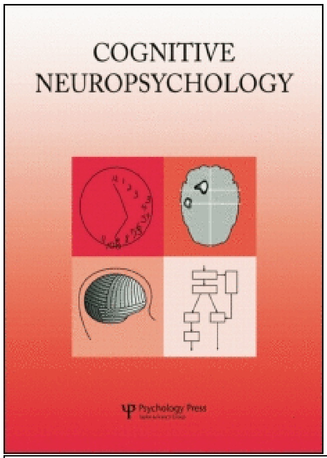

Cognitive Neuropsychology
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713659042

Visuospatial planning in the travelling salesperson problem: A connectionist account of normal and impaired performance

Simone Cutini [a]; Andrea Di Ferdinando [b]; Demis Basso [cd]; Patrizia Silvia Bisiacchi [a]; Marco Zorzi [a]
[a] University of Padova, Padua, Italy
[b] National Research Council, Rome, Italy
[c] University of Pavia, , Pavia, Italy
[d] University of Pisa, Pisa, Italy

First Published on: 06 October 2007

To cite this Article: Cutini, Simone, Di Ferdinando, Andrea, Basso, Demis, Silvia Bisiacchi, Patrizia and Zorzi, Marco (2007) 'Visuospatial planning in the travelling salesperson problem: A connectionist account of normal and impaired performance', Cognitive Neuropsychology, 25:2, 194 — 217

To link to this article: DOI: 10.1080/02643290701606408
URL: http://dx.doi.org/10.1080/02643290701606408

COGNITIVE NEUROPSYCHOLOGY, 2008, 25 (2), 194–217

Psychology Press
Taylor & Francis Group

Visuospatial planning in the travelling salesperson problem: A connectionist account of normal and impaired performance

Simone Cutini
University of Padova, Padua, Italy

Andrea Di Ferdinando
National Research Council, Rome, Italy

Demis Basso
University of Pavia, Pavia, Italy, and University of Pisa, Pisa, Italy

Patrizia Silvia Bisiacchi and Marco Zorzi
University of Padova, Padua, Italy

Planning is a fundamental cognitive function frequently employed in common daily activities. The Travelling Salesperson Problem (TSP), in which participants decide what order between a number of locations optimizes total travel distance, is a paradigm that allows the study of planning and strategy choice. In the TSP, subjects adopt visuo-spatial heuristics to perform the task and operate a continuous monitoring to adapt their behaviour. We present a connectionist model of the TSP that simulates bottom-up and top-down influences observed in the execution of the task. The model accounts for the continuous monitoring observed in healthy participants, and, after a simulated lesion, it also accounts for the decrease of heuristic switching observed in frontal patients and in normal subjects under repetitive transcranial magnetic stimulation (rTMS) over frontal lobe.

Planning is a fundamental cognitive function that is frequently employed in common daily activities. It involves the ability to produce mental representations of future behaviour prior to acting and to reason about its consequences in order to properly choose among the possible courses of action (G. Cohen, 1988). As a complex form of human problem solving, planning requires the cooperation between several cognitive processes, including strategy formation, coordination of mental functions, recognition of goal attainment and storage of representations that can guide behaviour from the initial to the goal state (Carlin et al., 2000).

Correspondence should be addressed to Marco Zorzi, Dipartimento di Psicologia Generale, University of Padova, via Venezia 8, 35131 Padua, Italy (E-mail: marco.zorzi@unipd.it).

This study was supported by grants from Ministero dell'Istruzione Università e Ricerca (MIUR) (P.S.B. and M.Z.) and the University of Padova (M.Z.). D.B. was supported by Grant FIRB RBNE018ET9_003. We are grateful to Rick Cooper and to one anonymous reviewer for their helpful comments and suggestions.

http://www.psypress.com/cogneuropsychology DOI:10.1080/02643290701606408

Planning is often carried out in small units during task performance, rather than in a distinct stage devoted to building an entire plan before its execution (e.g., Basso, Bisiacchi, Cotelli, & Farinello, 2001; Phillips, Wynn, McPherson, & Gilhooly, 2001). Planning requires an incremental process in most real-world situations due to limitations in working memory and control processes. The incremental aspect of planning implies that the plan made before the execution is mainly incomplete or inconsistent (e.g., Hayes-Roth & Hayes-Roth, 1979). Initial decisions can be later modified to develop an efficient strategy—that is, an opportunistic combination of simple schemas that can be activated or inhibited when needed. Indeed, human planning is based on cognitive heuristics (Hayes-Roth & Hayes-Roth, 1979; Hirtle & Gärling, 1992; Murakoshi & Kawai, 2000) that can be defined as behavioural schemas that approximate the correct solution using fewer cognitive resources than does performing an exhaustive algorithm. An efficient strategy requires continuous monitoring during task performance in order to allow for on-line changes of heuristic.

The Tower of Hanoi (ToH), since its introduction as a task to study planning from the information-processing perspective (Simon, 1973), and the Tower of London (ToL; Shallice, 1982) are the most widely used tools to assess planning in cognitive studies. The ToH is a complex problem-solving task that has demonstrated sensitivity to prefrontal lobe function and dysfunction (e.g., Goel & Grafman, 1995). The specific executive processes recruited for successful performance (or, conversely, impaired in prefrontal dysfunction) have been a subject of debate, but it is generally agreed that this task taps planning, working memory, and inhibition (e.g., Goel & Grafman, 1995; Roberts & Pennington, 1996). The ToL, derived from the ToH, has gained high popularity among neuropsychologists. Indeed, the results obtained with the ToL led Shallice and his colleagues (Norman & Shallice, 1986; Shallice, 1982) to the development of their influential theory of executive functions. The ToL has proved extremely valuable for investigating executive functions and their disorders following brain damage, and it has been employed in a wide range of studies. However, a number of potential shortcomings of the ToH/ToL have emerged in the recent years with regard to the planning component of the task. First, although sensitive to frontal lobe damage, the ToL has been questioned with regard to its ability to reliably measure planning skills (Kafer & Hunter, 1997). Second, instructions and cueing given to the participants (e.g., on-line planning vs. full mental plan, or prior information about the minimum number of moves; Phillips et al., 2001; Unterrainer, Rahm, Leonhart, Ruff, & Halsband, 2003), forward-thinking (Owen, Downes, Sahakian, Polkey, & Robbins, 1990; Ward & Allport, 1997), and problem structuring (Goel & Grafman, 1995, 2000; Kaller, Unterrainer, Rahm, & Halsband, 2004) seem to strongly influence task performance.

A task that strongly involves planning and is also representative of many real-world situations is the Travelling Salesperson Problem (TSP): Given a space in which a set of interconnected towns is represented by locations on a map, the task consists in finding an itinerary that visits each town exactly once, returning to the starting town, ensuring that total travelled distance is as short as possible. The TSP is a paradigmatic example of nonpolynomial combinatorial optimization (Lawler, Lenstra, Rinnooy Kan, & Shmoys, 1985) that has been extensively studied by mathematicians and computer scientists but much less by psychologists. Nevertheless, there has been a growing interest in the analysis of human performance in TSP-like problems (Cadwallader, 1975; Gärling, 1989, 1994). The TSP task is thought to be a suitable tool to investigate planning because it can be solved with multiple close-to-optimal solutions that can be evaluated with respect to the single perfect solution (MacGregor & Ormerod, 1996). More specifically, it is reasonable to assume that the TSP involves spatial planning, a type of problem solving that requires optimizing the performance against several constraints, based on spatial elements in the environment. In comparison with other planning tasks, spatial planning requires a stronger interaction between central and peripheral processes: Visual, attentional, and motor issues play a fundamental role, in addition to reasoning, for determining the final behaviour.

Previous studies employing a visually presented TSP (MacGregor & Ormerod, 1996; Polivanova, 1974; Vickers, Butavicius, Lee, & Medvedev, 2001) have revealed that human performance is determined by global perceptual properties to which the visual system is naturally attuned. These properties have been shown to influence the choice of spatially based heuristics that are used to perform TSP tours (Barr & Feigenbaum, 1981; Hirtle & Jonides, 1985; MacGregor & Ormerod, 1996; McNamara, 1992). A variant of the TSP, first proposed by Hirtle and Gärling (1992), introduces a distinction between start- and end-point so that participants have to perform an open path instead of a loop. Behavioural data collected using a computerized version of the open-ended TSP, the Maps Test (Basso, 2005; Basso et al., 2001), showed that three distinct spatially based heuristics are mainly used by human participants to perform the TSP task: the nearest neighbour (NN) heuristic (Barr & Feigenbaum, 1981), the straight-line heuristic (Hirtle & Gärling, 1992), and the direction heuristic (Basso et al., 2001). The first states that each location is recursively chosen on the basis of the minimum local distance from the current position. The straight-line heuristic states that a set of collinear points will be taken in order along the line, rather than starting in the middle. This heuristic has been observed in specific partial configurations, in which points approximately formed a line; it has been frequently observed in conjunction with heuristics based on following a specific direction, such as the direction heuristic. The latter takes place when subjects start from a location placed on a border and reach the next locations following one of the main spatial axes (horizontal or vertical) and a direction (up or down for vertical axis, left or right for horizontal axis). It has been introduced as a modification from the zig-zag heuristic described by Hirtle and Gärling (1992), because a definition provided on the basis of human reference points has been proven to be more suitable in the description of the performance (Bryant, Tversky, & Franklin, 1992). In the Maps test (Basso et al., 2001), the starting city is located typically in the upper left corner, and the end city is located in the bottom right corner. Thus, the horizontal movement was described as representative of a direction right heuristic (DR), whereas the vertical movement was described as a direction down heuristic (DD). The same configuration of TSP is used in the present study to allow us a direct comparison with the behavioural data. Figure 1 shows an example of application of the three heuristics (NN, DR, and DD) considered in the present study.

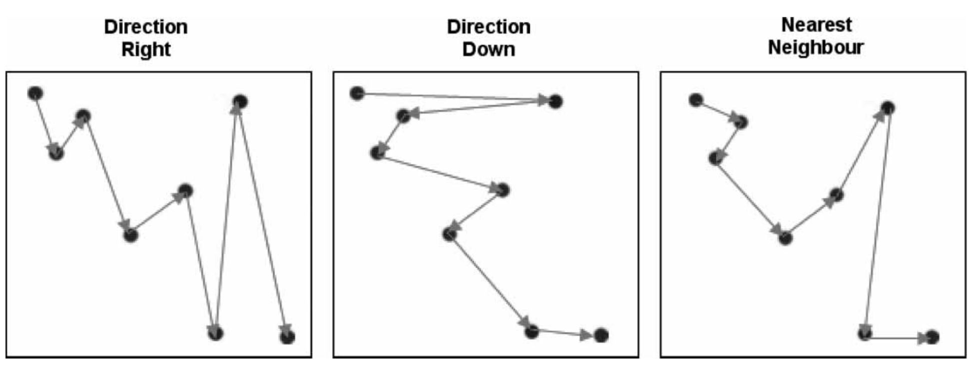

Figure 1. *The figure shows three different tours of the same TSP pattern. Each tour is representative of the use of a unique heuristic along the whole pathway. These heuristics have been implemented in the model using three different saliency maps that bias the choice of the order in which cities are visited.*

The fundamental role of planning in the TSP is confirmed by studies that investigated the effect of lesions or transient neurodisruption of the prefrontal cortex upon performance in the Maps test (Basso et al., 2001; Basso et al., 2006). One key aspect of the TSP task regards the coordination of different heuristics. The use of a single heuristic is inappropriate for most of the maps, and a change of heuristic is necessary to obtain a close-to-optimal solution. Indeed, this is what Basso et al. (2001) observed in the performance of healthy participants. However, the ability to change heuristic during the pathway to optimize performance was markedly impaired in frontal traumatic brain injured (fTBI) patients. A similar pattern of impaired performance was shown by healthy adults under inhibitory repetitive transcranial magnetic stimulation (rTMS) stimulation on the prefrontal cortex (PFC; Basso et al., 2006). Both frontal patients and healthy participants under rTMS did not show the normal pattern of continuous planning: Instead of switching heuristic during the execution of the task, they seemed to apply simple strategies based on only one heuristic (see Figure 2). These results are consistent with the notion that the PFC is a crucial brain area for planning processes and strategy formation (for recent neuroimaging evidence, see Fincham, Carter, van Veen, Stenger, & Anderson, 2002; Newman, Carpenter, Varma, & Just, 2003) and that its lesion is associated with planning deficits (Grafman, 1989, 1995; Lezak, 1995; Shallice, 1982, 1988).

One possible explanation for the finding that lesion or reversible neurodisruption of the PFC leads to a planning deficit in the Maps test is that this region would be crucial for the inhibition of the current heuristic (Basso et al., 2001) whenever a change is necessary to achieve a close-to-optimal performance. That is, the heuristic chosen at the beginning of the task was likely to be kept until the end with no signs of any further consideration of the possible alternative options. These findings are consistent with the presence of perseverative behaviour in frontal patients (Duncan, 1986; Luria, 1980). This behavioural rigidity has been explained in the attention to action (ATA) model by Norman and Shallice (1986; see Cooper & Shallice, 2000, for a computational model) as a

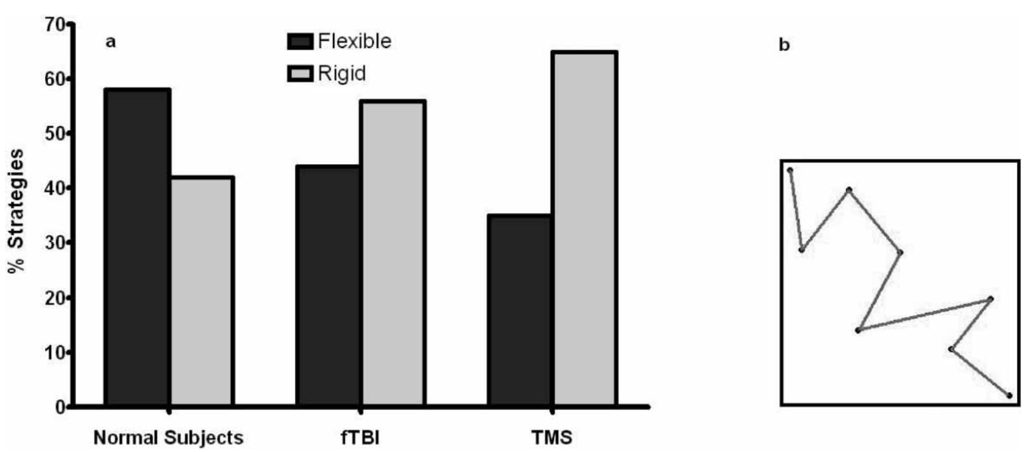

Figure 2. *(a) Percentage of strategies used in the open version of the TSP task (data replotted from Basso et al., 2001). From left to right: normal subjects, traumatic frontal lobe brain-injured patients (f TBI), and healthy participants under repetitive transcranial magnetic stimulation (rTMS). Flexible strategies imply the use of at least two different heuristics to perform a given tour. Rigid strategies are the result of using a single heuristic for the whole tour. The results highlight that while normal subjects often use flexible strategies, subjects under rTMS and traumatic frontal lobe brain injured patients use rigid strategies more often than flexible strategies. (b) An example of the open-version TSP.*

disruption of the supervisory attentional system (SAS). The SAS is thought to perform the inhibition of the behavioural schema automatically selected for the execution and the switch to another schema that emerged as more suitable to the actual situation. In the study of Basso and colleagues (2001), the impairment observed in fTBI patients was interpreted as a failure in controlling and modifying the plan in the execution phase, rather than pure lack of planning. According to this account, patients achieved an acceptable solution because the contention scheduling process was preserved. They therefore selected an appropriate heuristic based on the initial spatial analysis of the TSP configuration. However, a SAS failure did not allow them to modify the initial plan to optimize the execution of the task. As a consequence, their performance was far from optimal because the behavioural rigidity prevented heuristic switches. Therefore, the main drawback of this deficit is a lower level of optimization, which is indexed by longer tours in comparison to the performance of healthy controls.

In summary, there are three fundamental aspects of the TSP task that deserve consideration: (a) the incremental aspect of spatial planning, (b) the presence of visuo-spatial heuristics triggered by bottom-up processes that influence behaviour during the execution of the task, and (c) the crucial role of the frontal lobe to endorse flexible performance. The goal of the present work was to develop a computational model to simulate the cognitive mechanisms underlying the human performance in the TSP task. The main aim was to replicate in the model the incremental aspect of planning, with the interaction of bottom-up and top-down processes, and its disruption following a simulated lesion. A challenging aspect of our modelling enterprise was to implement all these features in a connectionist model that dispenses with the use of explicit rules to guide behaviour. The descriptive adequacy of the model was tested in terms of its fit to the behavioural data from both healthy participants and patients with frontal lobe lesions.

A CONNECTIONIST MODEL OF THE TSP

The computational model is composed by three interconnected modules, with a broad hierarchical organization and feedback connections, which loosely simulate the occipito-parieto-frontal circuit involved in the TSP task (see Figure 3). These components comprise: (a) a visual module, in which the input pattern is processed by Gabor filters (Jones & Palmer, 1987) to simulate the processes responsible for visuo-spatial analysis and perceptual grouping; (b) a competitive selection module that simulates the internal dynamics for the choice of the heuristic; (c) a spatial module encoding the to-be-visited locations and controlling the execution of the pathway in a sequential manner. Moreover, the presence of saliency maps, recurrent connections, and inhibitory mechanisms allows us to simulate the incremental aspect of visuo-spatial planning and the interaction of bottom-up and top-down processes.

Descriptive overview

In building the model we adopted a nested incremental modelling approach (see Perry, Ziegler, & Zorzi, 2007). This strategy, often neglected in psychology, consists in building a new computational model by combining the best features of previous models. Therefore, two main components of our model, the visual module and the spatial module, were simply taken from state-of-the-art computational models of vision and action (Di Ferdinando, Casarotti, Vallar, & Zorzi, 2005; Lee, 1996; Pouget & Snyder, 2000). One advantage of this approach is that, in spite of the complexity of the model, most of the parameters are predetermined and do not influence the ability of the model to fit the human data in this particular task.

Input to the model consists of a digital image displaying the points that constitute the TSP configuration. The patterns reproduced those used for the computerized version of the TSP task (the Maps test; Basso et al., 2001), where the starting

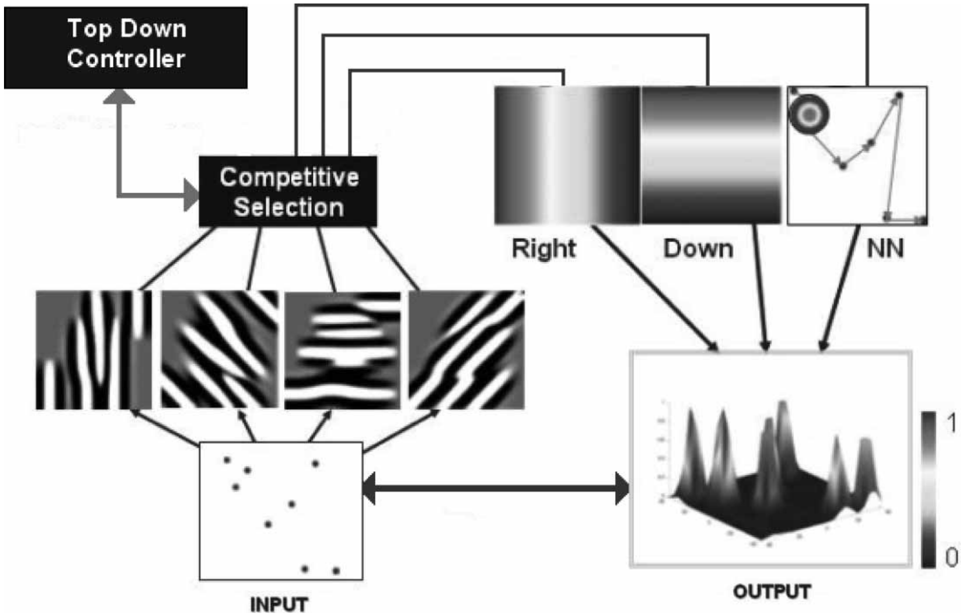

Figure 3. *The architecture of the model. The figure shows the different modules as well as their connectivity.*

point is located in the top-left corner, and the end-point is located in the bottom-right corner of the display. The output consists in a sequence of spatial commands, each encoding the next goal position in space (that is, the position of the next point to be visited). Thus, the model simulates both spatial and temporal aspects of task execution.

One important component of our modelling enterprise was to provide a computational account of the generation of spatial heuristics and of their influence on the planning process. Perceptual mechanisms clearly assume a critical function as the source of bottom-up influence in the generation of spatial heuristics and their successful use. MacGregor and Ormerod (1996) argued that the detection of the minimum path is an innate and natural tendency determined by the human visual system. Along with their suggestion, we hypothesized that the selection of the most appropriate heuristic for a given pattern is highly determined by its spatial configuration. A pattern elicits a particular response depending on contextual information, such as the strength of spatial relationships between the constitutive elements. This hypothesis is therefore linked to perceptual organization. Perceptual organization can be defined as the ability to impose structural organization on sensory data, so as to group sensory primitives arising from a common underlying cause (Carreira et al., 1998).

The neural substrates of perceptual grouping reside in the primary visual areas of the cortex. Simple cells in area V1 respond as linear spatio-temporal filters, and their receptive fields have been successfully modelled with Gabor filters (Daugman, 1988; Jones & Palmer, 1987; Lee, 1996), a set of Gaussian kernels modulated by a sinusoidal planewave. In our model the input patterns are processed with a set of Gabor filters to provide a computational account of the neural mechanisms involved in perceptual grouping (also see Carreira et al., 1998). Gabor filter processing provides the extraction of the salient features of the patterns, in particular their orientation. Humans solve the visually presented TSP essentially by applying spatial heuristics to the representation of the problem during the execution of the task. Therefore, we have assumed that the information deriving from the salient directional features extracted with the Gabor filters plays a crucial role in the selection of the heuristics. However, it is important to point out that the visual module is not specifically tied to the current model of the TSP. In fact, it is simply a general-purpose model for simulating low-level vision.

Visual processing provides information regarding the spatial-directional characteristics of the pattern formed by the points that constitute the TSP configuration. However, directional information provided by different visual orientation maps must be somehow compared to compute the principal axis of orientation. This has been achieved through a competitive selection process, implemented with a self-organizing, competitive-learning network. Competitive learning (Rumelhart & Zipser, 1985) sorts patterns sharing similar properties into the same category, and it can be viewed as a clustering technique. The network, presented with the input features detected by the visual-processing module, discovered three main categories of input images. The final version of the competitive network had therefore three output nodes, each encoding one specific image category. Each category was then associated to one of the three spatial heuristics.

Heuristics are selected by the competitive process based on bottom-up, salient perceptual information, but they must be turned into a signal that biases the execution of the pathway. The biasing signal in the model is provided by a saliency map, which is simply a gradient of activation influencing the spatial target map. Thus, the three heuristics (NN, DR, and DD) have been implemented in terms of different saliency maps. Activation of the spatial target map is driven by the retinal (input) image but it is modulated by the saliency map, so that a particular area is enhanced according to the selected heuristic: the upper region for DD heuristic, the left region for DR heuristic, and the space surrounding the last city visited for the NN heuristic. Spatial locations are represented on the spatial target map by Gaussian-shaped hills of activity (i.e., population coding; Pouget, Dayan, & Zemel, 2000, for a review). Lateral connections ensure that only one hill of activity, encoding the location of the next city to be visited, becomes fully active on the spatial map during processing. Note that the spatial map was not specifically designed for the TSP model, but it was taken from a previous computational model of visually guided movements (Di Ferdinando et al., 2005; also see Pouget & Snyder, 2000, for a similar approach); accordingly, no parameters of the spatial map were manipulated to implement the TSP model.

Every time a city is visited, the corresponding population code is subsequently suppressed in the spatial target map so that it will not be selected again during the sequential selection process. Moreover, inhibition spreads to the input map via a feedback connection to decrease the saliency of the visited city (i.e., its activation in the input map is reduced to 50%). As a consequence, Gabor filter reprocessing at the next time-step has the potential of triggering a different heuristic.

The sequential behaviour of the model arises from the competitive dynamics that are intrinsic in the spatial map, whereas the specific order in which cities are selected depends on biasing the competition through an activation gradient (i.e., the saliency map). The idea that sequentially ordered behaviour involves a stage of parallel activation of a set of responses has a long history and indeed was central to Lashley's (1951) influential arguments against associative chaining (see Houghton & Hartley, 1995, for discussion). Note that biased activation competition is central to the competitive queuing approach to serial order (Houghton, 1990), which has proven very effective to simulate sequentially ordered behaviour in both normal and pathological conditions (see Houghton & Hartley, 1995, for a review, and Botvinick & Plaut, 2004, 2006, for an alternative approach).

The capacity to change heuristic during the execution of the task—and thus to produce flexible behaviour—is guaranteed by the top-down controller (TDC), which simply has the ability to reset the activation of the three heuristic units in the competitive selection module. There are at least two competing hypotheses regarding the on-line control of heuristic choice. The first hypothesis is that participants adopt a constant replanning approach, in which a new heuristic is chosen at each step regardless of the past choices (e.g., Koenig, Furcy, & Bauer, 2002). The second hypothesis is that the switching is driven by a mismatch detection mechanism that triggers replanning only if the currently selected heuristic is making poor progress (e.g., Onaindía, Sapena, Sebastia, & Marzal, 2001). We tested the two accounts in different versions of the model to assess whether one of them would provide a better fit to the human data. Thus, in the constant replanning model the competitive selection module (i.e., the active heuristic) is reset after each step during the execution of the pathway. In contrast, the mismatch detection model performs a new step based on the heuristic used in the previous step, and the TDC resets the competitive selection module only when a mismatch between the current heuristic and the optimal heuristic (determined by reprocessing the visual input up to the level of the competitive selection module) is detected.

Implementation of the model

Visual processing of the TSP patterns

The input to the visual module consists of a 161 × 161 pixels image representing the pattern. Each city is represented by a black circle with a diameter of 8 pixels coded with ones, whereas empty space is coded with zeros. The image is then processed by a family of Gabor filters (see Appendix for mathematical details). Gabor filters are band-pass filters with tuneable centre frequency, orientation, and bandwidth, which can model the response of simple cells in the primary visual cortex (Lee, 1996). A set of eight Gabor filters tuned to different orientations was used to convolve the input image to obtain eight orientation maps. The use of a small number of filters was motivated by the nature of the task and by the need of simplifying the model (see Appendix for further discussion).

To calculate the strength of the directional features extracted by the different filters, we collected the highest value from each orientation map. This corresponds to a nonlinear MAX operation (see Riesenhuber & Poggio, 1999) over the units belonging to the same map, which effectively provides orientation information that is invariant of spatial position. Moreover, the structure of the visual module is consistent with the hypothesis of Field (1994) that oriented edge detectors constitute a sparse representation of the images. This means that for any image, only a few of the features are needed to represent that particular image, and that over an ensemble of images a particular feature will seldom be significantly active.

Competitive selection among heuristics

The output of the visual module is sent to a competitive network (Rumelhart & Zipser, 1985) that provides an unsupervised categorization of the input pattern based on its visuospatial (i.e., directional) properties. The input layer, composed of eight units, encoded the normalized highest value of the eight visual maps—that is, the strength of a particular orientation axis. The output layer encodes the categories discovered in the training data by the competitive learning algorithm (see Appendix for details). Each output unit of the competitive network sends inhibitory connections to all other output units and one excitatory connection onto itself. This implements a winner-takes-all mechanism ensuring that only one output unit becomes active for a given input. The output units of the competitive network are named "heuristic units" because each unit is later associated to one specific spatial heuristic (see below).

The training set consisted of 100 images with patterns of 6 to 10 cities; each pattern was generated using a pseudorandom procedure that produced a structure consistent with the TSP configurations used by Basso et al. (2001). The network discovered three categories of input images. Inserting more than three units in the competitive layer did not produce substantial changes in the results, as most of the patterns (95%) were still classified by three units. We associated each category (i.e., output unit) to one specific spatial heuristic. We observed that most of the participants executing the TSP task tended to select the direction right (DR) heuristic when the cities were principally distributed along the diagonal axis from top-left to bottom-right, while the direction down (DD) heuristic was often chosen when the cities were principally distributed along the opposite diagonal axis. In intermediate or ambiguous situations, subjects tended to use the nearest-neighbour (NN) heuristic. Notably, two output units of the competitive network after learning were mainly driven by the visual maps tuned to 45° and 135°, respectively. The 45° orientation corresponds to a top-left to bottom-right direction, whereas the 135° orientation corresponds to a top-right to bottom-left direction in the input image. Accordingly, we connected the two units to the DR and DD heuristics, respectively. The last output unit was not driven by any specific orientation map, suggesting that the patterns that it responds to are not characterized by a predominant orientation axis. Accordingly, it was connected to the NN heuristic, which is typically used by participants when confronted with TSP configurations whose spatial layouts lack a clear directional component.

Saliency maps

The three heuristics are implemented by different saliency maps that bias the execution of the pathway with a specific gradient of activity. Each saliency map has the same size as the spatial map (see below), and each unit in the spatial map is activated by the corresponding unit of the saliency map. The saliency map of the NN heuristic is implemented as a Gaussian-shaped hill of activity centred on the last visited city, whereas the saliency maps of the DR and DD heuristics consist of linear gradients that cover the entire visual space (Figure 3). The effect of a saliency map on the spatial map is to relatively enhance the activation of one city by reducing (in accordance with the specific gradient) the activation of the other cities on the map. The competition among units in the spatial map (see below) produces a single winning location that corresponds to the most salient city, which constitutes the next target.

All saliency maps consist in a gradient of activation with a value of 0.3 in the position of maximum enhancement. Thus, the DR saliency map has a value of 0.3 at the extreme left side that decreases linearly to zero at the extreme right side. The DD saliency map employs an equivalent gradient from the upper to the lower side of the spatial target map. Finally, the NN saliency map is represented by a broad Gaussian-shaped hill of activity centred on the last visited city, with a peak of 0.3 and a width of 15°.

Human participants in the open-ended TSP used by Basso and colleagues (2001) were instructed to start with the top-left point and finish with the bottom-right point. This constraint is implemented in the model through a small gradient of activation that provides a "default" bias to the spatial map. That is, the upper left corner is enhanced by 0.05, and the biasing activation decreases linearly to 0 at the bottom right corner. This small bias ensures that the tours performed by the model always start with the top-left point and finish with the bottom-right point but it is completely orthogonal to the (stronger) biases caused by the heuristic saliency maps.

Spatial target map

The spatial target map is composed of 21 × 21 units with lateral connections. Activation of the spatial target map is driven by the retinal (input) image but it is modulated by the saliency map (see Appendix for mathematical details). Spatial plans are represented at population code level (see Figure 4). Each city is represented by a Gaussian-shaped hill of activity, and the competition between units belonging to different populations is resolved over time in favour of one single population. The next city to be visited is therefore represented by the winning population. Its exact location is decoded through a simple vector method (Salinas & Abbott, 1995).

Note that goal locations are coded in retinal coordinates on the spatial map. This frame of reference, appropriate only for eye movements, was chosen for the sake of simplicity. A coordinate

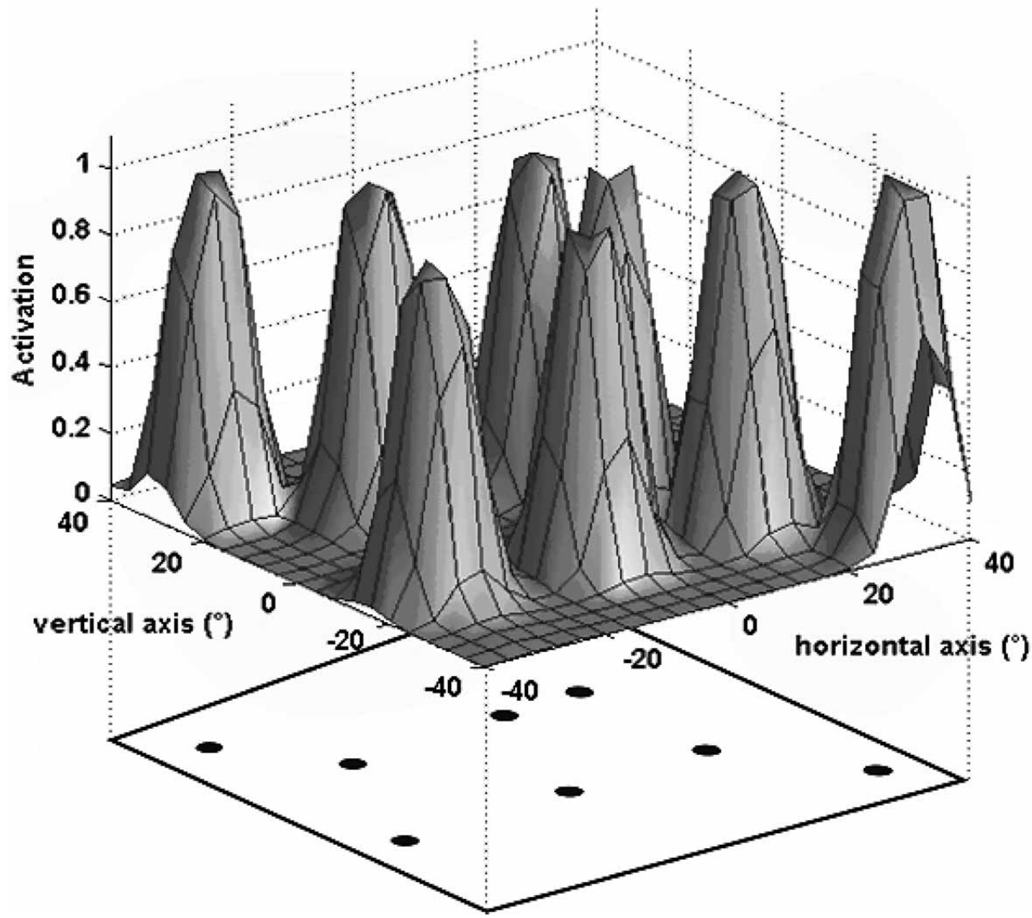

Figure 4. *The activation of the spatial map for one 8-point TSP pattern resulting from the retinal input only (i.e., a biasing signal from the saliency map is not present). Each location activates a population of units that is defined by a broad Gaussian-shaped hill. The different populations compete until there is only one winning population (while all the others are inhibited). This is the next location to be visited.*

transformation into a head-centred or hand-centred motor system would simply require the addition of one intermediate layer of "basis function" units that combine, in a multiplicative way, the retinal signal with the posture signals encoding the position of the eye and the hand (see Pouget & Snyder, 2000, for a review).

Top-down controller

The top-down controller (TDC) provides the required flexibility of the model to make it capable of a change of heuristic during the execution of the task. In the constant replanning version, the TDC resets the competitive selection module after each step. This means that a "new" heuristic must be selected before each step. In contrast, the mismatch detection version of the TDC triggers a reset only when the heuristic used to perform the last step is not optimal. More specifically, a change of heuristic is made possible only after the execution of one step in which the active heuristic did not match the optimal heuristic indicated by visual analysis. Thus, the mismatch detection TDC can switch the heuristic at least one step later than the constant replanning TDC.

Parameters of the model

Table 1 lists all parameters of the model. As previously mentioned, most of the values are identical to those used in the studies that describe a specific component of the model. The parameters of the visual model were not manipulated and simply reflect the minimal set of visual filters (cf. Riesenhuber & Poggio, 1999) at a medium-scale spatial resolution. The spatial target map was entirely taken from the work of Di Ferdinando et al. (2005) without any parameter change. The lateral connections in the competitive network were simply set to values that resulted in efficient winner-takes-all behaviour, and the learning rate was set to 0.1 without any manipulation. Thus, only one parameter is specific to the current work, which is the peak of activation in the saliency maps. However, this parameter was not systematically manipulated but it was simply set to a value that ensured a proper biasing of the competitive selection between targets in the spatial output map. In summary, it is important to stress that no parameter was manipulated for data-fitting purposes.

Simulations of normal performance

The model operates in a sequential manner, and it performs the task by choosing one city at every step. For each TSP configuration presented as input to the model, the visual image is analysed with Gabor filters to detect the most influential spatial features, and then the competitive selection module selects the heuristic that is most appropriate for the input pattern. The winning heuristic is implemented in terms of a saliency map, whose activation influences the spatial map determining the city to be visited. The spatial map represents all the locations to be visited through population codes; competition at the spatial level results in selection of the most activated population code, which corresponds to one particular goal location (the forthcoming city that will be visited), and in the inhibition of the other populations codes (all other locations). At the end of each step, the units in the spatial map corresponding to the selected (visited) city are inhibited, and the activation of the same city in the visual pattern is reduced via the inhibitory feedback loop (see Figure 3). This allows a possible change of heuristic: Indeed, the visual input is processed again, and a different heuristic might emerge from the competitive selection. This process takes place at every single step; therefore, heuristics may be changed several times during the execution of a single path.

However, the switching process requires the intervention of the TDC, which has the role of resetting the currently selected schema (i.e., the specific heuristic). A top-down influence

Table 1. *List of the parameters used in the model*

	Parameter	*Value*	*Taken from*
Visual module	x	0	Riesenhuber & Poggio, 1999
	y	0	Riesenhuber & Poggio, 1999
	ω_o (radial frequency)	.57 radians per unit length	
	θ (wavelet orientation)	0, $1/4\pi$, $1/2\pi$ and $3/4\pi$ radians	Riesenhuber & Poggio, 1999
	k	$K \approx \pi$	Lee, 1999
Competitive selection module	η (learning rate)	0.1	
Spatial module	σ (width of the Gaussian)	5°	Di Ferdinando et al., 2005
	A_E	10	Di Ferdinando et al., 2005
	σ_E	15	Di Ferdinando et al., 2005
	A_I	9	Di Ferdinando et al., 2005
	σ_I	105	Di Ferdinando et al., 2005
	Saliency maps	Peak activation of 0.3	

Note: See text for explanation.

becomes clearly visible in the most difficult tasks (i.e., for a large number of cities) where it provides flexibility to the behaviour. If the previously selected heuristic is not appropriate on the basis of the spatial analysis of the remaining unvisited locations, the TDC allows an on-line change to the plan. As previously mentioned, we evaluated the performance of the model using the two different versions of the TDC (constant replanning vs. mismatch detection). Moreover, to investigate whether the incremental aspect of planning is a necessary component to adequately describe human performance, we evaluated a version of the model in which one single heuristic was used for the execution of the entire path. In summary, we compared five different versions of the model: (a) constant replanning model, (b) mismatch detection model, (c) fixed DD model, (d) fixed DR model, and (e) fixed NN model.

These five models were tested on eight different TSP patterns. These were the most frequently tested maps across several experiments performed by Basso et al. (2001) and Basso (2005), so that the performance of each model could be compared with the behavioural data. We collected the tours executed by the model in its different versions for a comparison with those executed by the healthy adults (the number of participants used for this analysis was 140). We chose to inspect the tours at a global level, instead of analysing the single movements during the execution of the tour. Note that the comparison between model and human data at the level of single step is not a stringent one because the likelihood of finding the same single step is much higher than the likelihood of finding the same whole pathway. For each pattern, we compared the tour produced by the model with the tours performed by human participants and ranked the model's tour according to the frequencies observed in the human data.

Results

Table 2 reports the ranks for each tour and model version. A rank of 1 indicates that the tour performed by the model corresponds to the most frequently observed tour across human participants. Some pathways were not classifiable (N.C.) because they did not match any of the tours observed across our sample of human participants. In the case of Pattern 8, none of the models was able to provide a classifiable solution. It is worth noting, however, that the human solutions to Pattern 8 (constituted by 11 locations) was so widely variable that their ranking would be unreliable. The number of possible solutions increases exponentially with the number of locations to be visited; this turns into a greater variability of performance because there is also an increased number of close-to-optimal solutions. This contention is supported by a strong positive correlation ($r = .83$, $p < .005$) between the number of points in a map and the number of different solutions provided by human participants.

Inspection of Table 2 reveals that most of the pathways chosen by the constant replanning model correspond to the most frequently observed tours in the experiments of Basso and colleagues (2001; Basso, 2005). For half of the patterns, the tour executed by the constant replanning model matched the most frequent tour observed in healthy participants. Overall, six out of eight pathways were the first or second most frequent tour in the human data. This is a valuable result, considering that the open-ended TSP is a (non polynomial) NP-complete problem, which has an exponential increase of possible solutions in relation to the number of points. A nonparametric Friedman analysis of variance (ANOVA) on the rank of the tours produced by the five models revealed a significant effect of model version, $\chi^2(4) = 13.143$; $p < .05$. The mean ranks were 1.94, 2.94, 2.88, 4.31, and 2.94 for constant replanning, mismatch detection, fixed NN, fixed DR, and fixed DD, respectively. Pairwise comparisons showed that performance of the DR model was worse than any other model, but no other comparison reached the significance level. However, it is worth noting that the constant replanning model showed the best mean rank. The mean rank of the mismatch detection model was identical to that of the fixed DD model.

A qualitative analysis of the tours performed by the different models confirms that the constant replanning model offers the best fit to the human

Table 2. *Ranks and frequencies of the tours chosen by all versions of the model for each pattern in a comparison with the human data of Basso et al., 2001; Basso, 2005*

Map		Ranking and Percentage				
		Constant Replanning	Mismatch Detection	NN	DR	DD
	1	2 (16.8)	2 (16.8)	2 (16.8)	3 (5.46)	1 (68)
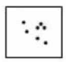	2	1 (38.66)	1 (38.66)	3 (14.71)	7 (2.52)	1 (38.66)
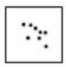	3	1 (50)	1 (50)	1 (50)	3 (10)	1 (50)
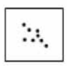	4	1 (25.6)	2 (18.9)	2 (18.9)	7 (2.9)	2 (18.9)
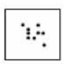	5	2 (23.1)	6 (2.5)	6 (2.5)	3 (15.5)	N.C.
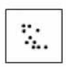	6	1(11.3)	1(11.3)	1(11.3)	10 (2.1)	7 (5.5)
	7	6 (11.3)	N.C.	12 (1.3)	N.C.	18 (0.4)
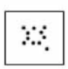	8	N.C.	N.C.	N.C.	N.C.	N.C.

Note: Each pattern depicted in the first column. A rank of 1 indicates that the pathway chosen by the model corresponds to that most frequently observed across human participants. N.C. (not classifiable) indicates that the tour performed by the model has never been observed across the sample of human participants. NN = nearest neighbour. DR = direction right. DD = direction down.

performance. Specifically, both the mismatch detection and the fixed NN models produced two crossed tours, whereas the constant replanning model did not produce any crossed tour (note that the fixed DD and fixed DR cannot, by definition, produce any crossing). The occurrence of crossings in the tours produced by human participants is extremely rare. Indeed, van Rooij, Stege, and Schactman (2003) hypothesized that humans try to avoid crossed lines when solving the TSP because they are sensitive to the fact that tours with crossed lines are nonoptimal (or, alternatively, that optimal tours have no crossings). Although this position has been criticized by MacGregor, Ormerod, and Chronicle (2004), it is generally agreed that optimal solutions have no crossings.

In summary, the constant replanning model is superior to the simpler models based on single heuristics, as well as to the model that allows for heuristic changes through a mismatch detection mechanism. Thus, having established that the constant replanning model provides the best fit to the data on healthy human participants, we investigated the effect of a lesion to the TDC to assess its ability to account for the impaired performance shown by fTBI patients.

Simulations of impaired performance

Simulation of the impaired performance of frontal patients was obtained by lesioning the TDC in the constant replanning model. The TDC was impaired by lowering its capacity to reset the competitive selection module. Lesions of different degrees of severity were simulated by progressively increasing the inefficiency of the top-down controller from 10% to 70%. A set of 10 frontally damaged networks was obtained by setting the lesion severity to 10%, 15%, 20%, 25%, 30%, 35%, 40%, 50%, 60%, and 70%.

The tours performed by the lesioned model were compared with those observed in fTBI patients. Tour rankings were computed as for the simulations of unimpaired performance. Moreover, we carried out an analysis of the degree of flexibility exhibited by the model by looking at the selection of heuristics during the execution of the task. In previous studies the classification of the heuristic used to move from one city to the next at each step in the pathway was obtained considering the distance between that city and all other unvisited cities (see Basso et al., 2001). If the next selected city was the closest on the horizontal axis, the move was considered as the result of a DR heuristic. Alternatively, if the next city was the closest on the vertical axis, the move was considered as the result of a DD heuristic. Finally, if the next city was the closest in terms of absolute distance (i.e., regardless of the direction), the move was considered representative of a NN heuristic. In the model, the choice of heuristic was directly assessed by recording the winning unit in the competitive module at every step. The presence of multiple heuristics in solving a given TSP configuration was taken as an index of flexible behaviour. Specifically, a flexible strategy was defined as a problem solution in which the participants (or the model) operated at least one heuristic switch; otherwise, the strategy was classified as rigid. Therefore, we assessed the flexibility of the lesioned model in comparison to the unimpaired model and its effect on tour optimization. Performance of 10 different "normal participants" was obtained by introducing some variability in the competitive process that leads to heuristic selection for both the simulations. Gaussian noise (mean 0, variance .05) was added at each processing step to the heuristic nodes. Performance on each map was therefore collected for each of the 10 different runs of the model.

Results

Each tour performed by the model after TDC lesions was ranked according to the patients' data, and the modal rank was calculated for each pattern. The tours of the impaired model were generally consistent with the performance of the frontal patients (see Table 3).

A comparison of the type of strategy (flexible vs. rigid) employed by the normal model and the lesioned model revealed a significant difference, $\chi^2(1) = 13.97$, $p < .001$. Rigid strategies were more frequent than flexible strategies in the lesioned model, whereas the unimpaired model showed the opposite trend. Overall, this pattern mirrors the results obtained by Basso et al. (2001) in their study of healthy participants and frontal patients (see Figure 5).

Figure 6 presents a comparison of two representative pathways performed by the constant replanning model and the lesioned model (30%

Table 3. *Ranks and frequencies of the tours performed by the lesioned model, for each pattern (depicted in the first column) in a comparison with the human data of frontally damaged patients (Basso et al., 2001)*

MAP	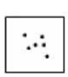	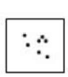	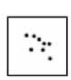	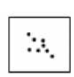	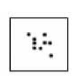	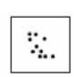	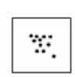	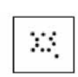
RANK	2	1	2	3	3.5	3	1	N.C.

Note : Each pattern depicted in column head.

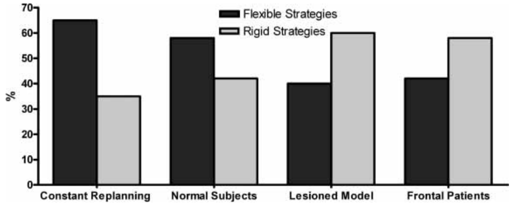

Figure 5. *Overall strategy chosen by the normal models and by the lesioned models, compared to the strategies of healthy participants and frontal patients. A strategy was classified as flexible if the heuristic was changed at least once during the execution of a given pathway.*

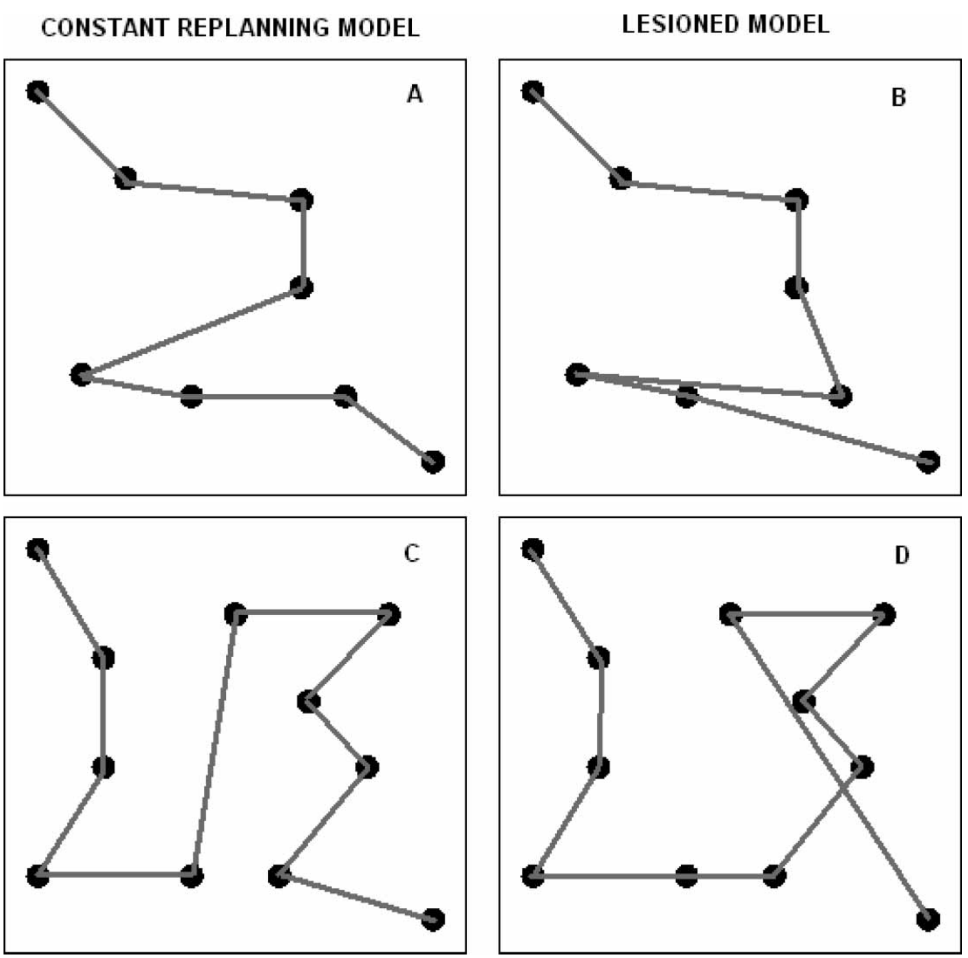

Figure 6. *Tours performed by the model on Maps 4 and 8. The left panels (A and C) show the performance of the normal model, whereas the right panels (B and D) show the pathways executed by a 30% lesioned model. Note that the tour depicted in D contains a crossing, which is a classical sign of nonoptimization and is rarely observed in neurologically intact participants.*

lesion). As can be noted, the pathways were markedly different because of the lack of heuristic changes. The tour executed by the lesioned model contains a crossing, which is a clear sign of nonoptimal performance (as discussed in the previous section).

Another fundamental difference between normal and impaired model is revealed by their different levels of optimization. As suggested by Graham, Joshi, and Pizlo (2000), an appropriate measure of the level of optimization is the ratio of the tour length (RTL), which is the ratio between length of the tour performed by the subject and the length of the shortest tour. In the present study (as in Basso et al., 2006), the RTL index for a given pathway X was computed as the ratio of excess tour length to optimal tour length:

$$\mathrm{RTL}(X) = [\text{tour length}(X) - \text{optimal tour length}(X)]/\text{optimal tour length}(X)$$

Mean RTLs for normal versus impaired model were significantly different (0.051 vs. 0.117, respectively; $t = -26.84$, $p < .001$), showing that TDC damage results in a lower level of optimization. These result mirror those of Basso et al. (2001), who found that the optimization level of frontal patients was significantly inferior to that of healthy controls (see Figure 7).

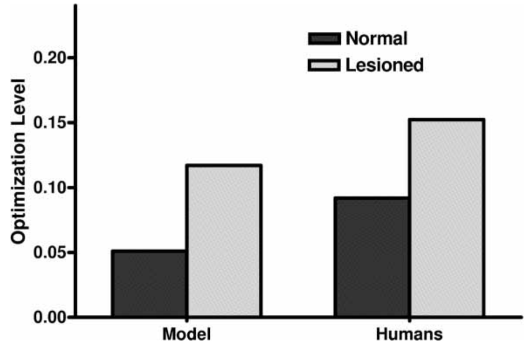

Figure 7. *Mean ratio of the tour length (RTL) for model and human participants, indexing the optimization level achieved in the TSP task. From left to right: normal model (dark grey), lesioned model (light grey), neurologically intact participants (dark grey), and frontal patients (light grey). Human data from Basso et al., 2001.*

GENERAL DISCUSSION

The TSP is a famous problem-solving task which gained notoriety about two centuries ago among mathematicians and physicists as one of the most fascinating optimization problems (for a review see Schrijver, 2005). Nevertheless, in the latest years there has been a growing interest for the TSP among cognitive scientists. Mathematicians and computer scientists have developed a large number of algorithms for solving the TSP that give an approximation of the optimal solution in a reasonable amount of time. Our model differs fundamentally from these studies in its purpose, because we have focused our attention on the simulation of the human cognitive processes involved in the solution of the TSP rather than on the optimization aspects of the problem. Indeed, humans facing the TSP problem carry out the task in a sequential manner, showing the ability to change heuristic during the pathway when needed to optimize their performance (Basso et al., 2001).

The simulations presented in this paper closely mirrored human performance in both normal and pathological conditions. This suggests that the model is able to capture the basic cognitive processes involved in the human solution of the TSP. The core of the model's ability to perform visuospatial planning resides in the bottom-up selection of a visuospatial heuristic in the competitive selection module. The spatial analysis performed by Gabor filters is the source of bottom-up influences in the model. The competitive network receives input from the visual module, and it provides an unsupervised categorization of the input pattern based on the spatial-directional characteristics of the pattern formed by the points that constitute the TSP configuration. Categorization into three classes emerged spontaneously during learning (through an unsupervised Hebbian learning algorithm) as a result of the exposure to a sample of TSP patterns.

Thus, the directional features of the input pattern are the main determinants of the heuristics selected in the particular path used to solve the TSP. The remarkable similarity between model and humans in choosing a movement path suggests that perceptual grouping and sensitivity to the spatial-directional characteristics of the visual pattern is a fundamental component of visuospatial planning in the TSP.

From a behavioural point of view, the key to the model's ability to generate plausible pathways resides in two main features: the selection of the most appropriate heuristic given the contextual information and the incremental monitoring process, which allows a change of heuristic when the ongoing one fails to fit to the sensorial information. Indeed, the most intriguing characteristic of the model regards its capacity to switch between heuristics. This is a fundamental characteristic that gives psychological plausibility to the model. Human participants execute the TSP in an iterative manner; the incremental process is less resource demanding than global planning because subjects do not need to generate a comprehensive plan resolving the entire situation but only the next appropriate action.

The hypothesis that the performance of neurologically intact participants is crucially dependent on incremental planning and on the interaction of bottom-up and top-down processes is highlighted by the comparison between five versions of the model that differed only for the strategic component. First, the use of a single, fixed heuristic proved to be inadequate for the solution of the most complex problems. Fixed-heuristic versions of the model performed plausible tours when confronted with maps containing a small number of points, but their performance broke down as the number of points increased. Specifically, we observed pathways that did not match any of the tours produced by human participants. This result confirms that a flexible use of heuristics is a fundamental aspect of human performance in the TSP (cf. Basso et al., 2001).

Flexible behaviour, however, requires a monitoring system that allows for heuristic switches. We contrasted two alternative hypotheses regarding the operation of the top-down controller: constant replanning versus mismatch detection. The first hypothesis is that a new heuristic is chosen at each step regardless of the past choices (e.g., Koenig et al., 2002), whereas the second suggests that the switching can take place only if the currently selected heuristic is making poor progress (e.g., Onaindía et al., 2001). The two different implementations of the TDC produced a markedly different performance. First, the pathways executed by the constant replanning model provided a better match to the tours that were most frequently observed across a large number of human participants. Second, and more important, the constant replanning model did not produce any crossed tour, whereas the mismatch detection model executed two crossed tours. The latter indicate a nonoptimal performance and are rarely observed in human performance (MacGregor, Chronicle, & Ormerod, 2004; van Rooij et al., 2003). This suggests that a continuous monitoring process, which allows for on-line changes of heuristic whenever the current one is not suitable, is a more viable model of the incremental planning ability displayed by neurologically intact participants.

The importance of flexible, incremental planning is also supported by the empirical data on patients with frontal traumatic brain injury (Basso et al., 2001) as well as normal subjects under rTMS over the frontal lobe (Basso et al., 2006; see Figure 2a). Both the lesion and the reversible neurodisruption of the frontal cortex caused a conspicuous decrease of flexible strategies that incorporate heuristic switches. The same pattern was observed in the simulations when the TDC of the constant replanning model was lesioned in a way that decreased its ability to reset the competitive selection module. Note that the frontal patients still produced acceptable solutions of the TSP in the Basso et al. (2001) study. The damaged model performed in the same way because of the preserved bottom-up mechanism: The simplest TSP problems often do not require a change of heuristic, and thus the performance of the lesioned model is indistinguishable from that of the normal model.

However, in the most complex patterns the intact model efficiently switched heuristic when a change was appropriate, whereas the damaged model often perseverated by keeping active the heuristic selected at the beginning of the pattern. Thus, TDC damage caused a loss of flexibility and adaptivity in the behaviour that turned into a poorer level of optimization.

It is worth noting that the remarkable flexibility of human cognitive control was simulated in the model solely by the interaction between the competitive selection module and the TDC. This reflects the functional role of the PFC hypothesized by J. D. Cohen, Braver, and O'Reilly (1996), who suggest that the PFC maintains the relevant features in an activation-based working memory, providing a top-down support (or biasing) of the corresponding perceptual processing and action selection pathways (also see O'Reilly, Noelle, Braver, & Cohen, 2002).

However, there is an alternative explanation for why heuristic switching might fail after frontal lobe damage. If monitoring is seen as a parallel process that could interrupt behaviour before the application of a heuristic, a monitoring failure would prevent the system to invoke the TDC.[1] This would cause the same loss of flexibility that we attributed to (and modelled as) impaired inhibitory function of the TDC. Thus, further empirical work is necessary to distinguish between these two alternative accounts of the rigid behaviour exhibited by fTBI patients (or healthy participants under rTMS) in the TSP. Notably the monitoring failure account would predict a loss of flexibility in neurologically intact participants when the attentional demands are increased by simultaneously performing a secondary task.

Relation to other models

The most important previous computational model of TSP from a psychological point of view was developed by MacGregor et al. (2000). Their model focuses on the human solution of the closed version of the TSP (see Introduction): It performs the task with a sequential procedure, and it is basically designed to conform in a general way to a convex hull approach. The model of MacGregor and colleagues is entirely driven by conditional rules implemented at each step. One example of its operations is provided by the following list of steps: "Apply the insertion criterion to identify which unconnected interior point is closest to the current arc—apply the insertion criterion to check whether the closest node is closer to any other arc—if not, proceed to Step 5—if it is, move to the end node of the current arc." Although this model represented a great development with respect to the previous conventional attempts to model human performance on TSP, some aspects still seem to require an improvement. Even if the results show a good fit between the model and the human solution, the model of MacGregor is unlikely to produce a particularly good fit to human solutions to highly patterned TSPs. This is because the model does not incorporate a mechanism that is sensitive to factors such as proximity of interior points and regularity of their arrangement. Moreover, recent evidence suggests that the heuristics used to solve the closed version of TSP cannot explain the human performance in the open-ended TSP (Chonicle, MacGregor, Ormerod, & Burr, 2006). Thus, the model of MacGregor et al. (2000), at least in its present form, could not be used to account for the human data on the Maps test of Basso et al. (2001).

It is important to note that a fundamental difference between our model and the model of MacGregor et al. (2000) resides in the nature of the computational mechanisms leading to the solution. In our model, there are no explicit rules that guide the performance, and the perceptual mechanisms simulated with Gabor filters allowed us to account for the perceptual components of the human solution of TSP. To the best of our

[1] We thank Rick Cooper for suggesting this view of monitoring as well as the account of rigid behaviour based on a monitoring failure.

knowledge, there is no other model that has successfully simulated human performance in TSP without using explicit rules.

Our model shares several conceptual properties with the attention to action (ATA) model of Norman and Shallice (1986) and the computational model of action planning developed by Cooper and Shallice (2000). The competitive selection module operates in a way that is similar to the contention-scheduling mechanism, whereas the TDC could be regarded as a sort of supervisor attentional system (SAS). In our model, the role of the TDC is to inhibit the previously selected heuristic to promote a flexible behaviour. Notably, a lesion to the TDC resulted in performance characterized by a rigid strategy that lacked heuristic changes, with a pattern that mirrored the behaviour exhibited by frontal patients.

Limitations of the model and future directions

Overall, the model provides a good match to both normal and impaired human performance in the TSP task. One possible criticism, however, is that the complexity of the model and the number of parameters outstrip the complexity of the data to be explained. This is far from being the truth. First of all, the complexity of the model is the result of the assembly of components taken from other computational models that are unrelated to the TSP but are necessary to build a comprehensive model of the task. Indeed, the complex nature of the TSP and the number of different cognitive processes involved called for a nested modelling approach (see Perry et al., 2007). Therefore, two main components of our model, the visual module and the spatial module, were simply taken from state-of-the-art computational models of vision and action (Di Ferdinando et al., 2005; Lee, 1996; Pouget & Snyder, 2000). As a result of this, most of the parameters in the model are not free (see Table 1), and they have an effect only at the module level (i.e., they are not determinant for an optimal solution of the TSP). In fact, the combination of the visual and spatial components simply forms a model of visually guided movements towards visually salient stimuli. Two other components, the competitive selection module and the TDC, are required to produce a solution of the TSP and form the core of the model's ability to perform incremental visuospatial planning.

A second possible criticism regards the complexity of the data to be explained. At a first glance, an analysis of the tours at a global level would seem less accurate than an analysis of the single movements during the execution of the tour. What should be kept in mind, however, is that the overall solution is much more meaningful than the single movements. Indeed, looking at the match between model and human data at the level of single step is not very fruitful, because the probability of finding a matching movement is much higher than the probability of finding the same pathway in its totality. In fact, for a tour with rank "1", the overall pathway of the human participants and that of the model is exactly the same, indicating a maximum concordance between the two solutions. Thus, this approach provides a rigorous and fine-grained analysis of the performance.

Nonetheless, a limitation of the current model is that it cannot capture the variability of human solutions to a single TSP pattern. Modelling individual differences was clearly beyond the scope of the current work, but this issue could now be tackled by combining behavioural and computational investigations. Simulations suggest that the observed variability of human solutions cannot be accounted for by simply adding noise during processing. Thus, individual differences might be captured in the model only by considering different "cognitive styles" (e.g., Witkin & Goodenough, 1981), such as predispositions towards using a particular heuristic (Bisiacchi, Basso, & Cotelli, 1999).

The last possible criticism to the current model concerns its limitation to the open-ended TSP task. In this regard, we note that the extension to the closed-TSP would require minimal modifications of the model. In particular, the two versions of TSP seem to involve different types of heuristics (Chronicle et al., 2006). Moreover, although the model was specifically designed to provide a solution of the TSP, it could be extended to other tasks that require sequential visuo-motor scanning. For example, the Trail-Making Test B (TMT-B: Reitan, 1958) requires alternately connecting digits (1–13) and letters (A–L) that are randomly located on a sheet of paper. The TMT-B is frequently used in neuropsychology to assess the executive function of mental set shifting. A simulation of the TMT-B test would require only few modifications of our model. Finally, the competitive selection and TDC components of our model could be used to simulate other nonspatial tasks that involve cognitive control. One advantage of our nested incremental modelling approach (see Perry et al., 2007) is that these components can be easily untied from the other parts of the model and reused in a different context. Note that only the competitive selection module and the TDC would be essential to account for the rapid switching in dynamic categorization tasks, as in the Wisconsin Card Sorting Test (WCST; Grant & Berg, 1948). However, extension to the WCST would clearly require substantial modifications that range from the nature of input and output representations to the structure of the learning phase (Rougier, Noelle, Braver, Cohen, & O'Reilly, 2005).

CONCLUSIONS

The present work is an effort to simulate the nature of the computational mechanisms underlying the human performance on visually presented TSP. Consistent with the connectionist approach to human cognition, our model dispenses with the use of explicit rules to guide behaviour in a complex problem-solving task. The simulations highlighted the fundamental role of perceptual grouping and sensitivity to the spatial-directional characteristics of the visual pattern in visuospatial planning while performing the TSP. Moreover, the model allowed us to assess the role of incremental planning and to test different hypotheses regarding the on-line monitoring of performance. The behaviour of the model seems to capture the fundamental aspects of human skilled performance and to mirror the impairment and behavioural rigidity typical of frontal lobe patients after a simulated lesion. We believe that the model provides a useful platform for designing new empirical studies that aim at a more fine-grained analysis of human performance in the TSP task because it can be used to make predictions regarding both healthy subjects and clinical populations.

Manuscript received 3 January 2007
Revised manuscript received 26 July 2007
Revised manuscript accepted 1 August 2007
First published online 6 October 2007

REFERENCES

Barr, A., & Feigenbaum, E. A. (1981). *The handbook of artificial intelligence* (Vol. 1). Stanford Los Altos, CA: W. Kaufmann.

Basso, D. (2005). *Involvement of the prefrontal cortex in visuo-spatial planning.* PhD thesis, Department of Psychology, University of Rome "La Sapienza". Available from: http://padis.uniroma1.it/getfile.py?recid=334

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., Ferreri, F., Bisiacchi, P. S., Olivetti-Belardinelli, M., et al. (2006). The role of prefrontal cortex on visuo-spatial planning: A repetitive-TMS study. *Experimental Brain Research*, *171*, 411–415.

Bisiacchi, P. S., Basso, D., & Cotelli, M. (1999). Cognitive strategies for sequential visuo-spatial decisions. In M. Zanforlin & L. Tommasi (Eds.), *Research in perception. Proceedings of the meeting in honour of F. Metelli* (pp. 179–181). Padova, Italy: Logos Press.

Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: A recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, *111*, 395–429.

Botvinick, M., & Plaut, D. C. (2006). Short-term memory for serial order: A recurrent neural network model. *Psychological Review*, *113*, 201–233.

Brown, T. H., & Chattarji, S. (1995). Hebbian synaptic plasticity. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 454–459). Cambridge, MA: The MIT Press.

Bryant, D. J., Tversky, B., & Franklin, N. (1992). Internal and external spatial frameworks for representing described scenes. *Journal of Memory and Language*, *31*, 74–98.

Cadwallader, M. (1975). A behavioral model of consumer spatial decision making. *Economic Geography*, *51*, 339–349.

Carlin, D., Bonerba, J., Phipps, M., Alexander, G. E., Shapiro, M. B., & Grafman, J. (2000). Planning impairments in frontal lobe dementia and frontal lobe lesion patients. *Neuropsychologia*, *38*, 655–665.

Carreira, M. J., Orwell, J., Turnes, R., Boyce, J. F., Cabello, D., & Haddon, J. F. (1998). Perceptual grouping from Gabor filter responses. In P. H. Lewis & M. S. Nixon (Eds.), *Proceedings of the Ninth British Machine Vision Conference* (pp. 336–345).

Chronicle, E. P., MacGregor, J. N., Ormerod, T. C., & Burr, A. (2006). It looks easy! Heuristics for combinatorial optimization problems. *Quarterly Journal of Experimental Psychology*, *59*, 783–800.

Cohen, G. (1988). *Memory in the real world*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: Recent developments and current challenges. *Philosophical Transactions of the Royal Society of London. Series B*, *351*, 1515–1527.

Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, *17*, 297–338.

Daugman, J. (1988). Complete discrete 2D Gabor transform by neural networks for image analysis and compression. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, *36*, 169–1179.

Di Ferdinando, A., Casarotti, M., Vallar, G., & Zorzi, M. (2005). Hemispheric asymmetries in the neglect syndrome: A computational study. In A. Cangelosi, G. Bugmann, & R. Borisyuk (Eds.), *Modelling language, cognition and action* (pp. 249–258). Singapore: World Scientific.

Duncan, J. (1986). Disorganization of behaviour after frontal lobe damage. *Cognitive Neuropsychology*, *3*, 271–290.

Field, D. J. (1994). What is the goal of sensory coding? *Neural Computation*, *4*, 559–601.

Fincham, J. M., Carter, C. S., van Veen, V., Stenger, V. A., & Anderson, J. R. (2002). Neural mechanisms of planning: A computational analysis using event-related fMRI. *Proceedings of the National Academy of Sciences*, *99*, 3346–3351.

Gärling, T. (1989). The role of cognitive maps in spatial decisions. *Journal of Environmental Psychology*, *9*, 269–278.

Gärling, T. (1994). Processing of time constraints on sequence decisions in a planning task. *European Journal of Cognitive Psychology*, *6*, 399–416.

Goel, V., & Grafman, J. (1995). Are the frontal lobes implicated in "planning" functions? Interpreting data from the Tower of Hanoi. *Neuropsychologia*, *33*, 623–642.

Goel, V., & Grafman, J. (2000). Role of the right prefrontal cortex in ill-structured planning. *Cognitive Neuropsychology*, *17*, 415–436.

Grafman, J. (1989). Plans, actions and mental sets: Managerial knowledge units in the frontal lobes. In E. Perecman (Ed.), *Integrating theory and practice in clinical neuropsychology* (pp. 93–137). Hillsdale, NJ: Lawrence Erlbaum Associates.

Grafman, J. (1995). Similarities and distinctions among current models of prefrontal cortical functions. *Annals of the New York Academy of Sciences*, *769*, 337–368.

Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The Traveling Salesman Problem: A hierarchical model. *Memory & Cognition*, *28*, 1191–1204.

Grant, D. A., & Berg, E. A. (1948). A behavioural analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology*, *38*, 404–411.

Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science*, *3*, 275–310.

Hirtle, S., & Gärling, T. (1992). Heuristic rules for sequential spatial decisions. *Geoforum*, *23*, 227–238.

Hirtle, S. C., & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory and Cognition*, *3*, 208–217.

Houghton, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation* (pp. 287–319). London: Academic Press.

Houghton, G., & Hartley, T. (1995). Parallel models of serial behaviour: Lashley revisited. *Psyche*, *2*(25). Retrieved September 24, 2006, from: http://psyche.cs.monash.edu.au/v2/psyche-2–25-houghton.html

Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, *58*, 1233–1258.

Kafer, K. L., & Hunter, M. (1997). On testing the validity of planning/problem-solving tasks in a normal population. *Journal of the International Neuropsychological Society*, *3*, 108–119.

Kaller, C. P., Unterrainer, J. M., Rahm, B., & Halsband, U. (2004). The impact of problem structure on planning: Insights from the Tower of London task. *Cognitive Brain Research*, *20*, 462–472.

Koenig, S., Furcy, D., & Bauer, C. (2002). Heuristic search-based replanning. In M. Ghallab, J. Hertzberg, & P. Traverso (Eds.), *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems* (pp. 294–301). Menlo Park, CA: AAAI Press.

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior*. New York: Wiley.

Lawler, E. L., Lenstra, J. K., Rinnooy Kan, A. H. G., & Shmoys, D. B. (1985). *The traveling salesman problem: A guided tour of combinatorial optimization*. New York: John Wiley and Sons.

Lee, T. S. (1996). Image representation using 2D Gabor wavelets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *18*, 959–971.

Lezak, M. (1995). *Neuropsychological assessment* (3rd ed.). New York: Oxford University Press.

Luria, A. R. (1980). *Higher cortical functions in man.* New York: Basic Books.

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

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

MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2000). A model of human performance on the traveling salesperson problem. *Memory & Cognition*, 7, 1183–1190.

McNamara, T. (1992). Spatial representation. *Geoforum*, *2*, 139–150.

Murakoshi, S., & Kawai, M. (2000). Use of knowledge and heuristics for wayfinding in an artificial environment. *Environment and Behaviour*, *32*, 756–774.

Newman, S. D., Carpenter, P. A., Varma, S., & Just, M. A. (2003). Frontal and parietal participation in problem solving in the Tower of London: fMRI and computational modeling of planning and high-level perception. *Neuropsychologia*, *41*, 1668–1682.

Norman, D., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. Davidson, G. Schwartz, & D. Shapiro (Eds.), *Consciousness and self regulation: Advances in research and theory* (Vol. 4, pp. 1–18). New York: Plenum.

Onaindía, E., Sapena, O., Sebastia, L., & Marzal, E. (2001). SimPlanner: An execution-monitoring system for replanning in dynamic worlds. In P. Brazdil & A. Jorge (Eds.), *Proceedings of progress in artificial intelligence knowledge extraction, multi-agent systems, logic programming, and constraint solving* (pp. 393–400). Berlin, Germany: Springer-Verlag.

O'Reilly, R. C., Noelle, D., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex and dynamic categorization tasks. Representational organization and neuromodulatory control. *Cerebral Cortex*, *12*, 246–257.

Owen, A. M., Downes, J. J., Sahakian, B. J., Polkey, C. E., & Robbins, T. W. (1990). Planning and spatial working memory following frontal lobe lesions in man. *Neuropsychologia*, *28*, 1021–1034.

Perry, C., Ziegler, J. C., & Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: The CDP+ model of reading aloud. *Psychological Review*, *114*, 273–315.

Phillips, L. H, Wynn, V. E., McPherson, S., & Gilhooly, K. J. (2001). Mental planning and the Tower of London task. *Quarterly Journal of Experimental Psychology*, *54A*, 579–598.

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.

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*, 125–132.

Pouget A., & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience*, *3*(Suppl.), 1192–1198.

Reitan, R. (1958). Validity of the Trail-Making Test as an indication of organic brain damage. *Perceptual Motor Skills*, *8*, 271–276.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*, 1019–1025.

Roberts, R. J., & Pennington, B. F. (1996). An interactive framework for examining prefrontal cognitive processes. *Developmental Neuropsychology*, *12*, 105–126.

Rougier, N. P., Noelle, D., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (2005). Prefrontal cortex and the flexibility of cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences*, *102*, 7338–7343.

Rumelhart, D. E., & Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science*, *9*, 75–112.

Salinas, E., & Abbott, L. F. (1995). Transfer of coded information from sensory to motor networks. *Journal of Neuroscience*, *15*, 6461–6474.

Schrijver, A. (2005). On the history of combinatorial optimization (till 1960). In K. Aardal, G. L. Nemhauser, & R. Weismantel (Eds.), *Handbook of discrete optimization* (pp. 1–68). Amsterdam: Elsevier.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London, Part B*, *298*, 199–209.

Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge, UK: Cambridge University Press.

Simon, H. A. (1973). The structure of ill-structured problems. *Artificial Intelligence*, *4*, 181–202.

Unterrainer, J. M., Rahm, B., Leonhart, R., Ruff, C. C., & Halsband, U. (2003). The Tower of London: The impact of instructions, cueing, and learning on planning abilities. *Cognitive Brain Research*, *17*, 675–683.

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

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

Ward, G., & Allport, A. (1997). Planning and problem-solving using the five-disc Tower of London task. *The Quarterly Journal of Experimental Psychology*, *50A*, 49–78.

Witkin, H. A., & Goodenough, D. R. (1981). *Cognitive styles: Essence and origins.* New York: International Universities Press.

APPENDIX

Mathematical details of the model

Visual module

Visual processing is based on the family of Gabor filters derived by Lee (1996), which satisfies both mathematical and neurophysiological constraints:

$$\psi(x,y,\omega_o,\ \theta) = \frac{\omega_o}{\sqrt{2\pi\kappa}} \cdot \exp\left\{\left(\frac{\omega_o^2}{8k^2}\right) \cdot [4(x \cdot \cos\theta y \cdot \sin\theta)^2 + (-x \cdot \sin\theta + y \cdot \cos\theta)^2]\right\} \cdot \left\{\exp[i(\omega_o \cdot x \cot\cos\theta + \omega_o \cdot v \cot\sin\theta] - \exp\left(\frac{k^2}{2}\right)\right\} \quad 1$$

where x and y represent the centre of the wavelet, ω_o is the spatial frequency in radians per unit length, and θ is the wavelet orientation in radians. K is a constant set to π (Lee, 1996). The real and imaginary parts of the complex function produce two filters, referred to as odd and even. In the present work, we varied only the wavelet orientation (0, 1/4;π, 1/2;π, and 3/4;π) for a total of eight filters (four even and four odd), while the spatial frequency was fixed at 0.57 radians. Four orientations constitute the minimal set of filters and are sufficient to provide rotation and size invariance (Riesenhuber & Poggio, 1999). We use a single, low-frequency bandwidth because it is more suitable to detect the main directional features of the entire stimulus. Note that adding more spatial frequencies, and thus more Gabor filter maps, did not improve the performance of the model.

Competitive selection module and top-down controller

Activation of each heuristic unit y_i at time t in the competitive selection module is obtained by summing the feedforward activation from the input layer I_i (visual module) and the recurrent input from the lateral connections in the heuristic layer.

$$y_i^t = I_i + \left(\sum_k y_k^{(t-1)} * w_{ik}\right) \quad 2$$

where y^{t-1} is the activation of the kth heuristic unit (including itself) at the previous time step $(t - 1)$, and w_{ik} indicates the weight of the corresponding lateral connection. The latter are fixed to the values of 0.2 and -0.1 for self-excitatory and lateral inhibitory connections, respectively. The feedforward activation from the visual module to each ith heuristic unit is calculated as follows:

$$I_i = \sum_j w_{ij} x_j \tag{3}$$

where x_j is the activation of the jth input unit, and w_{ij} is the weight of the corresponding connection.

The activation equations are run iteratively until one heuristic unit wins the competition. Note that the relaxation is driven by the lateral connections because I_i remains constant. To speed up the relaxation process we terminate the competition whenever one unit reaches a value of 0.9 (instead of the maximum of 1.0). At that point, the winning unit is set to 1, whereas all other units are set to 0.

During the learning phase, the feedforward weights between jth input units and the winning heuristic unit y are updated according to a Hebbian learning rule (Brown & Chattarji, 1995):

$$\Delta w_j = \eta(x_j w_j) y \tag{4}$$

where η is the learning rate (set to 0.1). Note that competitive learning (Rumelhart & Zipser, 1985) sorts patterns sharing similar properties into the same category, and it can be viewed as a clustering technique.

The effect of the TDC is simply to reset to zero the activation of all heuristic units in the competitive selection module. The efficiency of the TDC in resetting the heuristics nodes is decreased after a simulated lesion. That is, a residual activation (proportional to the severity of the lesion) persists in the heuristic nodes. For example, for a 20% lesion the residual activation of the winning heuristic unit corresponds to 80% of the activation at $t = 1$. The residual activation y of the winning unit is calculated as follows:

$$y = \frac{I_i * L}{100} \tag{5}$$

where L is the severity of the TDC lesion expressed as a percentage. Thus, when L is set to zero there is no residual activation.

Spatial module

Activation of each unit in the spatial target map is calculated as follows:

$$O = f\left[R_i(1 - S_i) + \sum_j W_{ij} O_j\right] \tag{6}$$

where R_i is the input from retinal (input) units (see below) and S_i is the activation value of the ith unit in the saliency map. The rightmost term of the equation computes the recurrent input resulting from the lateral connections W with the other units. Finally, $f(x)$ is a squashing function that bounds the activation in the [0, 1] range:

$$f(x) = \begin{cases} \dfrac{2}{1 + e^{-x}} - 1 & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

The retinal input R_i to each unit in the spatial target map is calculated as follows:

$$R_i = \exp\left(-\frac{d_{ri}^2}{2\sigma^2}\right) \tag{8}$$

where d_{ri} is the distance between the centre of the retinal receptive field (r_{xi}, r_{yi}) of the spatial unit and the retinal coordinates (r_x, r_y) of the visual target—that is, $d_{ri}^2 = (r_x - r_{xi})^2 + (r_y - r_{yi})^2$; σ is the width of the Gaussian (set to 5°). The receptive field centres were spread uniformly between $-40°$ and $+40°$ on both x and y axes, with increments of 4°.

The spatial target map contains symmetric lateral connections with fixed-value inhibitory weights that depend on the distance between neurons:

$$W_{ij} = \min\left[0, A_E \exp\left(\frac{d_{ij}^2}{2\sigma_E^2}\right) - A_1 \exp\left(-\frac{d_{ij}^2}{2\sigma_I^2}\right)\right] \tag{9}$$

where d_{ij} is the distance between the two neurons. The connections weights cannot have a positive value. A_E and σ_E are always higher than A_I and σ_I, respectively.

The exact location represented by the spatial target map is decoded through a simple vector method (Salinas & Abbott, 1995):

$$(O_x, O_y) = \left(\frac{\sum_i O_i O_{xi}}{\sum_i O_i}, \frac{\sum_i O_i O_{yi}}{\sum_i O_i}\right) \tag{10}$$

where (O_x, O_y) is the location of the planned movement, O_i is the activation value of the ith unit in the spatial target map, and O_{xi} and O_{yi} are the field centre coordinates of the ith unit.