

Human Performance on the Traveling Salesman and Related Problems: A Review

James N. MacGregor¹ and Yun Chu²

Invited

Abstract:

The article provides a review of recent research on human performance on the traveling salesman problem (TSP) and related combinatorial optimization problems. We discuss what combinatorial optimization problems are, why they are important, and why they may be of interest to cognitive scientists. We next describe the main characteristics of human performance on the TSP and related problems, and discuss the main theoretical explanations that have been offered. We then review some related developments in animal studies, spatial cognition, and neuropsychology. The article closes with a brief look at possible future directions.

Keywords:

traveling salesman problem, human performance, cognition

¹University of Victoria; ²University of Hawaii Manoa

1. What Are Combinatorial Optimization Problems?

Combinatorial optimization problems involve the search for an optimal solution among a discrete collection of possible solutions. Typically, the number of possible solutions grows exponentially with the number of elements in the problem, so “scanning all objects one by one and selecting the best one is not an option” (Schrijver, 2003, p. 1). This challenge posed by combinatorial optimization problems is vividly illustrated by a historical example. In 1962, Procter & Gamble ran a contest in which 1 prize of \$10,000 and 54 prizes of \$1,000 were awarded for finding the shortest round-trip route through 33 U.S. cities, starting and finishing in Chicago (Applegate et al., 2006). The total prize had the equivalent purchasing power of approximately half-a-million dollars today. The apparent simplicity of the challenge stands in sharp contrast with the relative magnitude of the prize.

Although simple to state and understand, the problem posed by Procter & Gamble is an instance of a well-known combinatorial optimization problem, the TSP (for an extensive review, see Lawler et al., 1985). The number of possible round-trip routes for TSPs is equal to $(n - 1)!/2$, where n is the number of cities. For the Procter & Gamble problem, this number calculates to about 1.32×10^{35} , or 132 billion trillion trillion possible routes. If one entered the competition, and programmed a computer to find every single route and select the optimal (shortest), it would take about 417 billion trillion years (assuming the computer found 10,000 routes per second). By then, of course, the value of the first prize would have diminished considerably.

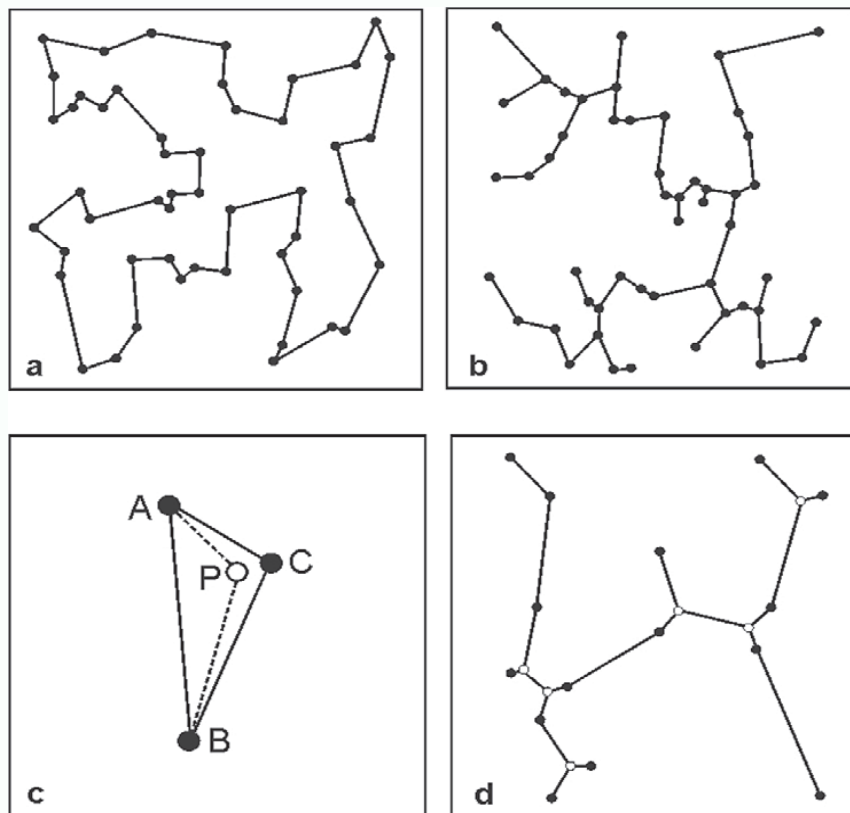
The TSP is ubiquitous in science and industry, being fundamental to such applied problems as

- routing of school buses to pick up children: time and money is saved if the bus travels the minimum distance (Schrijver, 2003)
- home delivery of meals to elderly: meals remain hot and money is saved if the shortest route is found (Schrijver, 2003)
- scheduling of machines to drill holes in a circuit board: in manufacturing processes, millions of drilling operations may be necessary. If the distance traveled by the robot drilling arm is minimized, production is made more efficient (Sangalli, 1992)
- analysis of the structure of crystals by X-ray diffraction: the necessary repositionings of the diffractometer are analogous to cities in a TSP. Minimizing the total repositioning movement makes analysis most efficient (Bland & Shallcross, 1987)
- storage and picking of stock in warehousing: how to select an optimal storage strategy so that subsequent retrieval of stock items is done with maximum efficiency (Dallari, Marchet, & Ruggeri, 2000)

- genome sequencing: software for TSPs can overcome difficulties in combining partial maps of genomes, and do so better than software specialized for genome research (Agarwala et al., 2000).

Very many other combinatorial optimization problems exist, and some are illustrated in Figure 1 which, in addition to an example TSP, presents instances of a minimum spanning tree problem (MSTP) and generalized Steiner tree problem (GSTP). The latter problems are discussed at a later point. While all combinatorial optimization problems have in common a large number of possible solutions, which makes it impractical to search exhaustively for the optimal one, a subset of problems is known to be especially challenging. These problems are termed NP-hard (where NP stands for nondeterministic polynomial). We review their importance next.

Figure 1. (a) TSP instance with optimal solution, (b) MSTP instance with optimal solution, (c) the Fermat point (P) for a simple Steiner tree problem (shown as open circle), (d) GSTP instance with the interpolated open points (open circles) and the optimal solution. (From Burns, Lee, & Vickers, 2006).



2. Why Are NP-hard Problems Important, and Why Should Human Performance Be of Interest?

Many combinatorial optimization problems, as illustrated above, are resistant to solution by brute-force computation beyond a certain n , precisely because of the exponential growth of the number of possible solutions. Instead, it is desirable to find algorithms that are able to discover the solution in an acceptable time frame. Some combinatorial optimization problems can be proven to be solved by such algorithms in polynomial time (e.g., as a function of the cube, or fifth power, etc. of the number of elements in the problem). These problems are thus known as polynomial-time solvable (P). Other combinatorial optimization problems are classified as NP-hard. One feature of NP-hard problems is that no polynomial-time algorithm currently exists for any of these problems, and it is widely believed that none exist. The TSP is one such NP-hard problem.

Two major challenges for mathematics have stemmed from the division between P and NP-hard problems. First, it is known that if a guaranteed polynomial-time algorithm were to be found for a single NP-hard problem, then all NP-hard problems would be solvable in polynomial time (in effect, $P = NP$ and so the class of NP-hard problems would cease to exist). This first challenge still stands. Second, finding algorithms that give very-close-to-optimal (but not necessarily optimal) solutions to NP-hard problems is of major importance, given that the optimal solution cannot—at present—be guaranteed. Major progress has been made in this area over the past several decades. One spectacular recent success is the solution of a TSP comprised of the 24,978 cities of Sweden (www.tsp.gatech.edu/sweden/index.html). Despite the sophistication of the methods used in this solution, the total CPU time required was in excess of 90 years.

It is evident that efficient algorithms for NP-hard problems in applied settings are desirable. One question that has only recently attracted attention is whether untrained humans can provide solutions to NP-hard problems, and if so, whether the solutions are close to optimal. This question has been studied most closely in the context of the two-dimensional Euclidean version of the TSP.¹ MacGregor and Ormerod (1996) demonstrated that in the case of 10- and 20-city problems, human subjects outperformed simple construction algorithms by an order of magnitude. Untrained human subjects, when instructed simply to draw the best route “by eye,” without any real-world constraints being imposed, typically provided solutions within 1% of optimal. Since these earlier studies, additional work has provided some evidence to show that human performance remains remarkably good even with considerably larger problems. Dry et al. (2006) used randomly generated TSPs of 10, 20, . . . , 120 points and demonstrated a linear relationship between number of cities and solution time. The quality of performance showed a similar linear relationship: even for TSPs of 120 cities, human performance was only approximately 11% above

¹ Unless otherwise stated, this is the version referred to throughout the article.

optimal. The limits to human performance, in terms of the number of cities in a problem, remain to be tested empirically: it is possible, for example, that larger problems might cause catastrophic worsening of performance.

The quality of human performance with TSPs has invited several questions that are likely to be of importance for three distinct fields of inquiry. First, of interest to computer science, is it possible that humans implement algorithms that are computationally less complex than standard mathematical algorithms? Second, of interest to operations research, is it possible that human performance may be capitalized upon in applied settings? Third, of interest to psychologists and neuroscientists, how do the cognitive systems of the brain deal with TSPs, and what underpins the high level of performance observed?

3. The Characteristics of Human Performance with NP-hard Problems

3.1. The Standard TSP

To our knowledge, the first systematic investigation of human performance with TSPs was published in Russian (Polivanova, 1974).² This paper contrasted different methods of problem presentation: tables of intercity distances versus a two-dimensional plot with the cities represented as dots. The number of cities was varied from 4 to 10. On a number of measures (number of solution steps, time taken, and number of optimal solutions), participants performed better when the cities were visually represented. When the number of cities was small (four and five), the advantage of the visual presentation was reduced. This is still one of the only papers in the literature explicitly contrasting differences between methods of problem presentation, and it makes the compelling point that visual processes appear to be important for humans in solving TSPs. The only other study of which we are aware used a variant of the TSP to compare several forms of problem representation, including two-dimensional plots versus tables of pair-wise distances (Gärling, 1989). The results were consistent with Polivanova's.

Recent psychological research was catalyzed by MacGregor and Ormerod (1996), who, as mentioned above, provided evidence that for 10- and 20-city TSPs, humans were able to draw shorter tours than simple construction algorithms, and often found solutions very close to optimal. Participants in their experiments had no particular knowledge of the TSP, and were simply asked to draw a freehand tour based on their best judgment.³ MacGregor and Ormerod (1996) further demonstrated that (i) there was little evidence

² An earlier paper, by Michie, Fleming, and Oldfield (1968) reported somewhat anecdotal data from five human participants attempting TSPs in various conditions either aided or unaided by a computer interface. It is not possible to compare the performance of their subjects to optima.

³ We recently became aware of an earlier article (Hill, 1982) that compared human performance with that of several heuristic procedures. The results were highly consistent with those cited in the present review.

of systematic individual differences in solving either 10- or 20-city problems, and (ii) the number of cities interior to the convex hull of the problem was highly correlated with response uncertainty, a measure of problem complexity. Participants also typically connected boundary points in order and avoided crossed arcs—both features are necessary for an optimal solution to TSP. (The convex hull can be visualized as an elastic band stretched around the nodes representing cities so that each node is enclosed. Nodes in contact with the elastic band lie on the boundary.)

Other studies have since replicated the basic finding that unaided human performance on visually presented problems is of high quality. MacGregor, Ormerod, and Chronicle (1999) showed that participants generated tours to a 48-city problem that were, on average, 10% longer than the optimal tour. Graham, Joshi, and Pizlo (2000) examined problems of sizes 6, 8, 10, 20, and 30, finding (i) good performance, with 50 to 90% of trials finding the optimal solution for 6-, 8-, and 10-city problems and tour length varying from approximately 1 to 3% above optimal across all problems, and (ii) a linear relationship between number of cities and performance. As with MacGregor and Ormerod (1996), human performance was frequently superior to some simple construction algorithms. Vickers et al. (2001, Experiment 1) used two groups of participants, one of which was given instructions to optimize tour length and the other of which was given instructions to draw the most aesthetically pleasing route through the cities. The optimization group produced significantly shorter tours on average than the aesthetic group. Furthermore, the best human solutions were optimal for 10- and 25-city problems, unlike a comparator elastic net algorithm, and for 40-city problems, human solutions were only approximately 2.5% above optimal, compared to 8.5% for the elastic net. An additional manipulation showed faster solution speeds for problems with more cities on the convex hull. Unlike the work of MacGregor and Ormerod (1996), Vickers et al. (2001) found evidence for individual differences in the form of positive correlations between performances on different problems. Furthermore, in Experiment 2, they demonstrated significant correlations between various measures of solution quality and scores on Raven's Advanced Progressive Matrices. van Rooij et al. (2006) compared the performance of children aged approximately 7 and 12 with that of adults, using problems of 5, 10, and 15 nodes. Even the youngest group of children found solutions within 7% of optimal, on average. Nevertheless, average performance improved with age: for the adult group, 42% of solutions were optimal, compared with 27% for the 12-year-old group and 20% for the 7-year-olds. As noted above, Dry et al. (2006) used a substantial set of problems, of 10, 20, ..., 120 cities, and discovered a linear increase in solution time with number of cities. Performance in terms of percentage above optimal tour length was also excellent, with tours being only about 11% above optimal even for the 120-city problems. Finally, Burns, Lee, and Vickers (2006) demonstrated significant relationships between general intelligence measures and performance on TSPs.

To summarize the available findings on human performance with TSPs, therefore, it

seems that there is considerable agreement that humans (i) require a visual representation of the problem; (ii) are good at the task, frequently outperforming simple computational algorithms in terms of finding tours close to optimal; and (iii) demonstrate an approximately linear decrement in performance (either tour distance or speed of completion) with increasing number of cities. There are equivocal findings about whether individual differences in performance exist, and also about the relative influence on performance of cities on the convex hull and interior to the convex hull.

A study by Chronicle et al. (2008) addressed the issue of individual differences by comparing the different approaches of two research groups. One group (MacGregor, Ormerod, Chronicle), typically using relatively small problems in a paper-and-pencil format, consistently reported negligible individual differences in performance. A second group (Lee, Vickers et al.), using a wider range of problem sizes combined with computer presentation, consistently reported larger, reliable, individual differences. Chronicle et al. combined the moderately large problems generated and tested by the latter group with the paper-and-pencil data collection procedures of the former. The results indicated the presence of individual differences at a level between those previously reported by the two groups. In addition, the pattern of results indicated that individual differences become more prevalent as problem size increases.

3.2. Related Problems

While most studies of human performance with optimization problems have been conducted using standard TSP-type problems, a number have introduced variations. In some cases, the variations have involved relatively minor changes to the instructions or task requirements. In other cases, quite different optimization tasks have been involved.

3.2.1. Instructional Variants of TSPs

Typically with TSP problems, participants are instructed that the task is to find the shortest route through the set of points, returning to the starting point, so that valid solutions form a closed tour. A variation on this approach is to use an open TSP instruction, which omits the latter requirement, so the task becomes that of finding the shortest path through a set of points, without returning to the start (Hirtle & Gärling, 1992; MacGregor & Ormerod, 1996). Although this may appear to be a minor variation, it has theoretical implications that we discuss at a later point.

Gärling (1989) conducted an experiment using several different five-node problems. A major manipulation compared performance on variations of tasks where a nearest neighbor strategy would and would not lead to the optimal solution. When the tasks were presented spatially, the large majority of participants in both conditions drew the optimal path. In fact, more did so when the optimal did not conform to the nearest neighbor strategy (94%) than when it did (71%). However, when the problem was presented as a

table of intercity distances rather than as a spatial plot, a majority of participants in both conditions selected the path consistent with the nearest neighbor strategy, whether or not it was optimal. The results are consistent with those of Polivanova's (1974) concerning the effects of problem representation on solution quality. They also support the growing body of evidence that, when the problem is presented spatially, human participants employ a strategy beyond a simple local heuristic of connecting to the next nearest point.

Hirtle and Gärling (1992) used open problems in a 2 x 3 within-subjects design with two levels of node distribution (clustered versus isotropic) and three levels of number of nodes (6, 10, and 18). In each problem, start and end points were given. Participants were presented with four versions of each problem, the original plus three others formed by rotation and translation. With 24 participants, this provided 96 solutions for each of the 6 problems. The authors were primarily interested in the heuristics employed by participants, and did not report the factorial analysis of variance implied by the experimental design. However, the results were reported in sufficient detail to allow reproduction of the cell means, which are shown in Table 1. In each cell, the upper number reports the mean percentage above the optimal (or best known) solution of human solution and the lower number shows the percentage of human solutions that were optimal. At a descriptive level, performance generally deteriorated with increasing number of nodes (with the exception of the 10-node distributed instance) and was somewhat better for clustered than for distributed instances (again with the same exception). The latter result is consistent with the proposition that people attend to clusters of dots when forming a solution.

Table 1. Percent above optimal (upper rows) and percent of optimal solutions (lower rows) for open problems with different numbers of nodes and different degrees of clustering (from data reported in Hirtle & Gärling, 1992)

	Number of Nodes		
	6	10	18
Problem Type			
Clustered	1.4% 90%	2.8% 50%	5.9% 19%
Distributed	3.3% 32%	2.7% 73%	9.6% 8%


Basso, Bisiacchi, and colleagues have used a variation of the standard TSP referred to as the Maps test. The main purpose of the research has been to identify how heuristics are used by human solvers and how this may be affected by brain damage or other forms of functional impairment (Basso et al., 2001; Basso et al., 2006). These latter findings will be described later, while in the present section we will describe the task itself and some

of the more general findings. Problems are presented on a computer screen displaying a schematic map of city blocks, typically with seven vertical and five horizontal roads. Instances are small, usually varying from four to nine nodes plus indicated start and end points displayed at intersections. Participants use keyboard arrow keys to move a cursor along roads to indicate a path between points. Using the Maps task, Basso and colleagues have reported results indicating that people may switch among several heuristics during the process of a single solution, and that inhibiting switching results in performance decrements (Basso et al., 2006; Cutini et al., 2008). Recent research from this group has also reported gender differences (Cazzato et al., 2010). Males and females completed 30 problems and performances were compared on a number of measures, including path lengths, response times, and heuristics employed (the latter were inferred by matching participant path segments to segments produced by four heuristics). The results indicated that males produced significantly shorter paths than females and were more likely to switch heuristics during problem solving.

The Maps task differs from the standard TSP not only in using an open path, but in using city-block instead of Euclidean space. (In the city-block norm—also known as the “Manhattan” and “L₁” norm—the distance between two points is measured along axes at 90 degrees.) This latter difference is potentially important, since a city-block metric may affect perceptual characteristics of optimal solutions. To illustrate, consider some of the optimal city-block solutions to the closed TSP shown in Figure 2. Assuming that each side of the square shown in the solution on the left is two units long, then the optimal city-block solution is eight units. This is the length of the other two solutions shown. However, both of them violate principles of optimal solutions in Euclidean space. The center solution has a crossing, while the solution to the right involves retracing of line segments. Any Euclidean solution with either of these characteristics must be non-optimal (van Rooij, Stege, & Schactman, 2003).

While the studies described so far examined the open task in its own right, others have compared performance on open versus closed versions. In one case, Vickers et al. (2003) report an experiment comparing performance on 10 randomly generated 30-node

Figure 2. Three optimal solutions to a closed TSP in city-block metric.




stimulus arrays under standard (closed tour) instructions with performance on the same 10 arrays under open instructions. The results indicated that performance on closed versions, at 4.8% above optimal, was consistently better than on open versions, at 7.1%. Similar results were reported by Chronicle, MacGregor, and Ormerod (2006), with performance on closed versions being better than on open versions in three experiments, although in one case, the difference did not reach significance. The implications of these differences are discussed below, when we describe the theoretical approaches to human performance.

As mentioned previously, another variation on the standard instruction to find the shortest tour was introduced by Vickers et al. (2001), in a condition where participants were instructed to draw a path that “looked most natural, attractive, or aesthetically pleasing” (p. 36). When compared with a second group, operating under the standard instruction to minimize tour length, the aesthetic group produced tours that were slightly longer. Nevertheless, the authors concluded that the results under aesthetic instructions corresponded to minimal solutions to a remarkable extent. The findings recall those of Pomerantz (1981), who reported that when subjects were asked to connect dot arrays to illustrate any patterns they perceived, they frequently connected the dots using shortest possible paths.

A further variation on standard instructions was reported by Chronicle, MacGregor, and Ormerod (2006), who investigated whether participants are as adept at finding *long* tours as they are at finding short tours (see Figure 3). The results using 10-point arrays indicated that they are not, by a substantial margin. A long tour took more time to produce and was significantly farther from the maximum path length than a short tour was from the minimum. Specifically, long tours required an average of 13 seconds and were almost 2 standard deviations away from the longest path, compared with short tours, produced in 6 seconds and less than half a standard deviation from the shortest. Human performance in producing long tours was also significantly worse than a simple farthest neighbor construction heuristic, to connect the next line segment to the most distant unconnected node. A second experiment using 15-node stimuli replicated the main findings, while controlling for the effects of increased visual clutter due to multiple line crossings under the longest path instructions.

To summarize the effects of instructional variations, the results indicate that participants continue to be effective in finding near-optimal tours, even when they are instructed to find “attractive, aesthetically pleasing” tours rather than shortest ones. On the other hand, when the requirement to return to the starting point is lifted, as in the “open path” variation, there is a reliable degradation in performance. An even more striking degradation occurs when the task requirement is changed from finding the shortest tour to finding the longest tour, where human performance appears to be significantly poorer even than the simple farthest neighbor construction heuristic. The theoretical implications of these differences are discussed in Section 4.

Figure 3. Longest and shortest tours for a TSP with $n = 10$ (from Chronicle et al., 2006).



3.2.2. Other Optimization Tasks

While the bulk of the work examining human performance with optimization problems has been conducted using TSPs, with and without instructional variations, a few studies have considered other types of optimization problem. Vickers et al. (2004) examined individual differences in performance on the TSP and two additional optimization problems, the MSTP and the GSTP, illustrated in Figure 1. The MSTP requires the construction of the shortest network that directly links all nodes, but does not require a continuous closed tour (branching is acceptable). The GSTP similarly requires the shortest path to be constructed, but allows the insertion of additional nodes (Fermat points). Edges incident to Fermat points must form three 120 degree angles. Vickers et al. (2004) used a 50-node TSP, 50-node MSTP, and 15-node GSTP, and showed that human performance was 9%, 6%, and 12% above benchmark respectively, in a sample of 48 students. Furthermore, they demonstrated significant correlations between performance on all three problems and Raven's Advanced Progressive Matrices. They concluded that there is evidence for individual differences in human performance with these optimization problems.

A study by Burns, Lee, and Vickers (2006) compared performance on three instances each of TSP, MSTP, and GSTP and on standard psychometric tests designed to measure various cognitive abilities. The results indicated significant correlations among TSP, MSTP, and GSTP performance and with measures of both low-level perceptual abilities, particularly involving spatial relations, and more high-level cognitive abilities, particularly involving verbal analytic reasoning. Burns, Lee, and Vickers (2006) proposed a structural equation model to account for the intercorrelations they observed between tests. This

structural equation model explained individual variation in combined TSP/MSTP/GSTP performance both in terms of a fluid intelligence factor, reflecting abstract reasoning ability, and a visuo-spatial intelligence factor, reflecting the ability to manipulate mental representations of objects.

Pizlo and Li (2005) examined human performance on the 15-puzzle (and variants), in which 15 sliding tiles are arranged on a 4 x 4 grid. The task is to rearrange the tiles from a scrambled initial state into an ordered goal state, but tiles may only be slid into the empty space adjacent to them. The 15-puzzle and its variants are known to be NP-hard problems. Although it was not possible for the authors to determine the optimal number of moves to solve the 15-puzzle, subjects typically reached solutions by visiting 200 or fewer states of the problem. This compares very favorably with the performance of algorithms designed to find the optimal solution, which visit between 540,000 and 6 billion states (Pizlo & Li, 2005, p. 1071). In a second experiment, Pizlo and Li examined performance on 5-, 8-, 15-, and 35-puzzles, with the aim of determining the complexity of the mental mechanisms involved in solving the puzzle. They showed a direct proportionality between number of moves and time to solution, suggesting that the rate at which participants worked was self-paced, rather than being dependent on the size of the problem space. They also demonstrated a low-complexity relationship between problem size and time to solution, suggesting that search of the factorially increasing problem space could not be occurring.

To the extent that problem solving involves both perceptual and cognitive processes, the 15-puzzle appears to lie farther to the cognitive end of the spectrum than the TSP, MSTP, and GSTP. Similarly, there are many other optimization problems with large state spaces (though not necessarily NP-hard) that have attracted the attention of psychologists and that have a strong cognitive component. These include optimal stopping problems, such as the Secretary problem (Burns, Lee, & Vickers, 2006; Ferguson, 1989), and sequential decision-making tasks, like bandit problems (Lee et al., 2010). Recently, an elaboration of the TSP—*TRAVELPLAN*—has been suggested as an alternative to platforms such as *Tailor-shop* for the study of complex problem solving from the European perspective (Ragni & Löffler, 2010). While based on the TSP, *TRAVELPLAN* includes additional constraints such as those imposed by schedules and preferences. While we recognize the wide range of optimization tasks and the many developments that are taking place in related fields, the present review focuses on those optimization tasks that contain a strong visuo-perceptual component.

While the study of optimization tasks other than the TSP has been fairly limited, the results to date suggest that human performance with these problems is also impressive. In addition, the results involving TSP, MSTP, and GSTP indicate the presence of systematic individual differences, which may involve both lower-level perceptual processes and higher-level analytical processes. This conclusion appears to be compatible with the

body of findings on TSP performance more generally. With relatively small problems, of 10 and 20 nodes, there is little evidence of individual differences in adult performance (MacGregor & Ormerod, 1996), suggesting that the role of lower-level processes may predominate. However, as problem size and complexity increase, so does the evidence of individual differences, suggesting a greater role of analytical processes (Chronicle et al., 2008). Similarly, comparisons of the performance of children with adults has led to the conclusion that very good TSP performance is attainable by perceptual processing alone, but that cognitive-analytical skills may still supply added value to what is provided by the perceptual system (van Rooij et al., 2006). Larger TSP instances, such as those of 50 nodes used by Vickers et al. (2004), may sufficiently draw upon analytical processes to allow for the ready detection of reliable individual differences. We next turn to the theoretical accounts that have been proposed to account for this pattern of findings.

4. Theoretical Accounts of Human Performance

Polivanova (1974) presciently suggested that TSP problem solving “is firstly based on the general appearance of the itinerary (considering such features as convexity, smoothness, the presence of obtuse angles, the absence of crossing lines, simplicity, aesthetic appearance) and secondly on complexes of points related by their degree of proximity to each other” (translated from p. 45 of the original). A number of theoretical positions have emerged that relate to one or other of her suggestions.

4.1. The Sequential Convex-hull Model

MacGregor and colleagues (MacGregor & Ormerod, 1996; MacGregor, Ormerod, & Chronicle, 1999; MacGregor, Ormerod, & Chronicle, 2000) have developed a model of human performance based on perception of the convex hull of the problem. MacGregor and colleagues argue that the convex hull is readily and rapidly available via low-level global perceptual processes. Several models of object perception depend on global precedence: the overall shape of a figure is perceived prior to the infill of details (Navon, 1977; Sanocki, 1993). In MacGregor et al.’s model, therefore, first, early perceptual processes allow a sketch of the convex hull to be generated, then interior points are connected sequentially according to an insertion criterion, in a clockwise or counterclockwise direction around the convex hull. This sequential convex-hull model provides good fits to human tour length data for problems with 10 to 100 cities (MacGregor, Ormerod, & Chronicle, 2000). Evidence for the importance of the convex hull for human performance comes from three sources. First, Ormerod and Chronicle (1999) showed that solution recognition is rapid, and is primed by short (but not long) presentations of the problem array. Both findings support the notion that global perceptual processing of shape underpins human performance. Second, increasing the number of cities interior to the convex hull increases human tour

length relative to optimality (i.e., reduces solution quality), but increasing the number of convex hull points does not affect human tour length (MacGregor, Chronicle, & Ormerod, 2004). Third, introducing manipulations that prevent the convex hull of the problem being used reduce solution quality substantially. Chronicle et al. (2006) showed that human performance on open TSPs, where the requirement was to start at the leftmost city and finish at the rightmost city of the array, was significantly impaired in comparison with closed TSPs. In the open version, although the convex hull is still perceptually available, it cannot be used to guide tour construction. On the other hand, the empirical support for the position has been questioned (Lee & Vickers, 2000; van Rooij, Stege, & Schactman, 2003), and qualitative discrepancies have been reported by Tak, Plisier, and van Rooij (2008) between their implementation of the model and human performance on a number of problems specially constructed to provide strong tests of the model.

4.2. Pyramid Models

Pizlo and colleagues (Graham, Hoshi, & Pizlo, 2000; Pizlo et al., 2006) have proposed a hierarchical clustering approach, implemented through a pyramid model consisting of layers in which the stimulus array is represented at a successively coarser level of resolution. In the original model (Graham, Joshi, & Pizlo, 2000), a TSP instance is represented as a high-resolution image at the bottom layer. At the next level, the resolution is reduced through Gaussian blurring, and this bottom-up process is repeated through each level of the hierarchy. At each higher level, the image is represented by intensity distributions having successively fewer modes. Following the representational phase, the solution phase begins at the lowest level in the hierarchy, at which there are no more than three modes, and an initial tour is produced connecting the modes. Then, at the next lower level, the modes of that level are inserted into the initial tour. This process is repeated in a top-down fashion resulting, at the lowest level, in a tour of the original TSP instance. The authors suggest that the model provides a plausible account of human solutions processes, based on the facts that (a) it provided a good fit to the performance of human subjects, (b) it is based on clustering of cities in a manner similar to that reported by subjects, (c) the architecture and processing resemble the neurological properties of the human visual system, and (d) its computational complexity (a quadratic function of the number of nodes) is not much greater than the linear function of n observed in human subjects. On the other hand, the results of Kong and Schunn (2007), reviewed at a later point, have been less supportive of the pyramid model.

A more recent implementation (Pizlo et al., 2006) uses the same general pyramid approach but includes components that mimic the differential effects of central and peripheral vision. Unlike the conventional pyramid approach, where the entire image is represented at the lowest layer, the foveating pyramid (FP) model has information about a limited part of the image around the “fixation point” corresponding to foveal vision. At

the next higher level, more of the image may be represented around the fixation point, but at a lower level of resolution. The FP produces successive approximations to a solution in a similar top-down fashion as the original pyramid model, except that connections at the highest level of resolution emerge in limited regions of the problem corresponding to the area of fixation. Once the cities in this area have been incorporated into the emerging solution, the model's "foveal attention" shifts (clockwise or counterclockwise) to the next neighboring region. Unlike the original pyramid model, the FP model has a parameter, k , corresponding to the magnitude of local search. This makes it possible to fit model solutions to individual data, rather than simply to the average of all participants. A further strength of the pyramid approach is its potential generality, and it has been successively applied to account for human solutions to the 15-puzzle, described previously (Pizlo & Li, 2005).

4.3. The Crossing-avoidance Hypothesis

MacGregor and Ormerod (1996) and Graham, Joshi, and Pizlo (2000) observed that human solutions rarely produced crossed arcs. Combining the results over experiments from several laboratories van Rooij, Stege, and Schactman (2003) reported that crossed arcs are indeed rare, occurring in about 6% of tours. It is known that optimal tours cannot contain crossed arcs (Flood, 1956), and on this basis van Rooij, Stege, and Schactman suggest that crossing avoidance may be a mechanism underlying human performance with TSPs, and may also have utility in accounting for performance with variants such as the open TSP (see above) where the convex hull cannot be utilized in the usual way. Vickers et al. (2003) systematically varied the number of possible intersections in 40-city problems, and found a main effect of number of intersections: performance in terms of tour length improved as the number of potential intersections increased. Furthermore, a measure of response uncertainty decreased as the number of potential intersections increased. Vickers et al. (2003) interpreted these findings as evidence in favor of crossing avoidance, and suggested a mechanism by which mental mechanisms could implement this strategy. They argued in favor of a hierarchical nearest neighbor process, which they characterize as "local-to-global." In this approach, clusters of nodes are first established based on nearest neighbor distances, then clusters are linked according to a nearest cluster algorithm. This process is conceptually related to a general model of visual perception, developed by Vickers, which emphasizes the importance of minimal structures (Vickers et al., 2003). Some data speak against the importance of crossing avoidance. For example, MacGregor, Chronicle, and Ormerod (2004) showed that a heuristic based on crossing avoidance alone was unable to account for the average performance of 45 participants on 6 TSP instances. In addition, performance on a contrived TSP with a single interior point was considerably better accounted for by a convex-hull model than a heuristic that combined crossing avoidance with a nearest neighbor decision rule (MacGregor, Chronicle, & Ormerod, 2004). In

addition, specially constructed stimuli that represent strong tests of the crossing-avoidance hypothesis have been found to increase the rate of crossings from the typical rate of 6% of tours reported by van Rooij, Stege, and Schactman (2003), to more than 40% of tours (MacGregor, Chronicle, & Ormerod, 2004; Tak, Plaisier, & van Rooij, 2008).

4.4. Global-local Models

Best (2005) has proposed a global-local model (GL-TSP) that appears to incorporate elements of both the sequential convex hull and pyramid approaches. In the initial global phase, groups of nodes are identified using a procedure related to hierarchical clustering. Then, an initial planned path is constructed using the clusters as nodes and applying a convex-hull cheapest-insertion procedure to connect them (as in the sequential convex-hull model). This completes the global phase. During the local stage, multiple paths through the nodes within the current cluster (and adjacent clusters, if necessary up to a maximum of six nodes) are evaluated in terms of their complete distance and the first node along the selected path (presumably the shortest path) is included. The local process then repeats from the new current node. Best (2005) tested the model by comparing model and human solution times and tour lengths using 10-, 20-, and 30-node problems, and concluded that the model provided a good fit. In particular, model solution times approximated the linear increases in human solution times with increasing problem size.

Kong and Schunn (2007) have also proposed a model that involves a global phase of clustering and cluster-connection, followed by a local phase, which they refer to as the *K*-means TSP model. In this model, clusters are identified using a *K*-means clustering algorithm. Next, a tour through the initial set of clusters is sketched. However, instead of using the convex hull, a Spline-curve is drawn through the cluster centroids, based on the shortest tour through them. This is calculated by brute force, and so the number of centroids is controlled to be five or fewer. The forming of the Spline-curve completes the global phase. Next, the nodes are projected onto the nearest point on the Spline-curve. Nodes are then connected together in order of adjacency of their projections on the Spline-curve. The *K*-means model was tested against several other models, including the sequential convex-hull and pyramid models, using data from 10 participants who solved 20 TSPs, ranging from 10 to 100 nodes. Comparisons between model and human results included number of crossings, mean accuracy, standard deviation of accuracy, and exact path correlations (reflecting the correspondence between model and human solutions at the level of node-to-node connections). The authors concluded that the sequential convex-hull model and *K*-means model were the best predictors of human performance overall. However, the simulation using the pyramid model does not appear to have varied the local search parameter to fit individual data, and consequently may not represent a test of the full model.

4.5. General Characteristics of Models

For each of the implemented models there are data suggesting that the model provides a good fit to human performance. Consequently, distinguishing between models may require tactics other than simple data fitting. One approach, suggested by Tak, Plaisier, and van Rooij (2008), is to devise and implement “strong” tests of models. While this is laudable for mature theories, a potential concern is that, at the present stage of development, the approach may eliminate *all* models. A further, and more serious, concern is that strong tests may be more likely to discredit well-specified models than poorly specified ones. An alternative approach to theory testing may be to identify the higher-order processes that different models hypothesize, and then test for those higher-order processes in human solutions. This would allow the testing of different *classes* of models, regardless of how well-specified they may be. For example, the different theoretical approaches vary in the ways in which they involve local versus global processes. While most models implicate both, Vickers et al. (2003) have characterized their approach as “local-to-global,” meaning that local processes are primary. This approach stands in contrast to the “global-to-local” direction of the sequential convex-hull and other models. MacGregor, Chronicle, and Ormerod (2006) have argued that if local processes, such as those based on point-to-point proximities, are primary, then there is no reason to expect that performance would change between versions that differ only on a global characteristic, such as whether the solution requires an open path or a closed tour. However, as noted above, results for open versus closed versions of TSPs have been found to differ systematically (Chronicle, MacGregor, & Ormerod, 2006; Vickers et al., 2003). In this respect, the evidence appears to favor global-to-local models.

At the same time, the global-to-local models differ among themselves in several respects. One is in their use of hierarchical clustering, which is implemented in all but the sequential convex-hull model. A second is in whether or not they abstract an outline tour that guides subsequent local decisions. All but the pyramid model do this. Finally, those that abstract an outline do so in different ways. In the *K*-means model the outline takes the form of a Spline-curve through the cluster centroids, which remains stable throughout subsequent processing. In the sequential convex-hull and GL models, the outline is based initially on the convex hull, but this shape evolves with the incorporation of each node. Kong and Schunn (2007) interpret the stability of the Spline-curve outline as an advantage of their *K*-means model, since maintaining a stable outline requires less processing capacity than an outline that changes. On the other hand, it is conceivable that having a general plan that can adapt to subsequent decisions could have other benefits. We pursue this issue below, when we review the evidence for the different general characteristics of models.

Within the global-to-local group, a major difference is in whether or not a model

generates a global outline. A relevant experiment has been reported by Kong and Schunn (2007). The experiment used 20 TSPs, varying in size from 10 to 100 nodes, to compare human performance under 3 conditions of global information. In the control condition, all nodes of a problem were displayed throughout the solution phase. A second condition, Global Preview, consisted of three phases. In phase 1, all nodes of a problem were displayed initially. A participant was provided pen and paper to draw any global information that might be needed during the solution phase. In the second phase, the nodes were clustered using a K -means algorithm, and the clusters were displayed as larger dots. Participants chose a path through the larger dots, which then determined the order of clusters in the third phase. In phase 3, all points were hidden, then clusters were presented one at a time in the order determined in phase 2. The participant connected the nodes within a cluster, then the next cluster was displayed. In the third condition (No Global Preview), clusters were presented one by one, as in the third phase of the Global Preview condition, but without seeing the full set of points initially. In other words, phases 1 and 2 were omitted. (In the third condition, it is unclear how the order of presentation of clusters was determined.)

The results indicated significant performance differences between all three conditions. The control group produced the shortest paths, followed by the Global Preview group, while the No Global Preview group produced the longest paths. The authors concluded that global information stored only mentally helps local search (consistent with the K -means model), while global information available throughout the task helps even more (consistent with the sequential convex-hull and global-local models). The experiment also found a significant difference between natural and random TSPs for the control condition only, which no model was able to duplicate. (Natural TSPs were real-world problems and tended to be more clustered than random TSPs.) However, a variation of the K -means model was able to produce similar results. Several questions of experimental control could influence the interpretation of the findings. For example, it is unclear whether performance was affected in the second and third conditions because participants were required to use clusters formed by a third party. Also, it seems likely that the amount of exposure to each stimulus problem may have decreased in order across the three conditions. Nevertheless, the study illustrates how it may be possible to design experiments to tease apart the contributions of global and local processes.

4.6. Theoretical Accounts of the Open Version

While the development and testing of theoretical models has been limited mainly to the classic version of the TSP, several researchers have theorized about the processes involved in the open version.

In an earlier section we reported a study by Hirtle and Gärling (1992) where participants solved open problems of 6, 10, and 18 nodes. A primary focus of the research was

to identify the heuristics used by participants. Based on previous research by Gärling, two heuristics were initially considered. The first was the familiar nearest neighbor (NN) heuristic. The second, referred to as the straight line heuristic (SL), posits that a set of collinear points will be connected in order from one end of the implied line, rather than starting in the middle. This heuristic is selective, and likely to apply only to sections of a problem. In addition, Hirtle and Gärling proposed a clustering heuristic. They proposed that nodes would first be clustered, then clusters would be ordered (using NN or SL), and finally, locations within clusters would be connected using the same heuristics by which clusters were ordered or, under conditions of high load, at random. A fourth heuristic was identified to account for some of the data, termed the zigzag heuristic (ZZ), in which the path moves vertically or horizontally back and forth from one edge of the space to the other. Overall, the authors found evidence of all four heuristics. With the six-node problems, many of the participants' solutions were completely consistent with the NN heuristic, although there were also instances consistent with the operation of the other heuristics. For the larger problems, however, no single participant solution conformed completely to the NN approach, although large sections of some solutions did.

MacGregor, Chronicle, and Ormerod (2006) used data from the open version of the task to assess the degree of fit of three different heuristics. One was the NN heuristic. A second was crossing avoidance (CA), which they implemented in terms of the following rule: from the start, connect the path to the closest point on the *x*-axis.⁴ Repeat until reaching the end point. (This is consistent with the ZZ heuristic of Hirtle and Gärling, 1992). The third was a convex-hull heuristic (CH), which conformed to the following rule: from the start, connect the path to the closest point that is not the end point that can be reached without crossing the convex hull around the points not yet on the path. Repeat until only the end point remains; then connect to it.

Heuristic performance was compared with participants in terms of overall path length, number of crossings, and the frequencies of node-to-node connections. While the results overall supported the convex-hull over the other two heuristics, MacGregor, Chronicle, and Ormerod (2006) reported finding evidence for all three. The presence of multiple heuristics is consistent with results described below, which compared performance of normal and lesioned subjects.

Cutini et al. (2008) have proposed a connectionist model to describe performance in the open TSP and to account for differences observed between healthy participants and patients with frontal traumatic brain injury (fTBI). While the model consists of a number of components, the core is a competitive-learning network that selects from among three available heuristics. The heuristics used were NN, DR (direction right, which connects to

⁴ The *x*-axis was chosen because the problems tested had lateral paths, starting on the left and ending on the right, or vice versa. The *y*-axis would be appropriate for vertically oriented paths. Also, we note that this is only one way that crossing avoidance may be implemented.

the nearest node on the x -axis, as in the CA heuristic, above) and DD (direction down, which connects to the nearest node on the y -axis). The model operates sequentially node by node and a unique characteristic is its capacity to switch heuristics for each decision. A version of the model, which repeats the competitive selection process at each step, was found to be superior to versions that maintained a heuristic either throughout the process or until a mismatch was signaled. That is, continuously monitoring and reselecting from among available heuristics provided the best model fit to the performance of healthy participants. In contrast, inhibiting the ability of the model to switch heuristics produced data that better fit the performance of FTBI patients. This and related extensions of the TSP task into neuropsychological research are further discussed below.

5. Broader Developments

In addition to the basic experimental research described above, studies involving the TSP have been reported within a diverse range of fields, including animal studies, spatial cognition, and neuropsychology. A sample of this research is summarized here.

5.1. Animal Studies

A similarity has often been noted between the TSP and foraging in animals, where the animal leaves a home base to visit several feeding sites before returning home. In the animal world, the costs of inefficient foraging routes may include decreased resource intake, increased energy outlay, and greater exposure to predators (Anderson, 1983). Consequently, using efficient routes is likely to be adaptive.

An early study tested young chimpanzees in a variation of the TSP where food rewards were placed in up to 18 locations in an open field. The experimenter placing food in the randomly selected sites was accompanied by a second experimenter carrying the animal to be tested. On release, the chimpanzee followed a path without regard to the one that it had been carried in and that approximated the shortest route (Menzel, 1973).

A similar study using vervet monkeys found that they were able to remember no more than six locations. Nevertheless, their choices indicated an ability to plan a foraging route up to three steps ahead and a sensitivity to shorter routes. One experiment used a diamond-shaped path with rewards hidden at each corner. When the route was one-way (open path), the optimal required crossing at the shortest diagonal in a zigzag path. For a round trip (closed tour) the optimal followed the perimeter of the diamond. The monkeys were sensitive to this variation, and in each version reliably selected the optimal route (Gallistel & Cramer, 1996). In a task consisting of eight target locations, MacDonald and Wilkie (1990) reported that yellow-nosed monkeys tended to minimize the total distance traveled in retrieving rewards from the targets.

Although most studies with animals appear to have used primates, several have

involved other species. Bureš, Burešová, and Nerada (1992) studied rats required to visit six target locations arranged in two different configurations, and reported a preference for visiting the targets in the shortest tour (> 40% of trials compared with an expected 5% by chance). In a rare, and possibly unique, comparison of cross-species performance, Gibson, Wasserman, and Kamil (2007) conducted a study of human subjects and pigeons using open path problems of three, four, and five targets. Pigeon solutions were reliably shorter than random solutions and longer than human solutions. In a further experiment, pigeon performance was substantially improved by increasing the costs of inefficient routes. Although this study presented stimuli using a computer screen, most animal studies have used natural terrains. While rarely used in human studies, natural or artificial terrains have been employed in several of the studies described below.

5.2. Spatial Cognition

There has been a long and continuing interest in TSP-type tasks in research that might be described as “spatial cognition,” encompassing a number of behavioral science disciplines, including behavioral geography, environmental psychology, and market research. The following section briefly summarizes a selection of these works.

Gärling and Gärling (1988) reported a study in which shoppers returning to a parking lot were interviewed about their route selection. Of those visiting more than two locations, the majority (69%) attempted to minimize the distance traveled, frequently by first selecting the most distant location on their route. The authors suggested that, in addition to minimizing distance traveled, this could reflect a goal to minimize the effort involved in carrying purchases. The example illustrates the potential contaminating influences of optimizing on multiple criteria when researching in a naturalistic setting.

Hui, Fader, and Bradlow (2009) used data from a large supermarket where shopping carts were fitted with a radio device that emitted a unique signal every five seconds and provided a plot of the location of a cart over time. To the extent that the cart location approximates the location of the customer using it, the results recorded the tour followed by a shopper from the entrance of the store to the checkout counter. The resulting dataset contained a total of almost 1000 tours and the purchase records associated with them. The latter provided information about the number and types of items selected by customers. Hui, Fader, and Bradlow used the data to address the questions of whether shoppers tend to be “optimal” in their shopping patterns and in what ways they depart from optimality. For each customer, an optimal path was calculated, defined as the shortest path that started from the store entrance, connected all of the recorded purchases and ended at checkout. The customer’s recorded path was then compared with the “optimal,” and two different types of deviation were noted. The first occurred when the order of picking up items deviated from the optimal order (“order deviation”). The second occurred when the shopper failed to take the shortest path from one product location to the next (“travel

deviation").⁵ Thus, any given path could be decomposed into three additive components, and where the total length equals the TSP-optimal path plus order deviation plus travel deviation. The findings showed that while total lengths departed considerably from the optimal, this was due mainly to travel deviations, suggesting that shoppers may have made substantial detours while en route to their next planned location. However, the order of visiting product locations was much closer to the optimal order, and the path length attributable to the order component of selected routes was only 15% above the optimal. Thus, the error rate for the order of selecting nodes appears to be of the same order of magnitude as in typical laboratory studies, which is impressive, given the increased complexity both of the context and of the data collection and analysis procedures.

Wiener, Ehbauer, and Mallot (2009) used a simulated shopping task where participants navigated between targets placed on a 5 x 5 meter grid on the floor of a large laboratory. A space of this scale is referred to as "vista space," as it can be fully apprehended from a single location. Vista spaces share characteristics of smaller, pictorial, spaces and larger, environmental, spaces. Like the former, they can be seen in entirety from a single viewpoint while, like the latter, spatial relationships change when moving through the space. In the experiments, 25 small cardboard pillars were arranged in a regular 5 x 5 grid, and 25 symbols were randomly distributed about the pillars so that each pillar had a different symbol visible on its top surface. Participants were given a "shopping list" of a subset of target symbols. The task consisted of navigating the shortest route connecting a given starting location with all of the target locations and returning to the start. Participants solved 36 TSPs ranging from 5 to 10 locations.

In Experiment 1, the tasks consisted of three different subtypes: NN-adequate, where the shortest path conformed to the NN heuristic; NN-inadequate, where it did not; and cluster/NN-ambiguous, where the targets fell into two distinct clusters and where the NN heuristic would generate multiple solutions. Overall, performance was similar to that found by previous research using "pictorial spaces" (e.g., presented on paper or LCD screen). Optimal paths were found on almost 50% of trials and the average percentage above optimal (PAO) was 6%. Male and female performance did not differ significantly (5% above optimal for males, 7% for females). There were no significant main or interaction effects on percentage above optimal involving the three different types of task. PAO on both the NN-inadequate and cluster/NN-ambiguous tasks was significantly better than predicted by a NN strategy. In the cluster tasks, participants showed a modest but significant preference to visit the larger cluster first (on 59% of trials compared to 50% by chance). The result is consistent with the findings of some animal studies (Cramer & Gallistel, 1997). Post-task interviews uncovered two common strategies. In one, participants

5 In paper-and-pencil and point-and-click versions of the TSP task, an "order deviation" occurs when nodes are selected in the wrong order, but since the participant is credited with the shortest distance between adjacent selections, the issue of "travel deviation" does not arise.

reported dividing the 25 potential locations into regions, assigning targets to regions, and plotting a coarse path through the regions. In the second, participants selected a subset of the target symbols based on some criterion (e.g., color) and plotted a coarse route through this subset of symbols. In both cases, the strategies appear to apply a hierarchical global-to-local planning scheme, where a coarse outline is plotted and then target locations are inserted at the local level during navigation through the space. A second experiment further supported the reliance of performance on a region-based hierarchical planning scheme.

To summarize, the results suggest that when TSP problems are presented in environments that go beyond “pictorial” space, performance remains quantitatively and qualitatively similar to that observed in the more commonly used paper-and-pencil and point-and-click representations. Path lengths are close to optimal and reflect a process that is superior to a local, nearest neighbor strategy. The strategies reported bear strong similarities to those that have been proposed to account for data collected in more artificial settings. However, whether such similarities will extend to larger, more topographically varied, environments, remains to be seen.

5.3. Neuropsychology

Vickers and Lee (1998) noted a similarity between the TSP and the Halstead-Reitan Trail Making Test (TMT), one of the most widely used neuropsychological instruments in assessing brain damage. Consistent with this observation, a number of neuropsychological studies have used TSP-type tasks as a basis for detecting differences between normal and functionally impaired performance.

In one study, Bisiacchi, Sgaramella, and Farinello (1998) compared the performance of four groups: young adults, two groups of older adults (between 70 and 80, and over 80), and a group of closed head injury patients (CHI). The task was a paper-and-pencil test, which simulated a shopping trip involving 10 errands to different destinations shown on a map. Although the goal was to complete the task using the shortest distance, it differed from a standard TSP because of the presence of certain constraints, such as different opening and closing times for some destinations. In addition to measuring performance, the researchers examined solutions for indications of three heuristics—global planning, NN, and clustering. Compared with the other groups, the CHI patients were found to perform more poorly on an overall “optimization” score, to revisit destinations more often, and to have more difficulty in applying the cluster and NN heuristics.

Basso et al. (2001) also compared a CHI group with uninjured participants, matched for age and education. Using 10-node problems, the researchers measured performance on their Maps test. The authors reported that in previous research with the method, normal participants appeared to plan only a small section of the route before starting and needed to make ongoing decisions during path execution. To capture this aspect of performance

in the present study, the authors tracked the variability in planning time prior to the execution of each step of the solution. Normal participants showed a relatively constant planning time at each step, whereas CHI participants showed much greater variability.

Basso et al. (2006) used the same task to compare the performance of normal adults operating with and without repetitive transcranial magnetic stimulation (rTMS), which is considered to exert inhibitory effects on the targeted brain region. The experiment used a counterbalanced within subjects design where half of the rTMS trials were applied to the frontal regions (areas F3 and F4—the active condition) and half to the left parieto-occipital position (area PO3—the placebo condition). The hypothesis was that spatial planning for strategy change would be inhibited in the active condition, resulting in longer path lengths and fewer shifts in the heuristics used during the execution of a path. Participants' solution path segments were analyzed for the presence of four possible heuristics: NN, cluster, direction down, and direction across (previously described in Section 4.6). The results indicated that, when in the active condition, participants did indeed produce longer paths and showed decreased heuristic switching during path execution. This same pattern of results has been observed on the same task when the performance of normal participants was compared with patients with traumatic frontal lobe brain injury (Cutini et al., 2008).

6. Future Directions

During the past 15 years, several groups of psychologists have turned their attention to human performance on optimization problems. Intriguingly, in this domain, untrained humans appear to have the ability to perform well on problems where the number of alternatives is extremely large, and certainly well beyond the capacity to search and evaluate. Furthermore, the mental mechanisms underpinning this performance appear to be of low complexity—a general finding (albeit with some exceptions) is that human performance deteriorates linearly as the problem size increases exponentially. The startling ability of humans to draw—for example—an optimal TSP tour without training and in a very modest time (typically a few seconds) when the number of alternatives solutions may number in the billions or trillions is in stark contrast to other areas of reasoning and problem solving. For example, the inability of humans to select the correct cards on the Wason selection task, a test of conditional reasoning, is legendary in experimental psychology (Johnson-Laird & Wason, 1977), even though the task requires the evaluation of only four alternatives. In the area of insight, the inverted triangle problem is solved on the first attempt by only 9% of participants, despite its relatively small state space (81,147 states, compared to 6.08×10^{16} for a 20-city TSP; see Chronicle, MacGregor, & Ormerod, 2004).

In the case of the TSP, although there are differences between the theoretical accounts of human performance existing in the literature, a common thread is that mental

mechanisms that exist to serve other perceptual needs are brought into service when the naive human is faced with a problem. Algorithms implemented by the visual system, to solve the complex problems of figure-ground segregation, boundary detection, and object recognition (Marr, 1982), may be at the root of human performance with this problem. Other optimization problems may also tap into visual system algorithms: for example, in the p-median problem the task may exploit the ability of humans to judge the center of mass of a planar object, which is known to be very accurate (Bingham & Muchisky, 1993). As yet, no algorithms have been put forward to explain performance on the MSTP and the GSTP, although it is reasonable to speculate that performance on both may relate to perception of object axes. It is more difficult to account for human performance on the 15-puzzle in terms of visual system algorithms.

In conclusion, knowledge of human abilities with optimization problems is exciting cognitive scientists for several reasons. First, optimization problems represent a new domain for those interested in the cognitive psychology of problem solving—one in which the challenge is to understand human success rather than failure. Second, investigating the nature of the cognitive processes involved may add to knowledge about algorithms implemented by the visual system. Third, the apparently low complexity of these algorithms may transfer to inform the development of new procedures for optimization in applied settings. Several challenges lie ahead, not the least of which are conducting experiments that will provide critical tests of competing theoretical positions, and accounting for the seemingly contradictory data concerning individual differences in optimization problem solving. Nonetheless, we anticipate that studies of human optimization will become a fertile interdisciplinary area over the next decade.

References

- Agarwala, R., Applegate, D. L., Maglott, D., Schuler, G. D., & Schäffer, A. A. (2000). A fast and scalable radiation hybrid map construction and integration strategy. *Genome Research, 10*, 350–364.
- Anderson, D. J. (1983). Optimal foraging and the traveling salesman. *Theoretical Population Biology, 24*, 145–159.
- Applegate, D. L., Bixby, R. E., Chvátal, V., & Cook, W. J. (2006). *The Traveling Salesman Problem: A Computational Study*. Princeton University Press.
- Basso, D., Bisiacchi, P. S., Cotelli, M., & Farinello, C. (2001). Planning times during travelling salesman's problem: Differences between closed head injury and normal subjects. *Brain and Cognition, 46*, 38–42.
- Basso, D., Lotze, M., Vitale, L., Ferrer, F., Bisiacchi, P. S., Olivetti-Belardinelli, M., Rossini, P. M., & Birbaumer, N. (2006). The role of prefrontal cortex on visuo-spatial planning: A repetitive-TMS study. *Experimental Brain Research, 171*, 411–415.

- Best, B. J. (2005). A model of fast human performance on a computationally hard problem. *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, pp. 256–261.
- Bingham, G. P., & Muchisky, M. M. (1993). Center of mass perception and inertial frames of reference. *Perception & Psychophysics*, *54*(5), 617–632.
- Bisiacchi, P. S., Sgaramella, T. M., & Farinello, C. (1998). Planning strategies and control mechanisms: Evidence from closed head injury and aging. *Brain and Cognition*, *37*, 113–116.
- Bland, R. E., & Shallcross, D. F. (1987). "Large Traveling Salesman Problem Arising from Experiments in X-ray Crystallography: a Preliminary Report on Computation," Technical Report No. 730, School of OR/IE, Cornell University, Ithaca, New York.
- Bureš, J., Burešová, O., & Nerad, L. (1992). Can rats solve a simple version of the traveling salesman problem? *Behavioral Brain Research*, *52*, 133–142.
- Burns, N. R., Lee, M. D., & Vickers, D. (2006). Are individual differences in performance on perceptual and cognitive optimization problems determined by general intelligence? *Journal of Problem Solving*, *1*(1), 5–19.
- Cazzato, V., Basso, D., Cutini, S., & Bisiacchi, P. (2010). Gender differences in visuospatial planning: An eye movements study. *Behavioural Brain Research*, *206*, 177–183.
- Chronicle, E. P., MacGregor, J. N., Lee, M., Ormerod, T. C., & Hughes, P. (2008). Individual differences in optimization problem solving: Reconciling conflicting results. *Journal of Problem Solving*, *2*, 41–49.
- Chronicle, E. P., MacGregor, J. N., & Ormerod, T. C. (2006). Optimizing and "pessimizing"—human performance with instructional variants of the travelling salesperson problem. *Journal of Problem Solving*, *1*, 74–82.
- Chronicle, E. P., MacGregor, J. N., Ormerod, T. C., & Burr, A. (2006). It looks easy! Heuristics for combinatorial optimisation problems. *Quarterly Journal of Experimental Psychology*, *59*, 783–800.
- Cramer, A. E., & Gallistel, C. R. (1997). Vervet monkeys as travelling salesmen. *Nature*, *387*(6632), 464–464.
- Cutini, S., Ferdinando, A., Basso, D., Bisiacchi, P. S., & Zorzi, M. (2008). Visuospatial planning in the travelling salesperson problem: A connectionist account of normal and impaired performance. *Cognitive Neuropsychology*, *25*, 194–217.
- Dallari, F., Marchet, G., & Ruggeri, R. (2000). Optimisation of man-on-board automated storage/retrieval systems. *Integrated Manufacturing Systems*, *11*, 87–93.
- Dry, M. J., Lee, M. D., Vickers, D., & Hughes, P. (2006). Human performance on visually presented traveling salesperson problems with varying numbers of nodes. *Journal of Problem Solving*, *1*, 20–32.
- Ferguson, T. S. (1989). Who solved the Secretary Problem? *Statistical Science*, *4*, 282–296.
- Flood, M. M. (1956). The travelling salesman problem. *Operations Research*, *4*, 61–75.
- Gallistel, C. R., & Cramer, A. E. (1996). Computations on metric maps in mammals: Getting

- oriented and choosing a multi-destination route. *Journal of Experimental Biology*, 199, 211–217.
- Gärling, T. (1989). The role of cognitive maps in spatial decisions. *Journal of Environmental Psychology*, 9, 269–278.
- Gärling, T., & Gärling, E., (1988). Distance minimization in downtown pedestrian shopping. *Environment and Planning A*, 20, 547–554.
- Gibson, B. M., Wasserman, E. A., & Kamil, A. C. (2007). Pigeons and people select efficient routes when solving a one-way “traveling salesperson” task. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 244–261.
- Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The travelling salesman problem: A hierarchical model. *Memory & Cognition*, 28, 1191–1204.
- Hill, A. V. (1982). An experimental comparison of human schedulers and heuristic algorithms for the traveling salesman problem. *Journal of Operations Management*, 2, 215–223.
- Hirtle, S. C., & Gärling, T. (1992). Heuristic rules for sequential spatial decisions. *Geoforum*, 23, 227–238.
- Hui, S. K., Fader, P. S., & Bradlow, E. T. (2009). Research note—the traveling salesman goes shopping: The systematic deviations of grocery paths from TSP optimality. *Marketing Science*, 28, 566–572.
- Johnson-Laird, P. N., & Wason, P. C. (1977). A theoretical analysis of insight into a reasoning task. In P. N. Johnson-Laird and P. C. Wason (Eds.), *Thinking: Readings in Cognitive Science* (pp. 143–157). Cambridge University Press.
- Kong, X., & Schunn, C. D. (2007). Global vs. local information processing in visual/spatial problem solving: The case of traveling salesman problem. *Cognitive Systems Research*, 8(3), 192–207.
- 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*. Wiley.
- Lee, M. D., Zhang, S., Munro, M., & Steyvers, M. (2010). Psychological models of human and optimal performance in bandit problems. *Cognitive Systems Research*. In press.
- Lee, M. D., & Vickers, D. (2000). The importance of the convex hull for human performance on the traveling salesman problem: A comment on MacGregor and Ormerod (1996). *Perception & Psychophysics*, 62(1), 226–228.
- MacDonald, S. E., & Wilkie, D. (1990). Yellow-nosed monkeys' (*Cercopithecus ascanius whitei*) spatial memory in a simulated foraging environment. *Journal of Comparative Psychology*, 104, 382–397.
- MacGregor, J. N., Chronicle, E. P., & Ormerod, T. C. (2004). Convex hull or crossing avoidance? Solution heuristics in the travelling salesperson problem. *Memory & Cognition*, 32, 260–270.
- MacGregor, J. N., Chronicle, E. P., & Ormerod, T. C. (2006). A comparison of heuristic and human performance on open versions of the traveling salesperson problem. *Journal of Problem Solving*, 1(1), 33–43.

- 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 & Cognition*, *28*, 1183–1190.
- Marr, D. (1982). *Vision*. W. H. Freeman.
- Menzel, E. W. (1973). Chimpanzee spatial memory organization. *Science*, *182*, 943–945.
- Michie, D., Fleming, J. G., & Oldfield, J. V. (1968). A comparison of heuristic, interactive and unaided methods of solving a shortest-route problem. In D. Michie (Ed.), *Machine Intelligence*. Elsevier.
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, *9*(3), 353–383.
- 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.
- Pizlo, Z., & Li, Z. (2005). Solving combinatorial problems: 15-puzzle. *Memory & Cognition*, *33*(6), 1069–1084.
- Pizlo, Z., Stefanov, E., Saalweachter, J., Haxhimusa, Y., & Kropatsch, W.G. (2006). Traveling salesman problem: A foveating pyramid model. *Journal of Problem Solving*, *1*, 83–101.
- Polivanova, N. I. (1974). On some functional and structural features of the visual-intuitive components of a problem-solving process. *Voprosy Psikhologii [Questions of Psychology]*, *4*, 41–51.
- Pomerantz, J. R. (1981). Perceptual organization and information processing. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual Organization* (pp. 141–180). Erlbaum.
- Ragni, M., & Löffler, C.M. (2010). Complex problem solving: Another test case? *Cognitive Processes*, *11*, 159–170.
- Sangalli, A. (1992). Why sales reps pose a hard problem. *New Scientist*, *12*, 24–28.
- Sanocki, T. (1993). Time course of object recognition: Evidence for a global-to-local contingency. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 878–898.
- Schrijver, A. (2003). *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag.
- Tak, S., Plaisier, M., & van Rooij, I. (2008). Some tours are more equal than others: The convex-hull model revisited with lessons for testing models of the Traveling Salesperson Problem. *Journal of Problem Solving*, *2*, 4–28.
- van Rooij, I., Schactman, A., Kadlec, H., & Stege, U. (2006). Perceptual or analytical processing? Evidence from children's and adult's performance on the Euclidean Traveling Salesperson problem. *Journal of Problem Solving*, *1*(1), 44–73.

- 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 & Cognition, 31*, 215–220.
- Vickers, D., Bovet, P., Lee, M. D., & Hughes, P. (2003). The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean travelling salesperson problems. *Perception, 32*, 871–886.
- Vickers, D., Butavicius, M. A., Lee, M. D., & Medvedev, A. (2001). Human performance on visually presented traveling salesman problems. *Psychological Research, 65*, 34–45.
- Vickers, D., & Lee, M. D. (1998). Never cross the path of a traveling salesman: The neural network generation of Halstead-Reitan trail making tests. *Behavior Research Methods, Instruments, & Computers, 30*, 423–431.
- Vickers, D., Lee, M. D., Dry, M., & Hughes, P. (2003). The roles of the convex hull and number of intersections in performance on visually presented travelling salesperson problems. *Memory & Cognition, 31*, 1094–1104.
- Vickers, D., Mayo, T., Heitmann, M., Lee, M. D., & Hughes, P. (2004). Intelligence and individual differences in performance on three types of visually presented optimization problems. *Personality and Individual Differences, 36*, 1059–1071.
- Wiener, J. M., Ehbauer, N. N., & Mallot, H. A. (2009). Planning paths to multiple targets: Memory involvement and planning heuristics in spatial problem solving. *Psychological Research, 73*, 644–658.

Acknowledgment

We are grateful to Ed Chronicle for his contributions to an earlier version of this paper.

Paper submitted: July 29, 2010

Paper accepted: October 23, 2010