THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY
2006, 59 (4), 783–800

Psychology Press
Taylor & Francis Group

# It looks easy! Heuristics for combinatorial optimization problems

**Edward P. Chronicle**
*University of Hawaii at Manoa, Honolulu, HI, USA*

**James N. MacGregor**
*University of Victoria, Victoria, British Columbia, Canada*

**Thomas C. Ormerod and Alistair Burr**
*Lancaster University, Lancaster, UK*

Human performance on instances of computationally intractable optimization problems, such as the travelling salesperson problem (TSP), can be excellent. We have proposed a boundary-following heuristic to account for this finding. We report three experiments with TSPs where the capacity to employ this heuristic was varied. In Experiment 1, participants free to use the heuristic produced solutions significantly closer to optimal than did those prevented from doing so. Experiments 2 and 3 together replicated this finding in larger problems and demonstrated that a potential confound had no effect. In all three experiments, performance was closely matched by a boundary-following model. The results implicate global rather than purely local processes. Humans may have access to simple, perceptually based, heuristics that are suited to some combinatorial optimization tasks.

How humans solve complex problems or make decisions in the face of uncertain information has long held the attention of psychologists. Debate has typically focused upon divisions between proponents of normative systems such as Bayesian inference and formal logic as the basis of human thought (e.g., Braine & O'Brien, 1991; Oaksford & Chater, 1994; Rips, 1994) and those who describe human thinking in terms of the operation of heuristics that guide the discovery of plausible solutions. The latter approach can be further divided between proponents of heuristics determined by properties of the space of possible solutions, such as means–ends analysis (e.g., Newell & Simon, 1972), representativeness, and availability (Tversky & Kahneman, 1973) and structural analogy (Anderson, 1993), and proponents of heuristics determined by the ecological

Correspondence should be addressed to Edward P. Chronicle, Department of Psychology, University of Hawaii at Manoa, 2340 Campus Road, Honolulu, HI 96822, USA. Email: chronicl@hawaii.edu
This paper is dedicated to the late Douglas Vickers (University of Adelaide, Australia), who reviewed the first version of this paper. We are most grateful to Sandy Aitkenhead and Peter Naish of the Open University, UK, for facilitating the testing of volunteers. We thank Ranbir Prihar for developing the programs for generating problems and finding optimal open paths. Iris van Rooij and Ulrike Stege of the University of Victoria, Canada, kindly helped us to confirm optimal paths for some problems. Cheryl Matthews of the University of Maryland, Baltimore County, assisted with the preparation of materials.

 
http://www.psypress.com/qjep DOI:10.1080/02724980543000033

environment of task performance, in particular the evolutionary characteristics of the problem solver (e.g., Gigerenzer, 2001). Progress in deciding between alternative accounts of the fundamental basis of human thought has been slow, partly because each of the three approaches has equal a priori applicability to the tasks that are generally studied, rendering the derivation of unique predictions extremely difficult.

In this paper we consider human performance on a class of optimization problems that differs from the kinds of decision-making and problem-solving task typically studied by psychologists, in that there is no known computational procedure that guarantees the discovery of a correct (i.e., optimal) solution in a reasonable time. The absence of an algorithmic system simplifies the psychological explication of performance on optimization problems in some respects, since there can be no normative account of human performance. However, alternative accounts of human performance on optimization problems have proposed either strategies for limiting the problem space (e.g., van Rooij, Stege, & Schactman, 2003; Vickers, Lee, Dry, & Hughes, 2003b) or heuristics that capitalize upon the preexisting perceptual competencies of the problem solver (e.g., Graham, Joshi, & Pizlo, 2000; MacGregor & Ormerod, 1996; MacGregor, Ormerod, & Chronicle, 2000; Ormerod & Chronicle, 1999). The aim of this paper is to contribute to ongoing debate about the fundamental basis of human thought by testing predictions from alternative heuristic accounts of performance on two specific optimization problems: the Euclidian version of the travelling salesperson problem (E-TSP), and a variant of that problem.

The E-TSP is an optimization problem that requires finding the shortest tour through a set of locations in the plane, returning to the starting location. A simple example of a TSP, together with its solution, is shown in Figure 1. TSPs are common in industrial and management settings and may underlie practical problems as diverse as circuit board drilling, X-ray crystallography, and the laying of ducting (Krolak, Felts, & Marble, 1971; Sangalli, 1992). Some examples that involve laser movements in chip manufacture may represent travelling salesman problems of up to a million nodes (Sangalli, 1992). Such problems present a challenge because of their combinatorial properties: The number of possible tours is given by $(L-1)!/2$, where L is the number of locations. As L increases it becomes unfeasible to find the shortest tour by computing all possible solutions. Rather, heuristic algorithms have been developed which produce acceptable solutions at low computational cost. Since the mid-1990s, cognitive psychologists have become increasingly interested in optimization problems such as the TSP. They enable straightforward laboratory study of an important class of human problem-solving heuristics: those that enable us to circumvent the impossible processing demands of finding a gold-standard solution, but still to perform reliably well.

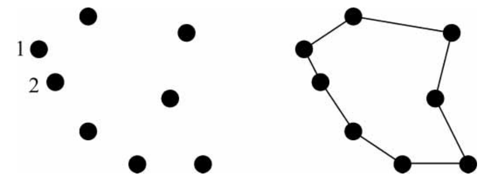


**Figure 1.** *A simple Euclidean TSP, left, with its optimal tour, right. Numbered locations are referred to in the text.*

We have previously demonstrated (MacGregor & Ormerod, 1996; MacGregor, Ormerod, & Chronicle, 1999) that untrained adults reliably outperform simple computer-based construction algorithms in problems where $L \leq 60$. For example, one such construction algorithm generates a tour by moving from its starting location to the nearest neighbour location, and so on for each subsequent location. Humans typically construct tours that are significantly closer to optimal than this algorithm. It should be noted that much more sophisticated computational algorithms are now available and are able to find solutions of excellent quality to, for example, a problem representing all of the 15,112 cities in Germany (Applegate, Bixby, Chvátal, & Cook, 2001). Such algorithms typically run on networks of computers and work by mathematically complex iterative tour improvement. Nonetheless, problems with between 20 and 60 locations still represent significant computational challenges. For example, a modern desktop computer finding 100,000 tours per second would still require a little over 19,000 years to find the optimal solution to even a 20-location problem by exhaustive search. The results of MacGregor and Ormerod (1996) and MacGregor et al. (2000) suggest that naïve participants overcome these challenges by employing a particular *convex-hull* heuristic to construct—rapidly and with little apparent effort—optimal or close-to-optimal tours. In order to explain the convex-hull heuristic, it is important to note that for any TSP, there is a set of locations that lie on a convex hull around the boundary of the problem. It may aid understanding to imagine all the locations in Figure 1 marked with pins, and an elastic band stretched around the problem. The pins touching the elastic band are the locations falling on the convex hull. We argue that the convex hull is perceptually salient because it represents the figural boundary—or global shape—of the TSP. Mathematically, the optimal tour does not intersect itself, from which it follows that it must connect the convex-hull points in order of adjacency (Quintas & Supnick, 1965). A heuristic based on following the convex hull is therefore entirely sensible for humans to adopt, as it will permit (but not guarantee) the discovery of the optimal tour. Ormerod and Chronicle (1999) argue that the psychological implementation of this heuristic capitalizes upon the natural perceptual processes of object identification, notably the identification of boundaries around objects (cf. Marr, 1983). In this respect, the convex-hull heuristic that we propose bears comparison with a class of evolutionary heuristics described by Gigerenzer (2001) as being "fast and frugal". In particular, it has the attributes of being a simple and straightforward account of a higher level cognitive ability that takes input from complex and massively parallel lower level processing (Gigerenzer & Todd, 1999).

A number of empirical findings have been reported that support our hypothesis of a perceptually based solution process guided by the convex hull. MacGregor and Ormerod (1996), for example, reported the following four findings. First, complexity varied with the number of interior points, not total points, consistent with the proposition that incorporating boundary points into a tour imposes little or no cognitive load. Second, boundary points were connected in order of adjacency, and self-intersecting tours were extremely rare. Third, solutions exhibited few "indentations", reflecting an apparent preference not only to connect boundary points in sequence, but to connect them to each other (an indentation occurs whenever at least one interior point is connected between two adjacent boundary points). Fourth, there was little evidence of individual differences in the quality of solutions, consistent with a fundamental, low-level process.

Additional evidence consistent with the convex-hull hypothesis has come from subsequent studies. Performance has been found to be superior when the interior points are closer to the boundary of the figure than when they are more distant (MacGregor et al., 1999). A model based on a convex-hull approach has been shown to produce tours with lengths that are within the range of human solutions (MacGregor et al., 2000). Ormerod and Chronicle (1999) demonstrated that participants were capable of judging the optimality of exemplars of TSP tours presented for a maximum of 2,000 ms on a computer screen. Furthermore, RTs for optimality judgements about a tour exemplar could be primed by short (100-ms) presentations of the set of locations alone, but not by long presentations (>500 ms). Ormerod and Chronicle (1999) argued that only global figural properties of TSPs were likely to be available from short primes, and that these were sufficient to make fast and accurate judgements about tours. More recently, MacGregor, Chronicle, and Ormerod (2004, Exp. 3) manipulated the figural properties of TSPs such that the convex hull had low or high perceptual salience. Participants' tours were significantly closer to optimal in the high salience condition than in the low salience condition.

Other laboratories have also reported evidence consistent with the convex-hull hypothesis. Studies agree that unaided human performance on visually presented problems is of a high quality (Graham et al., 2000, Vickers, Butavicius, Lee, & Medvedev, 2001), frequently outperforming computer-based heuristic algorithms (Graham et al., 2000). While the high quality of performance does not represent specific support for the convex-hull hypothesis, the fact that good solutions are produced quickly and with apparent ease argues in favour of some common, relatively automatic, underlying process. That the process may be perceptually based and related to Gestalt principles is further supported by the results of Vickers et al. (2001), who compared performance under two different instructional conditions. One group (O group) was instructed to draw the shortest tour, and the other (G group) to produce a tour that was "most natural, attractive or aesthetically pleasing" (p. 36). Although there were a number of differences in average performances under the two conditions, participants in both groups were equally likely to produce optimal tours; common tours within the O group were often produced by the G group also, and the shortest tour was approximately equally likely to occur in either group. When asked to describe their strategy, 7 of the 36 participants reported that "they began by forming a roughly circular perimeter around the points", while 8 stated that the tour "was immediately and intuitively obvious" to them (p. 40).

Our emphasis on the convex-hull heuristic as a primary cognitive tool for humans attempting TSPs has not gone unchallenged, however. Vickers et al. (2003b) have reported that increasing the number of points on the convex hull can lead to decrements in performance. This is a surprising result, and the extent to which it rests on particular properties of the stimulus materials used is unknown. A number of researchers have also suggested that the perception of local features or properties of the TSP stimulus may be more important than the global-level extraction of the convex hull. Van Rooij et al. (2003), for example, note that optimal TSP tours are known to have no intersections, and that tours with no intersections must connect convex-hull points in order of adjacency (Quintas & Supnick, 1965). From this mathematical given, van Rooij et al. (2003) suggest that a strategy of avoiding crossings may be primary, and that convex-hull following is a necessary consequence. This suggestion is made plausible by the empirical demonstration that only about 5% of tours constructed by naïve participants contain crossings. A crossing avoidance heuristic is comparable with the kinds of problem−space-limiting heuristics offered by Newell and Simon (1972) and Tversky and Kahneman (1973), though it differs somewhat in being problem specific rather than general to a wide range of problems. A similar crossing-avoidance explanation, which also incorporates elements of a nearest-neighbour heuristic, has been proposed by Vickers et al. (2003b). They describe their approach in terms of a "local-to-global" strategy in contrast to the "global-to-local" convex-hull hypothesis.

While the crossing avoidance heuristic is a reasonable suggestion in principle, we have argued that it is unlikely to be cognitively implementable in the form of exhaustive search, since this would impose extremely high demands on processing capacity. Further, we have demonstrated that tours produced solely by a strategy of crossing avoidance are significantly longer than the tours of adults, and also that the convex-hull heuristic better predicts observed tour length than does the crossing avoidance heuristic (MacGregor et al., 2004). Nonetheless, while these results would appear to rule out crossing avoidance as the sole explanation of human E-TSP performance, it is conceivable that a crossing avoidance strategy in concert with other local processes might be able to explain the various empirical findings. For this reason, we conducted three experiments designed to test between the convex-hull explanation and generic local-to-global approaches.

An intuitively appealing method of examining the convex-hull heuristic is to vary the problem instructions such that it is more or less applicable. One way in which this may be achieved is to remove the requirement that TSP tours return to the location at which they started, a suggestion originally made by MacGregor and Ormerod (1996). It is convenient to refer to such problems as "open" E-TSPs, and the standard versions as "closed" E-TSPs. It can be demonstrated that open problems have a similar degree of computational intractability to closed problems. For clarity, we henceforth use the term "path" to refer to solution attempts with open problems, reserving the term "tour" for closed problems.

For open problems, a convex-hull heuristic is inapplicable, since shortest paths do not necessarily follow the convex hull, as is the case with closed problems. Thus, if a perceptually based convex-hull heuristic underlies human performance, performance should be degraded on open problems. If, on the other hand, performance is guided by purely local decisions, such as crossing-avoidance, there is no apparent reason why removing the requirement to complete a closed tour would have any effect.

Vickers, Bovet, Lee, and Hughes (2003a, p. 875) attempted such a manipulation as follows: "Instructions for Open paths emphasized that the path could begin at any node (town), but had to finish at a different node (town)". Participants in this condition produced longer paths than those given the standard instruction (that the tour must start and end at the same location). Unfortunately, a number of methodological difficulties make this result difficult to interpret. First, consider the numbered locations in Figure 1.

If a participant decided to begin a path at Location 1 and finish it at Location 2, the nearest neighbour point, the convex hull is only minimally degraded. Clearly, start and finish locations may vary widely both within and between participants, and it is not therefore possible to know how reliably the open path instructions used by Vickers et al. (2003b) interrupted use of the convex-hull heuristic. Second, three aspects of the experimental procedure give cause for concern. Participants constructed solutions under both closed and open instructions, and although order of instruction was counterbalanced, it is not clear that carry-over effects were symmetrical, given the strong intrinsic tendency to use the convex hull if possible (MacGregor et al., 2004). Participants were also given feedback as to the quality of each path/tour, in the form of an optimal solution to each exemplar with their own path/tour superimposed in a different colour. This may have obscured the operation of naïve heuristics through learning additional or alternative strategies on the basis of the optimal solutions presented. Third, allowing participants to start and end on nodes of their choice introduced systematic biases against the open condition. This is because, in closed tasks, any starting-point must lie on the optimal tour. In contrast, in open versions, only two possible starting-points will lie on the optimal path (the points at each end of this path). Furthermore, open versions under these instructions are more complex than closed versions: With no start and end point provided, open paths have $L!/2$ possible solutions, $L$ times as many as closed versions. These differences could have made it more difficult to generate optimal paths in the open condition and represent confounding factors.

Finally, one of the dependent measures used by Vickers et al. (2003a) could have introduced a serious bias against the open condition. The measure was the difference between observed path length and the optimal path length, expressed as a percentage of the optimal, or "percent above optimal" (PAO). Normally, this is a useful measure of relative performance, one that has been reported both in the psychological and the operations research literatures (Golden, Bodin, Doyle, & Stewart, 1980; Graham et al., 2000; MacGregor & Ormerod, 1996). However, in the special case of comparing performance on closed and open versions of the same instance, a problem may arise with PAO. Consider, for example, a case where the length of the shortest open path is $x$, and the shortest closed tour is identical, except for the final arc, of length $y$, that connects the start and end points. The optimal closed path will then be of length $x + y$. Consider now the case of two identical suboptimal paths in the open and closed versions (except for the final closing arc), and that this path exceeds the optimal by $z$. Under these (quite possible) circumstances, PAO will be $z/(x + y)$ for the closed version and $z/x$ for the open version. That is, under these circumstances, the same quality of performance on closed and open versions of an instance always results in a better PAO score for the closed version.[1]

The present paper contrasts human performance on two versions of the same set of TSPs. We asked participants to draw either closed tours, in which the specified starting location was returned to, or open paths, in which the specified starting and finishing locations were always on opposite sides of the problem. The convex-hull hypothesis predicts that the instruction to produce a closed tour should give rise to closer-to-optimal solutions, as it is only with such an instruction that the heuristic can operate effectively. The manipulation that we implemented has two important features. First, participants in both the closed and open conditions must use start and end points that lie on the optimal path/tour, thereby removing a potential bias against open versions. Second, it provides for a stringent test of the convex-hull hypothesis, because the open problems that we employ are combinatorially simpler than the closed problems (they have $[L - 1]/2$ times fewer possible solutions) and should—on that basis alone—be easier.

## EXPERIMENT 1

Experiment 1 implemented the manipulation of start and finish points outlined above. Either participants were required to start and finish on the left-most location of the problem (closed problem type), or they started at the left-most location and finished at the right-most location (open problem type). We expected that the closed conditions would give rise to nearer-to-optimal solutions than would open conditions.

In addition, previous work suggested an orthogonal manipulation of the number of locations interior to the convex hull (MacGregor et al., 2000, 2004). We have proposed that, in closed conditions, the convex hull can always be derived perceptually by problem solvers, regardless of the number of locations falling on the convex hull. Varying the number of convex-hull locations does not therefore alter the difficulty of TSPs a great deal. The absolute number of interior points, however, should determine the number and complexity of indents from the convex hull. Consistent with this prediction, MacGregor et al. (2004) found that the quality of participants' tours decreased significantly as the number of interior points increased; however, the main effect of total number of points was not significant.

### Method

#### *Design*

Two variables were manipulated. The first, problem type, was assigned between groups and had two levels: open and closed. The second, interior points, was assigned within groups and also had two levels: one interior point and six interior points. For each level of interior points, six problem exemplars were produced. The order of presentation of the resultant 12 problems was pseudorandomized for each participant. The dependent variable was the length of path/tour produced.

#### *Participants*

A total of 50 participants volunteered for the experiment. All were adult cognitive psychology students participating in a residential summer school for undergraduates. No identifiers were collected, so the exact demographics of the participants are unknown. All participants declared themselves naïve to the TSP and had not been assigned any reading or study materials relevant to optimization problem solving during their cognitive psychology course.

#### *Materials*

TSPs were constructed by randomly generating rectangular coordinates under the following constraints. The coordinate space extended from 20

[1] We are grateful to the late Douglas Vickers for pointing out this potential problem with PAO.

to 620 pixels in the *x* direction and from 10 to 460 pixels in the *y* direction. Boundary points were selected from a rectangular border area extending inwards from the edges of the coordinate space by 60 pixels in the *x* direction and 45 pixels in the *y* direction. Potential boundary points were resampled if they fell within the convex hull. Interior points were randomly sampled from the range of coordinates falling inside the outer rectangular border. Six problem exemplars were produced with one interior point, and six exemplars with six interior points. All problem exemplars had a total of 12 points. Coordinates were then plotted with standard software, as circular dots, diameter 3 mm, within a 120 × 90-mm grid. Paper booklets of problems in pseudorandom order, prefaced by standardized instructions, were prepared for each participant. Examples of problems in each condition of the experiment are shown in Figure 2.

*Procedure*

Participants were tested either individually or in pairs, in a quiet room. Participants were asked to start and finish at the point(s) specifically indicated on each problem and to pass through each dot only once, completing the shortest route possible. They were encouraged simply to draw the path/tour that looked right to them. They were also asked to write "1" next to the first line that was drawn to indicate the path/tour direction. There was no time limit for completion of the booklets. Participants typically completed working on their booklets in 10–15 minutes.

## Results and discussion

For each path/tour produced by each participant, the sequence of points visited was coded against a template, and the total path/tour distance (ignoring any curvature of drawn lines) was derived trigonometrically. Of the 600 paths/tours, 8 involved incorrect starting or end points and were excluded from analysis.

The first analysis compared the number of paths/tours in each condition that were optimal. To do so, optimal solutions for each problem exemplar were found either by TRAVEL (Boyd, Pulleyblank, & Cornuéjols, 1987), a suite of algorithms for analysis of TSPs, or by exhaustive search. (For open versions, the optimal path was defined as the shortest possible path connecting the L points from the specified start point to the specified end point.)

Of the total 592 paths/tours, 231 (39%) were optimal: 155 (26%) in the closed condition, 76 (13%) in the open. For analysis, the results were expressed in terms of the number of optimal tours per participant. For the closed condition, the mean number of optimal tours per participant were 4.44 and 1.76 (out of a possible 6) in the one and six interior points conditions, respectively. For the open condition, the corresponding means were 1.68 and 1.36 (see Figure 3a). The results were analysed using a mixed analysis of variance, with number of interior points as the within factor and closed or open problem type as the between factor. The results indicated that all three effects were significant. The interaction effect indicated that the superiority in performance for closed tours over open paths was greater in the one than in the six interior points condition, $F(1, 48) = 68.37$, $MSE = 0.51$, $p < .001$. There was a main effect of interior points: Performance was superior in the one interior point condition, with an average of 3.06 optimal solutions, than in the six interior points condition, with 1.56, $F(1, 48) = 110.48$, $MSE = 0.51$, $p < .001$. There was also a main effect of problem type, closed problems resulting in more optimal solutions than open (3.10 vs. 1.52), $F(1, 48) = 38.66$, $MSE = 1.61$, $p < .001$.

The results for the closed condition supported previous findings, that performance is superior for problems with fewer interior points (MacGregor et al., 2004; MacGregor & Ormerod, 1996). The results also supported the present hypothesis, that performance is better in closed than in open conditions. However, the effect was largely limited to the one interior point condition. This may reflect the nature of the effect, or it may be a result of the relative insensitivity of the present performance measure. As task difficulty increases, the number of people

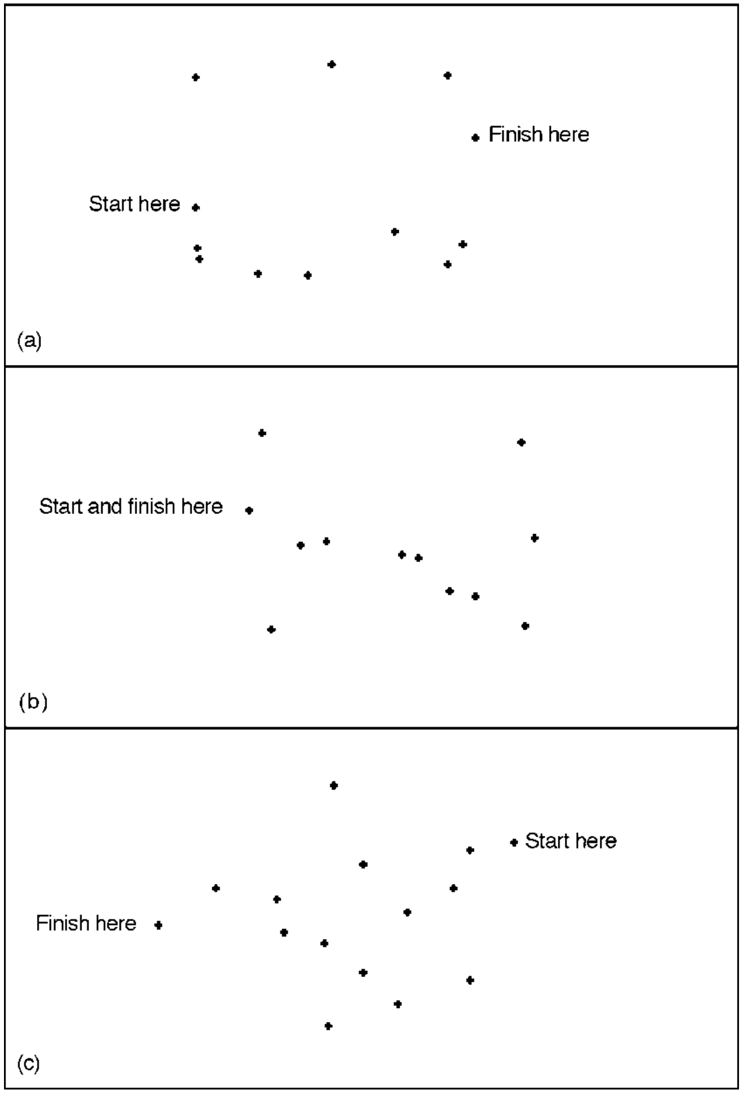


**Figure 2.** *Examples of the TSPs used. Panels (a) and (b) relate to Experiment 1; (a) shows a problem with $L = 12$ and 1 interior point, requiring an open path; (b) shows a problem with $L = 12$ and 6 interior points, requiring a closed tour. Panel (c) relates to Experiment 2, and shows a problem with $L = 15$ and 9 interior points, requiring an open path starting at the rightmost location.*

finding the optimal solution decreases, so that at some level of task difficulty there will be a floor effect, and the proportion of optimal solutions will cease to be a useful indicator. We may have reached this level in the six interior points condition in the present case. What is required is a measure that reflects the variations in quality among nonoptimal solutions. There are a number of possible measures that may achieve this. The first is simply the lengths of solutions themselves. However, because closed tours involve one arc more than open paths, they will tend to be longer for this reason alone. A second possible measure is the difference between the observed path/tour length and the optimal length for that problem. Again, however, since

A

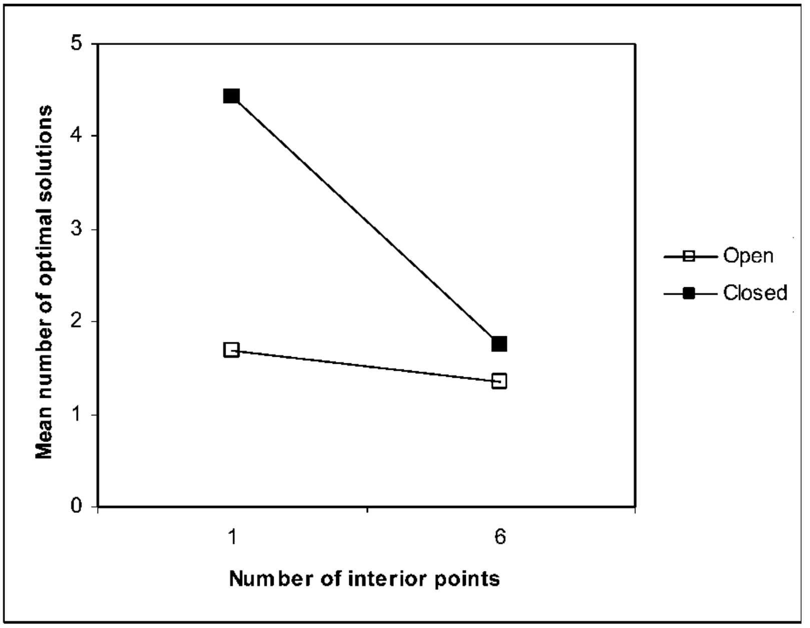


B

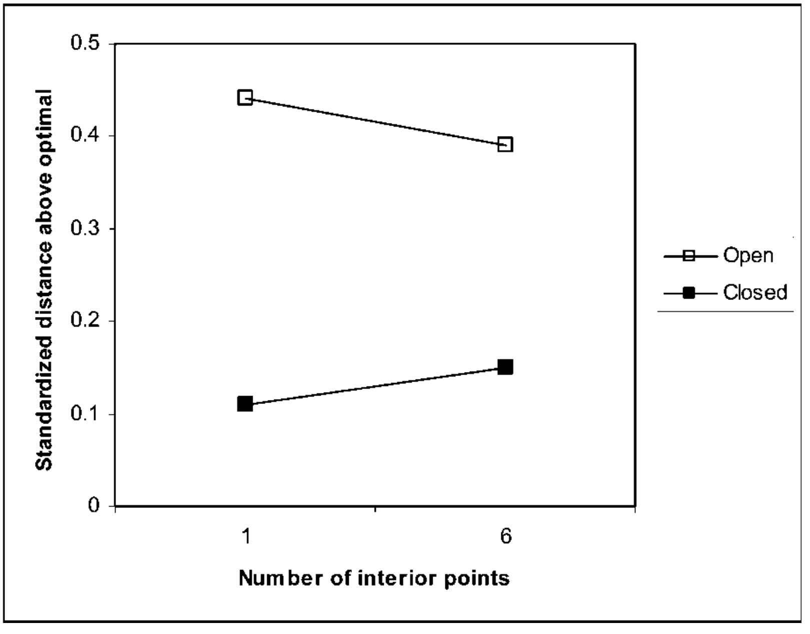


**Figure 3.** *A. Mean number of optimal path/tours by problem type and number of interior points in Experiment 1. B. Path/tour lengths in Experiment 1 (expressed as standardized distance above optimal, $Z_{DAO}$) by problem type and number of interior points.*

optimal closed tours were longer than the corresponding open paths, this may introduce a degree of bias. A third option is the percentage above optimal (PAO). However, for reasons presented in the Introduction, this may favour closed over open tasks. A fourth is the *standardized distance above optimal* ($Z_{DAO}$): the difference from optimal divided by the standard deviation of all possible path lengths (Vickers et al., 2001). We chose this latter measure, where a score of 0 represents the best possible performance.

Scores were calculated as follows. For each problem in the open and closed conditions, the standard deviation of a random sample of 10,000 solutions was found. These were used as estimates of the corresponding population standard deviations. The optimal distances were subtracted from each participant's score, to provide a distance above optimal. These distances were then divided by the appropriate standard deviation, and the results averaged across the 6 one interior point problems and the 6 six interior points problems, yielding two scores for each of the 50 participants.

As before, a mixed analysis of variance was performed, with results similar, but not identical, to those using the number of optimal solutions. For the closed condition, the mean standard distances above optimal were 0.11 and 0.15 for the one and six interior points conditions, respectively. For the open condition, the corresponding results were 0.44 and 0.39 (see Figure 3b). Again, the main effect of problem type was significant, $F(1, 48) = 41.12$, $MSE = 0.05$, $p < .001$, reflecting significantly superior performance in the closed condition. However, the main effect of number of points was nonsignificant, $F(1, 48) = 0.01$, while the interaction effect was borderline, $F(1, 48) = 3.67$, $MSE = 0.01$, $p = .06$. The results based on $Z_{DAO}$ confirmed our main finding, that performance in the closed condition was superior to that in the open. In the closed condition, tours were only 3.17% longer than the optimal on average, compared with 9.57% for open paths. The trend of the interaction effect, while not quite significant, was consistent with that obtained using number of optimal paths: that while performance in the closed condition tended to deteriorate with increasing interior points, performance in the open condition did not. Nevertheless, open paths shared one consistency with closed tours, in that the incidence of crossed arcs was uniformly low. For the open paths, there were 7 cases of crossed lines (2.4% of all paths); for closed tours there were 12 (4%). This difference was not significant (Fisher's Exact Test, $p = .35$).

We next examined whether the convex-hull model (MacGregor et al., 2000) provided a useful first approximation to whatever heuristic(s) may be involved in the construction of closed tours. Figure 4 compares the lengths of the model's solutions with those of the participants, across the present 12 problems. For the majority of problems, it can be seen that the model's solution closely matches the participants' modal solution.

The present results have a number of limitations. One is that the problems were generated systematically to have certain properties, and this nonrandom nature of the stimuli may have influenced the findings. Clearly, this limits the generalizability of the findings. A second is that the starting-point in each case was located at the left side of the display. As most Western participants are accustomed to working from left to right (in reading or writing), performance in the open condition of Experiment 1 was potentially influenced by the direction in which paths were constructed. Experiment 2 was designed to avoid these limitations.

## EXPERIMENT 2

Experiment 2 retained the problem type manipulation of Experiment 1 and orthogonally varied the starting-point of the problem so that half of the paths/tours started at the right of the page, and half at the left. In addition, Experiment 2 increased the number of locations to $L = 15$, with the number of interior points set at 9. The reasons for this variation were twofold. First, it was important to ensure that the effect of problem type seen in Experiment 1 was not an

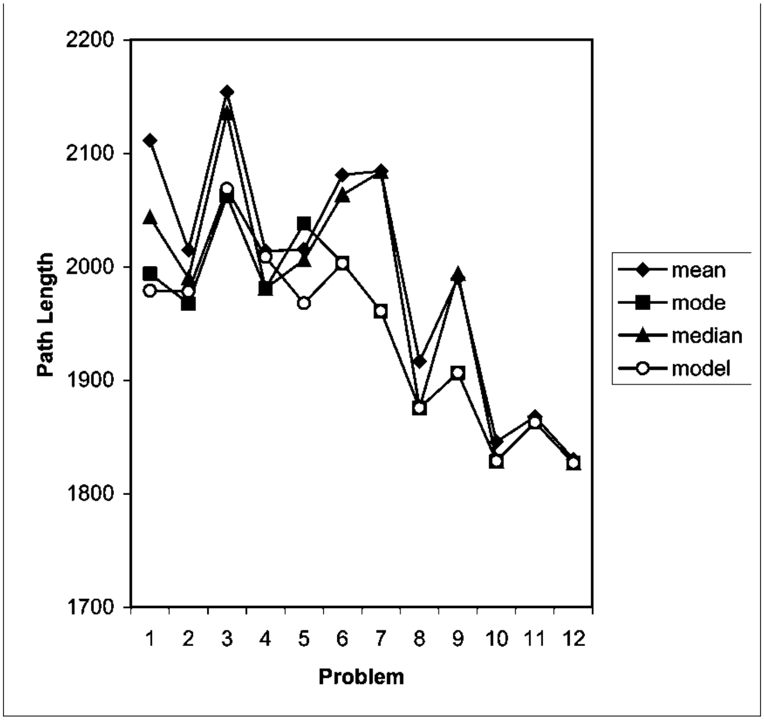


**Figure 4.** *Tour lengths generated by the convex-hull model compared with mean, median, and modal tour lengths for the 12 problems of Experiment 1.*

artefact of the particular sets of locations on the convex hull (particularly in the one interior point TSPs, the convex-hull locations formed a prominent quasi-rectangular shape). To further generalize the findings, Experiment 2 therefore used randomly generated point-sets. Second, it was felt desirable to demonstrate that the problem type effect would generalize to more complex problems: The TSPs used in Experiment 2 have approximately 43.5 billion more possible solutions than those used in Experiment 1.

## Method

### *Design*

Problem type was varied between participants as in Experiment 1 (open, closed). Starting-point was also varied between participants: Half started paths/tours at the left-most location (left) and half at the right-most location (right) of each problem. A total of 12 participants were assigned to each condition of the experiment. Each participant constructed paths/tours for the same set of six problems; the order of presentation was pseudorandomized.

### *Participants*

A total of 48 new participants, who had not taken part in Experiment 1 and were unaware of the aims and purposes of our research programme, volunteered for Experiment 2. All other participant details remain the same as those previously.

### *Materials*

TSPs for Experiment 2 all had nine interior points and six convex-hull points, and they were generated as follows. Sets of 15 locations were randomly sampled from within the rectangular coordinates $20 \leq x \leq 620$ and $10 \leq y \leq 460$ pixels, until a problem was generated with six points on the convex hull. This procedure was repeated until a total of six such problems had been obtained. For each of the six problems, coordinates were then plotted with standard software, as circular dots, diameter 3 mm, within a 150 × 110-mm grid. Paper booklets of problems in pseudorandom order, prefaced by standardized instructions, were prepared for each participant. An example problem is shown in Figure 2.

*Procedure*

The procedure was identical to that used in Experiment 1.

## Results and discussion

As in Experiment 1, the first analysis examined the frequency of optimal paths/tours. This was markedly lower than that in Experiment 1, at 14% compared with 39%, reflecting the increased complexity of the task. The resulting floor effect made it unsurprising that there were no significant differences between conditions, with 21 optimal solutions in the closed condition and 20 in the open.

We again conducted analyses using the more sensitive measure of path/tour length expressed as the standardized distance above optimal ($Z_{DAO}$), derived from participants' solutions as for Experiment 1.

Scores were submitted to a 2 × 2 analysis of variance, with problem type (open, closed) and starting-point (left, right) as the two between factors. There was no significant effect of starting-point, $F(1, 44) = 0.08$, or interaction effect, $F(1, 44) = 1.45$, $MSe = 0.01$. While performance was slightly superior in the closed condition, with an average $Z_{DAO}$ of 0.21 compared with 0.25 for the open condition, the result did not approach significance, $F(1, 44) = 1.64$, $p > .20$.

Experiment 1 indicated that while performance on closed tours tended to deteriorate as the number of interior points increased, performance on open paths did not. A trend of this kind suggests that as the number of interior points increases, performance on closed and open problems will coincide. It may be that the increase to nine interior points in the present experiment created exactly these conditions. Alternatively, the effect of the first experiment may have arisen from the artificial nature of the stimuli, disappearing when randomly generated stimuli were used. Experiment 3 further investigated these possibilities.

As in Experiment 1, we also examined the fit between our model and the observed data from participants. The results for closed tour construction were again consistent with the operation of a uniform heuristic, and, as before, the convex-hull model proved to be a promising candidate. Figure 5 compares the model's tour lengths across the six problems with the participants', and it indicates a good approximation.

*Individual differences*. Previously, we reported a failure to find evidence for individual differences on closed problems with 10 and 20 points (MacGregor & Ormerod, 1996). In contrast, Vickers et al. (2003a) reported substantial correlations in people's performance across both closed and open problems. They suggest that the apparent discrepancy between our results and theirs may be because the problems we used were both simpler than theirs and generated in a constrained rather than random way. The present results allow a closer examination of these possibilities, since the problems used in Experiment 2 were randomly generated. Also, they appear to have been of a comparable level of difficulty, in terms of performance, in that our open and closed versions had average PAO scores of 6.1% and 4.7%, respectively, compared with corresponding values of 7.1% and 4.8% for Vickers et al. (2003a).

In the present case, the rank-order correlation in PAO between each pair of problems was calculated separately for the participants from the closed and open conditions. For the closed condition the resulting 15 correlation coefficients ranged from $r = -.33$ to .66, with an average of .10. Of the

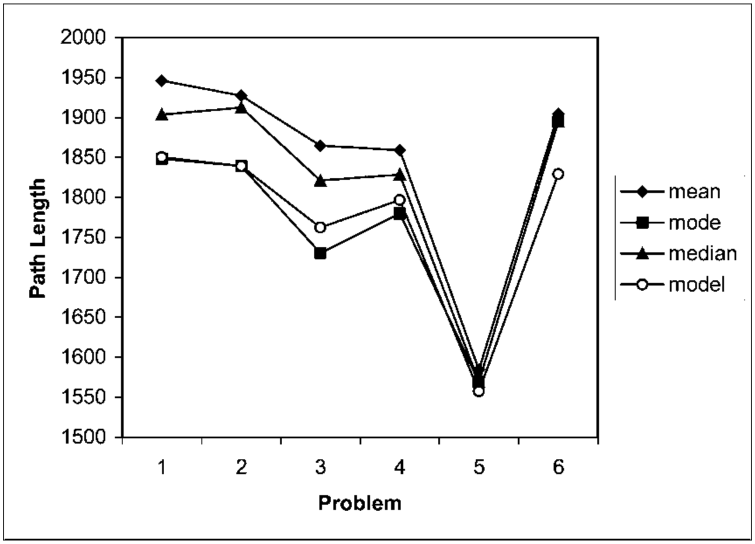


**Figure 5.** *Tour lengths generated by the convex-hull model compared with mean, median, and modal tour lengths for the 6 problems of Experiment 2.*

15, 2 were significantly greater than 0. For the open condition, the correlations ranged from −.29 to .31, with an average of −.01. None of the 15 coefficients was significantly different from 0. The results stand in sharp contrast with those of Vickers et al. (2003a), who report average correlations across problems ranging from .47 to .70. It seems likely to us that the discrepancy has arisen because of procedural differences. A number of features of the Vickers et al. (2003a) procedure may have allowed participants to improve performances across trials. Participants worked online, were permitted to erase and redraw arcs, were required to correct invalid paths/tours that had too few or too many links, and were given feedback after each trial, both by having the optimal path/tour displayed over their own solution and by having their score displayed as a percentage above optimal. In contrast, our participants completed the tasks without any form of corrective feedback. It seems possible, therefore, that the individual differences reported by Vickers et al. (2003a) reflect differences in learning or strategy and that such individual differences disappear when experimental conditions encourage relatively rapid and unsupported task performance.

## EXPERIMENT 3

The first experiment used highly constrained stimuli with relatively few interior points (one and six) and observed systematic differences in performance on closed and open instances. Experiment 2 used randomly generated stimuli with relatively many interior points (nine), and no systematic differences resulted. The present experiment used randomly generated stimuli, as in Experiment 2, with relatively few interior points (four), as in Experiment 1. Should the results replicate those of Experiment 2, it would suggest that the differences observed in Experiment 1 were due to the artificial nature of the stimuli. Should the results replicate those of Experiment 1, it would suggest that the failure to find an effect in Experiment 2 was due to the increased number of interior points.

### Method

#### *Design*

A single variable, problem type, was manipulated. It had two levels, open and closed, assigned between groups. The same eight problem exemplars were used in each of the open and closed conditions. The order of presentation of the eight problems was pseudorandomized for each participant.

*Participants*
A total of 40 undergraduate students in introductory psychology workshops at the University of Hawaii at Manoa volunteered to participate. No identifiers were collected, so the exact demographics of the participants are unknown. The use of human participants in Experiment 3 was approved by the Committee on Human Studies of the University of Hawaii.

*Materials*
Potential TSP instances were sampled by randomly selecting 15 points from rectangular coordinates extending from 20 to 620 pixels in the $x$ direction and from 10 to 460 pixels in the $y$ direction. This was repeated until eight 15-node instances were found that each had four interior points. Coordinates were then plotted with standard software, as circular dots, diameter 3 mm, within a 200 × 150-mm grid. Paper booklets of problems in pseudorandom order, prefaced by standardized instructions, were prepared for each participant.

*Procedure*
Participants were tested in groups of between 10 and 16. A general overview was provided verbally by the experimenter, and participants then read standardized instructions printed in their booklets. Participants were then given a maximum of 8 minutes to complete the booklet of eight problems.

## Results and discussion

In the case of 3 participants, the majority of their eight solutions could not be scored, either because they failed to include all 15 points or because the paths were ambiguous. These were excluded from further analysis, leaving 18 participants in the closed condition and 19 in the open.

Of the remaining 290 scoreable solutions, 59 were optimal, 43 (30%) in the closed condition, and 16 (11%) in the open. For analysis, the results were expressed in terms of the number of optimal tours per participant. For the closed condition, the mean number of optimal tours per participant was 2.39 (out of a possible 8), for the open condition, it was 0.84. The difference between conditions was significant, $t(35) = 4.72$, $p < .001$.

We repeated the analysis, using the standardized distance above optimal as the dependent variable. The effect of problem type was significant, $t(35) = 4.43$, $p < .001$, showing better performance on closed tours, with a mean of 0.14 standard units above optimal, compared with a mean of 0.36 for open paths.

For closed tours, the convex-hull model again provided a good fit to participant solutions (see Figure 6).

*Individual differences*. Rank-order correlations of performance were found for all possible pairs of the eight instances. For the closed condition, the resulting 28 correlation coefficients ranged from −.54 to .55, with a mean of .01. Four were significant beyond the .05 level, two positive and two negative. For the open condition, the coefficients ranged from −.38 to .37 with a mean of .05. None of the 28 differed significantly from 0. The results were similar to those found in Experiment 2, revealing no evidence of systematic individual differences in performance.

# GENERAL DISCUSSION

The first experiment reported in this paper demonstrates that human performance on TSPs is degraded in open versions, where the requirement to complete a circuit is removed. It further suggested that the degree to which tours are degraded in open versions might be dependent on the number of interior points in the problem. Taken together, the results of Experiments 2 and 3 replicate the findings of Experiment 1 very closely. The problems used in Experiment 2 had relatively many interior points, and the open/closed manipulation did not give rise to significant differences in performance, whereas there was a robustly significant improvement for closed over

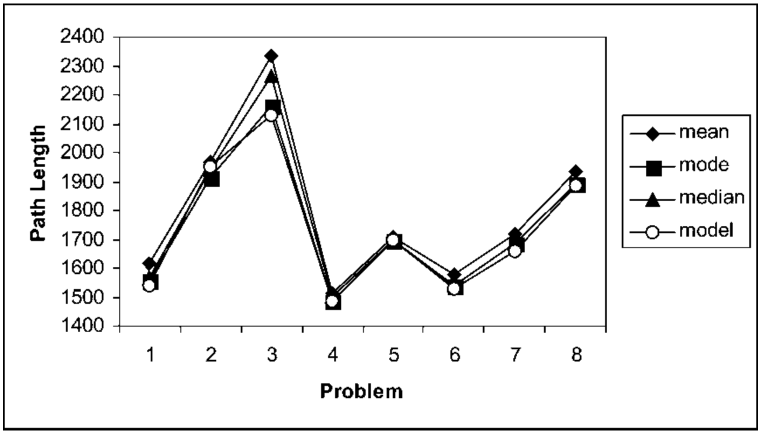


**Figure 6.** *Tour lengths generated by the convex-hull model compared with mean, median, and modal tour lengths for the 8 problems of Experiment 3.*

open problems in Experiment 3, in which the problems had relatively few interior points.

The results of all three experiments are consistent with our previous work, in which we demonstrate that the difficulty of closed problems is dependent on the number of interior points, rather than the total number of points in the problem (MacGregor et al., 2004). They further suggest that, by contrast, open problems may become slightly easier as the number of interior points increases. It is possible that human performance with open problems is itself guided by an heuristic, the nature of which is currently unknown, that becomes more reliable with increases in the number of interior points.

The experiments show, in addition, that the observed difference between closed and open versions cannot be explained by biases against the open versions that may have existed in previous research. Our explanation for the difference is that, in open versions, a convex-hull heuristic cannot be effectively used to guide path construction. The manipulation of heuristic use was achieved solely by changing instructions about where to start and finish, rather than by altering aspects of the TSP stimuli. The findings therefore challenge explanations based on local processes of point-to-point connection, since the experimental manipulation left the local point-to-point relationships unchanged. The findings converge with our previous work in which the ability of participants to use a convex-hull heuristic was disrupted using a different method. MacGregor et al. (2004, Exp. 3), laid out a 21-location TSP so that there was a sigmoid array of locations traversing the problem. This array of locations had the Gestalt property of good continuation. If a tour followed the sigmoid, it could not follow the convex-hull locations. MacGregor et al. (2004) then manipulated the perceptual salience of the convex-hull locations and found that the quality of tours improved with increasing convex-hull salience.

Why should a convex-hull heuristic appear to be so effective for human participants encountering TSPs? One of the visual system's most important natural tasks is to distinguish unitary objects from their backgrounds; vision scientists have argued that this requires an object boundary to be computed from a so-called raw primal sketch consisting of unconnected edges, terminations, and line segments (Marr, 1983). The visual system appears to be adapted at a physiological level to derive curvilinear contours from the synchronized output of individual, spatially localized orientation selective units (Field, Hayes, & Hess, 1993). Moreover, psychophysical detection performance is higher for closed rather than open contours, when these contours are defined by individual, separated units (Pettet, McKee, & Grzywacz, 1998): The visual system seems naturally to prefer closure. Neuro-computational models exist that satisfactorily account for such psychophysical phenomena (Yen, 1999).

Two intriguing possibilities arise, therefore: First, the visual system may implement potentially useful algorithms that have not yet been explored by computer scientists and others attempting to solve TSPs. There is some evidence that TSP tour construction is indeed guided by low-level heuristics. Findlay (2003) has developed an eye movement task in which a stimulus array consisting of discrete locations is searched. At each location, the presence or absence of a particular letter must be noted. At the finishing location, the participant must verify whether the number appearing there is the number of target letters seen. Under these circumstances, preliminary data indicate that the eye movement system programmes a close-to-optimal tour of the locations. It is therefore plausible to argue that the convex-hull heuristic operates at a low level in the visual system and that its output is available to several other neurocognitive systems. It should be noted, however, that the work of Graham et al. (2000) suggests that other heuristics, notably hierarchical clustering, are implementable by the visual system and may provide a reasonable fit to human performance on TSPs. As yet, there is no direct experimental evidence for or against such alternative heuristics.

Second, it seems likely that there will be other combinatorial optimization problems that humans perform well on, by virtue of input provided to higher level cognitive functions by their routine perceptual processes. This second suggestion has already received some support from the finding that humans are good at finding near-optimal solutions to the Euclidian p-median problem (Brusco, 2001). In this problem, a subset S of p locations from V, a set of n locations in the plane, is selected. Each of the remaining n-p locations must be joined to the nearest location in S. The goal is to select S such that the sum of the linear distances between selected and unselected locations is minimized. The task is analogous, say, to having purchased 50 vacant lots ($n = 50$) with a view to establishing a supermarket chain. If distribution depots must be built on 4 lots ($p = 4$), and stores on the remaining 46 lots ($n - p = 46$), it will be desirable to site the distribution depots such that the sum of the distances between each store and its assigned distribution depot is minimized. Brusco (2001) reports that human participants are capable of providing solutions to p-median problems within 2% of optimal and further comments that participants may "make a global assessment of the figure and begin with the visual partition of the [locations] into p groups".

Human performance with the simple context-free TSPs used in these experiments seems to be determined by a convex heuristic that fits neatly within the fast-and-frugal approach described by Gigerenzer and Goldstein (1996). Although consistent with this kind of heuristic, we do not suggest that a convex-hull strategy is of broad generality: For example, it is unlikely that such a strategy would yield solutions of high quality to TSPs that contained additional constraints that might be found in realistic navigational optimization problems (e.g., different node priorities, restrictions on order of node visitation). In these cases, we suspect, humans are likely to switch to problem-solving heuristics other than the convex-hull heuristic. Nonetheless, we suggest that optimization problems are likely to become fertile ground for cognitive science in the coming years. The problems themselves are of intrinsic interest, as it is often not computationally feasible to compute the gold-standard solution. Despite the complexity of the problems, human heuristics often give rise to close-to-optimal solutions, as we and others have demonstrated. The physiological implementation of heuristics such as the convex-hull heuristic seems—to us—likely to occur at a low level in the visual system. It may well be that some of the routine tasks of visual processing can be recast as issues of combinatorial optimization. If so, the low-level neurocomputational processes that have evolved to cope with the requirements of vision may provide the underpinnings for fast and frugal human heuristics that cope admirably well with the challenges posed by combinatorial optimization problems.

Original manuscript received 16 June 2004
Accepted revision received 18 January 2005
PrEview proof published online 10 August 2005

## REFERENCES

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Applegate, D., Bixby, R., Chvátal, V., & Cook, W. (2001). *Milestones in the solution of TSP instances.* Retrieved June 4, 2003, from Princeton University, Mathematics Department Web site: http://www.math.princeton.edu/tsp/milestone.html

Boyd, S. C., Pulleyblank, W. R., & Cornuéjols, G. (1987). TRAVEL: An interactive travelling salesman problem package for the IBM personal computer. *Operations Research Letters, 6*, 141–143.

Braine, M. D., & O'Brien, D. P. (1991). A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review, 98*, 182–203.

Brusco, M. (2001). *Human performance on the Euclidean p-median problem.* Unpublished manuscript.

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field". *Vision Research, 33*, 173–193.

Findlay, J. M. (2003). Eye scanning and visual search. In J. M. Henderson & F. Ferreira (Eds.), *The interface of language, vision and action: Eye movements and the visual world.* New York: Psychology Press.

Gigerenzer, G. (2001). The adaptive toolbox. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox.* Cambridge, MA: The MIT Press.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103*, 650–669.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

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 travelling salesman problem: A hierarchical model. *Memory & Cognition, 28*, 1191–1204.

Krolak, P., Felts, W., & Marble, G. (1971). A man–machine approach toward solving the travelling salesman problem. *Communications of the ACM, 14*, 327–334.

MacGregor, J. N., Chronicle, E. P., & Ormerod, T. C. (2004). Convex hull or crossing avoidance? Solution heuristics in the travelling salesperson problem. *Memory and Cognition, 32*, 260–270.

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

MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (1999). Spatial and contextual factors in human performance on the travelling salesman problem. *Perception, 28*, 1417–1427.

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

Marr, D. (1983). *Vision: A computational investigation into the human representation and processing of visual information.* San Francisco, CA: Freeman.

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

Oaksford, M., & Chater, N. (1994). A rational explanation of the selection task. *Psychological Review, 103*, 381–391.

Ormerod, T. C., & Chronicle, E. P. (1999). Global perceptual processing in problem-solving: The case of the travelling salesperson. *Perception and Psychophysics, 61*, 1227–1238.

Pettet, M. W., McKee, S. P., & Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research, 38*, 865–879.

Quintas, L. V., & Supnick, F. (1965). On some properties of shortest Hamiltonian circuits. *American Mathematical Monthly, 72*, 977–980.

Rips, L. J. (1994). *The psychology of proof: Deductive reasoning in human thinking.* Cambridge, MA: MIT Press.

Sangalli, A. (1992, December 12). Why sales reps pose a hard problem. *New Scientist*, 24–28.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*, 207–232.

van Rooij, I., Stege, U., & Schactman, A. (2003). Convex hull and tour crossings in the Euclidean travelling salesperson problem: Implications for human performance studies. *Memory and Cognition, 31*, 215–220.

Vickers, D., Bovet, P., Lee, M. D., & Hughes, P. (2003a). The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean travelling salesperson problems. *Perception, 32*, 871–886.

Vickers, D., Butavicius, M., Lee, M., & Medvedev, A. (2001). Human performance on visually presented travelling salesman problems. *Psychological Research/Psychologische Forschung, 65*, 34–45.

Vickers, D., Lee, M. D., Dry, M., & Hughes, P. (2003b). The roles of the convex hull and number of intersections in performance on visually presented travelling salesperson problems. *Memory & Cognition, 31*, 1094–1104.

Yen, S.-C. (1999). Cortical computation of perceptual salience. (Doctoral dissertation, University of Pennsylvania, 1999). *Dissertation Abstracts International, 59 (7-B)*, 3278.