problem.

3.2. Results

Participants' solution lengths were expressed as a proportion above the benchmark (PAB), or best known, solution length. Performance on the APM Set 2 was measured in terms of the number of correct items out of a total of 36.

3.3. Comparisons between performance on the TSP, MSTP, GSTP, and APM

The Pearson correlations between test 1 and test 2 performance on the TSP, MSTP, and GSTP were all positive (Pearson $r = 0.68$, 0.78 and 0.49, respectively; $P \leqslant 0.01$; $N = 48$ in all three cases). Accordingly, the two performance scores for each participant were combined to give an average performance score for each test. The means across all participants of these performance scores were 0.09, 0.06 and 0.12 (SDs = 0.08, 0.11 and 0.09) for the TSP, MSTP and GSTP, respectively.

Table 2 presents the Pearson correlations between performance on each pair of problems and between performance on each problem and on the APM. All the coefficients are modest to moderately strong, and highly significant statistically. There was a high degree of internal consistency among the six problem instances, with a value of 0.87 for Cronbach's α. When the performance measures for all three types of problem were summed, the Pearson correlation with APM was −0.51 ($P \leqslant 0.01$; $N = 48$).

3.4. Path complexity as a predictor of performance across the three types of problem

In an attempt to focus on a possible common basis for successful performance on the three optimisation problems, average measures of each participant's path complexity were calculated for each type of problem. Average path complexity correlated strongly with average PAB scores

Table 2
Pearson intercorrelations between PAB scores for the three types of problem and APM scores in Experiment 2

	TSP	MSTP	GSTP	APM	*N*
TSP		0.66	0.61	−0.46	48
MSTP			0.73	−0.44	48
GSTP				−0.46	48

for each problem (Pearson $r = 0.99$, 0.98 and 0.88 for the TSP, MSTP and GSTP, respectively; $P \leqslant 0.01$; $N = 48$, in each case). Measures of path complexity also showed reliable intercorrelations (Pearson $r = 0.64$, 0.53 and 0.69, respectively for the TSP:MSTP, TSP:GSTP and MSTP:GSTP comparisons; $P \leqslant 0.01$; $N = 48$, in each case). In addition, path complexity measures for each type of problem correlated reasonably well with APM scores (Pearson $r = -0.50$, -0.44 and -0.40 for the TSP, MSTP and GSTP, respectively; $P \leqslant 0.01$; $N = 48$, in each case).

4. Discussion

In contrast to the absence of individual differences in the TSP studies of MacGregor and Ormerod (1996) and MacGregor et al. (2000), the results of both experiments corroborate the finding by Vickers, Bovet et al. (in press), Vickers, Priess et al. (submitted for publication) that, provided visually presented TSP instances are sufficiently difficult, such tasks show reliable individual differences. Both experiments also replicate the finding by Vickers et al. (2001) that TSP performance is associated with scores on Raven's APM.

The results of the present study further indicate that there are consistent individual differences in two other, visually presented, optimisation tasks (the MSTP and the GSTP). These had been selected because they could be interpreted a priori as depending, respectively, more and less closely on the detection of least distances, which appears to be an important factor in performance on TSP tasks as well as in the perception of structure generally (Vickers, Bovet et al., in press; Vickers, Lee et al., in press; Vickers, Priess et al., submitted for publication). Performance on all three tasks was stable across two different instances of the same problem and was reliably correlated across successive instances. In addition, performance on all three tasks was strongly intercorrelated and had good correlations with APM scores.

These results suggest that performance on the three optimisation tasks and on the APM involves an important common factor. The question remains as to whether this factor is predominantly perceptual or cognitive. As shown by the high correlations between solution length and measures of path complexity, the TSP, MSTP, and GSTP all depend closely on the ability to utilise visual structure embodied in the pattern of least distances between the nearest neighbours in an array of similar elements. In turn, the path complexity measures of the extent to which participants make use of this structure correlate as highly with APM scores as do normalised measures of solution length. In our view, therefore, evidence from both experiments suggests that the common factor could be identified as the ability to perceive visual organisation or structure.

This conclusion is surprising, because, for each type of problem, there is an important, additional element, associated with the way the least distances are integrated into an overall minimal structure. In the case of the GSTP, in particular, it seemed a priori most likely that this element could be interpreted as an abstract, cognitive calculation. Because intelligence is frequently identified with the ability to organise cognitive processes in a flexible and efficient manner, we expected that performance on the GSTP would show the highest correlation with APM scores and that performance on the MST would show the lowest correlation with the APM. However, the three correlations are virtually identical.

A possible explanation for the high degree of intercorrelation between the three types of problem, as well as for their reliable correlations with both solution lengths and with APM scores, is that, although not incorrect, the analysis of nearest neighbour relations, proposed by Vickers, Priess et al. (submitted for publication), is less comprehensive than is theoretically possible as an account of the visual perception of structure. Specifically, any planar array of randomly positioned nodes defines a geometrical construction, known as a proximal polygon or Voronoi diagram (O'Rourke, 1994). The lines that are equidistant from each pair of nodes constitute the edges of two half-planes. The Voronoi diagram is composed of convex polygonal cells, formed from the intersections of these half-planes and each polygon can be interpreted as the region of influence around each of a given set of nodes. Each of the nodes in an array belongs to a unique Voronoi polygon, so that the Voronoi diagram contains all of the proximity information defined by the array.

Whereas the computation of TSP and GSTP solutions are intractable problems, the calculation of Voronoi diagrams is computationally feasible. This is extremely useful, because Voronoi diagrams (and a complementary construction, termed Delaunay triangulation) provide efficient methods of finding the nearest neighbours and the MSTs of a given set of points (O'Rourke, 1994). In addition, such constructions are being actively investigated as ways of finding rapid, near-optimal solutions to TSPs and GSTPs (Beasley & Goffinet, 1994).

The relevance of these developments for the study of human visual perception and cognition is that it is not implausible to suppose that the nervous system encodes information about the relative positions of an array of elements in a way that resembles Voronoi diagrams in being maximally comprehensive and efficient. If so, then the TSP, MST and GSTP tasks may all make use of this information and thus may not differ as much as we originally assumed.

So far as the relationship between performance on each of these three tasks and on the APM is concerned, the simplest explanation is that information about relative position provides a means of detecting structure in an array. In particular, it has been found to be effective in detecting and differentiating various forms of symmetry or transformational structure (Vickers, 2001). Consistent with the interpretation that the visual perception of structure plays an important common role in the APM and in the three optimisation tasks, DeShon, Chan, and Weissbein (1995) concluded that performance on the APM was determined by two general strategies: an analytic one, based on propositional representations, and a visuo-spatial process that appeared to be involved in recognising the nature of various transformations and extrapolating from them.

5. Conclusions

The results from the present experiment suggest a number of conclusions. First, they show that there are stable and reliable individual differences in performance on visually presented TSP tasks, and that these differences extend to other visually presented optimisation tasks. Second, the results confirm the findings of a previous study that showed a relationship between TSP performance and the APM. Third, they indicate that this relationship is shared by two other related, but distinguishable, tasks involving the detection or construction of minimal structures. Fourth, the results are consistent with the conclusion that the process common to the APM and to the detection of minimal structures is the detection and use of powerful information regarding the relative positions of stimulus elements.

At the same time, the results raise the possibility that either perception is even more intelligent than has been generally assumed or that cognition is more perceptually based. One way to investigate this would be to combine the study of solving visual optimisation problems with an examination of performance on the APM and on a 'purely perceptual' task, such as the detection of structure in Glass patterns (Glass, 1969), as well as of the ability to solve non-visual optimisation problems, such as the so-called Secretary Problem (Gilbert & Mosteller, 1966). Current research is being directed to this question.

Acknowledgements

The research on which this article is based was supported by an Australian Research Council Grant (DP0210851) to D. Vickers.

References

Ahuja, R. K., Magnanti, T. L., & Orlin, J. B. (1993). *Network flows: theory, algorithms, and applications*. New York: Prentice Hall.

Beasley, J. E., & Goffinet, F. (1994). A Delaunay triangulation-based heuristic for the Euclidean Steiner problem. *Networks*, *24*, 215–224.

Chi, M. T. H., Glaser, R., & Farr, M. J. (1988). *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum.

Cohen, J. (1994). The Earth is round ($P < 0.05$). *American Psychologist*, *49*, 997–1003.

DeShon, R. P., Chan, D., & Weissbein, D. A. (1995). Verbal overshadowing effects on Raven's Advanced Progressive Matrices: evidence for multidimensional performance determinants. *Intelligence*, *21*, 135–155.

Duncan, C. P. (1959). Recent research on human problem solving. *Psychological Bulletin*, *56*, 397–429.

Edwards, W., Lindman, H., & Savage, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, *70*, 193–242.

Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, *47*, 273–305.

Fisher, R. A., & Yates, F. (1963). *Statistical tables for biological, agricultural, and medical research* (6th ed). Edinburgh: Oliver & Boyd.

Gilbert, J. P., & Mosteller, F. (1966). Recognising the maximum of a sequence. *Journal of the American Statistical Association*, *61*, 35–73.

Glass, L. (1969). Moiré effect from random dots. *Nature*, *243*, 578–580.

Golden, B., Bodin, L., Doyle, T., & Stewart, W. (1980). Approximate traveling 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.

Hunter, J. E. (1997). Needed: a ban on the significance test. *Psychological Science*, *8*, 3–7.

Hwang, F. K., Richards, D. S., & Winter, P. (1992). *The Steiner tree problem*. Amsterdam: Elsevier.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association, 90 (430)*, 773–795.

Lawler, E. L., Lenstra, J. K., Rinooy Kan, A. H. G., & Schmoys, D. B. (1985). *The traveling salesman problem: a guided tour of combinatorial optimization*. Chichester, UK: Wiley.

Lee, M., & 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 and Psychophysics*, *62*, 226–228.

Lindley, D. V. (1972). *Bayesian statistics: a review*. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Lovett, M. C. (2002). Problem solving. In H. Pashler, & D. Medin (Eds.), *Stevens' handbook of experimental psychology* (3rd ed.), Vol. 2: Memory and cognitive processes (pp. 317–362). New York: John Wiley & Sons.

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

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.

O'Rourke, J. (1994). *Computational geometry in C*. New York: Cambridge University Press.

Press, W. H. (1992). *Numerical recipes in FORTRAN: the art of scientific computing* (2nd ed). Cambridge, UK: Cambridge University Press.

Raven, J. C., Court, J. H., & Raven, J. (1988). *Manual for Raven's progressive matrices and vocabulary scales. Section 4: Advanced progressive matrices*. London: Lewis.

Ray, W. S. (1955). Complex tasks for use in human problem-solving research. *Psychological Bulletin*, *52*, 134–149.

Resnick, L. B., & Glaser, R. (1976). Problem solving and intelligence. In L. B. Resnick (Ed.), *The nature of intelligence* (pp. 205–230). Hillsdale, NJ: Lawrence Erlbaum Associates.

Sivia, D. S. (1996). *Data analysis: a Bayesian tutorial*. Oxford: Clarendon Press.

Smith, M. U. (1991). *Toward a unified theory of problem solving*. Hillsdale, NJ: Erlbaum.

Stein, D. L. (1989). Spin glasses. *Scientific American*, *261*, 36–42.

Sternberg, R. J. (1982). Reasoning, problem solving, and intelligence. In R. J. Sternberg (Ed.), *Handbook of intelligence* (pp. 225–307). New York: Cambridge University Press.

Vernon, P. A., & Strudensky, S. (1988). Relationships between problem-solving and intelligence. *Intelligence*, *12*, 435–453.

Vickers, D. (2001). Towards a generative transformational approach to visual perception. *Behavioral and Brain Sciences*, *24*, 707–708.

Vickers, D., Bovet, P., Lee, M. D., & Hughes, P. The perception of minimal structures: performance on open and closed versions of visually presented Euclidean Traveling Salesperson problems. *Perception* (in press).

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

Vickers, D., Preiss, A. K., & Hughes, P. The role of nearest neighbours in the perception of structure and motion in dot patterns (submitted for publication).