Traveling salesman problem: The human case

Jan M. Wiener, Thora Tenbrink

In the field of human cognition, performance and optimization behavior in the TSP has mainly been investigated by means of visual versions in which humans are confronted with a number of dots on a computer monitor. Their task is to connect these dots by a straight line such that the resulting path is optimal with respect to overall length. Path planning tasks similar to the TSP are quite common also in everyday navigation, for example, in shopping routes. In this paper, we systematically disentangle the cognitive processes and the range of external factors that influence problem solving in tasks that resemble the classical TSP, covering the area so as to include everyday human navigation tasks. We identify those areas for which human heuristics and strategies are already known and work out hypotheses concerning the generalizability of results gained within particular subfields of the area.

1 The traveling salesman problem

Imagine a salesman who has the task of visiting a range of cities in order to sell his goods. He will probably start from a specific location to which he will also return after visiting each place just once. His route will be planned in a way that avoids detours and unnecessary effort; in short, the experienced traveling salesman will aim for the shortest possible route. Motivated from such scenarios, the well-known generalized traveling salesman problem (TSP) is formulated as follows: Given a number of target places and the costs of traveling from any target to any other target, what is the cheapest round-trip route that visits each target place once and then returns to the starting place? The TSP has been shown to be NP-complete; it becomes increasingly complex and computationally expensive with increasing numbers of targets. The only solution to the TSP that guarantees finding the optimal tour is to calculate and compare the length of all possible permutations, i.e. ordered combinations. Considering that the number of permutations is n! (the factorial of the number of places to visit), the TSP rapidly becomes impractical: while the number of permutations for visiting 6 places is 720, for 11 places it is already 39 916 800. However, several heuristics have been developed, some of which are introduced below, that greatly reduce computational effort while still resulting in near-optimal solutions. The TSP has received much attention since it was first introduced as a mathematical problem in the 1930ies. In fact, it is one of the most intensively studied problems in computational mathematics, yet no efficient solution method is known for the general case. From the point of view of cognitive science, this problem is of interest because humans quickly find good solutions to TSPs with, for instance, 11 target places, evidently without computing all possible permutations. Also, there are a number of related problem scenarios such as the "open TSP" in which the starting point is not visited again, the "orienteering problem" which involves different target weightings, further sequential path decision problems, and search and exploration issues that in crucial respects resemble the TSP. Nevertheless, in spite of the intuitive name of the TSP, there are few attempts to establish the relationship to human navigation in natural environments. In this article, we compare various different versions of the original computational TSP for the human case, including more naturalistic human behavioral variants in similar tasks (such as shopping routes), and discuss how different methodologies may shed light on different aspects involved in such a task.[1]

1.1 Strategies

Visual versions of the TSP. A great amount of attention and a broad range of experimental studies have been devoted to visual versions of the TSP (e.g., [11, 12, 18, 8, 19]). In these experiments, participants are usually confronted with a number of points on a computer monitor. Their task is to connect these points such that the resulting path is optimal with respect to the overall length. Generally, results from such experiments demonstrate that the participants are impressively good at solving visual TSPs, that inter-individual performance differences are rather small, and that performance is influenced both by the problem size and by the distribution of the targets. Performance is measured by relating the chosen solution to the optimal solution and is expressed as percentage above optimal (PAO).

There is an ongoing debate on the strategies participants apply in such experiments. It has, for example, been proposed that humans apply the **convex hull method** [11]. The convex hull is easily visualized by imagining an elastic band stretched open to encompass all dots; when released, it will assume the shape of the convex hull, touching all boundary dots of the TSP (the remaining dots are referred to as interior places/dots). Afterwards, the segment of the elastic band which is closest to an unconnected dot will be stretched to include that dot into the tour. This latter step is repeated until all points are incorporated in the overall tour. This heuristic has been further refined with respect to the exact method of incorporating interior points, either by identifying minimal distance (**cheapest insertion criterion**) or by relying on angles (**largest interior angle criterion**) [12]. The fact that a tour that follows the convex hull method is by definition free of crossings and that humans tend to avoid crossings is seen as one important piece of supporting evidence for the convex hull method [11]. This line of argumentation, however,

---

[1] Funding by the Volkswagen Foundation is gratefully acknowledged.

also works the other way around: if humans know that crossings will result in sub-optimal solutions they will consequently avoid crossings when solving TSPs, rather than following the convex hull method [18]. Tours or paths that result from convex hull strategies and crossing avoidance strategies are often very similar, if not identical.

Another strategy, the **hierarchical nearest neighbor** (hNN) method states that participants first establish clusters of several points based on NN distances, which they then link into a tour, using some variant of the nearest neighbor (NN) algorithm [19]. The NN-algorithm is one of the simplest procedures to solve TSPs quickly. From its current location, the start location in step one, the NN algorithm repeatedly visits the closest non-visited target location until all targets have been visited [7]. The NN-algorithm has been recognized for decades and is still regarded as one of the major options available to problem solvers in TSP related tasks as it is known to result in reasonably good solutions, at least, for small TSPs. The **hierarchical pyramid algorithm** assumes that from the original stimulus (point pattern) a series of images are generated which are increasingly blurred and compressed [8]. By these means, a hierarchy of images is generated in which neighboring dots collapse to clusters. The algorithm then starts with generating a tour upon an image which is so high in the hierarchy that only 3 clusters exist. By progressively moving to the next lower layer in the hierarchy further clusters, and eventually dots, are inserted into the tour.

Both of the latter approaches for solving visual TSPs, essentially reduce the complexity or computational effort of the planning problem by using hierarchical levels of abstraction that are used to plan coarse tours that are then refined. By first planning coarse tours a large number of (usually) suboptimal solutions is ruled out and search space is substantially reduced.

Navigational versions of the TSP. Path planning and wayfinding tasks similar to the TSP are quite common also in everyday navigation, for example, considering a typical shopping route on which multiple stores have to be visited. Pedestrian shopping behavior has been investigated with respect to distance minimization in multi-stop shopping routes [6]. Most shoppers first chose the location farthest away, probably to minimize effort to carry bought goods, and then minimized distances locally between shopping locations. This strategy of minimizing traveled distances between target places is similar to the NN-strategy.

Our own research uses a smaller-scale scenario involving 25 places marked by uniquely colored symbols, arranged on a regular grid in a large experimental room [20]. Participants were asked to visit a given subset of these places. Results show that the standard NN-strategy is not sufficient to explain planning behavior. Participants combined neighboring target places to clusters and preferred to visit large rather than small target clusters first (**cluster strategy**, cf. [21]). Furthermore, participants demonstrably used a **region-based planning strategy**, first planning a coarse path on the region level and then refining this plan during navigation by inserting close-by target places (for similar strategies in large scale spaces see [21]). Such hierarchical planning schemes allow for the reduction of computational effort and memory load during planning while still resulting in reasonably short paths. Our results are supported by linguistic evidence, i.e., by a close analysis of systematically elicited verbal representations of the task [16]. These were collected both during and after the experiment and reflected the participants' navigational strategies in several ways.

2 Visual and Navigational TSPs

Specific features of the task presentation and design may either highlight or obscure crucial factors of heuristical planning in TSP-like tasks, they may involve either more or less demands on memory, and they may trigger or lead to the neglect of certain cognitive strategies. In the following, we present a systematic account of the generalized features in which tasks (as employed in TSP research so far or observable in everyday chores) may differ. The most crucial difference is clearly posed by the contrast between visual and navigational versions of the TSP, which we take as our starting point for closer examination.

2.1 Scale

Navigational TSPs take place at a much larger scale than visual TPSs, which are usually presented as a simple graphic on a computer screen. This scale of space is often referred to as **figural space**. However, even between various versions of navigational tasks, there may be considerable differences in scale. Experiments range from indoor **vista space** (e.g., [20]) to **large scale space** such as outdoor scenarios [6]. In figural space all information is visually accessible at all times. In vista space (describing the space that is visible from a current location) this is also the case in principle; however, if the tasks requires active navigation, relevant information might, at least temporarily, be out of sight. In large scale navigational spaces, such as cities, information about the relevant spatial locations can usually not be perceived directly. Here, navigators either have to consult external representations of space, such as street maps, or the information has to come from cognitive-map like spatial memory (see also Section 2.4). Furthermore, when navigating vista or large scale spaces, spatial relations between the navigator and relevant locations constantly change. Accordingly, navigators permanently have to update the locations of places, deal with perspective changes, and, in contrast to tasks on the figural space, they usually do not have an overview over the environment as a whole.

TSPs in real world large scale spaces furthermore involve taking multiple alternative paths between target places into account during planning. In city-like environments it is usually not possible to navigate via straight lines between places; each subroute constitutes its own planning problem. The complexity involved here is well proven by the broad range of studies addressing wayfinding tasks that only concern a start and a single goal. Such studies, spanning research on route planning [2] as well as linguistic route descriptions [4], are typically carried out in large scale space. For experiments directly addressing TSP-like tasks, in contrast, this approach is largely unknown.

Up to now, the results in the literature are far from being systematic as far as the effects of scale are concerned (though see [13] for a detailed account). Currently, no clear predictions can be formulated as to the ways in which planning strategies and heuristics in TSP-like tasks will be affected by scale (and scale alone, considering that most relevant approaches differ in many other respects as well). It is conceivable, however, that the mere fact that the environment is or is not visible at one glance, or by looking around (within a room), or by traveling for a few minutes (as opposed to hours, as in travel planning), fundamentally influences planning strategies, even if the main task remains constant.

2.2 Goals and distractors

Navigation tasks typically involve a far lesser number of goals than visual versions of the TSP, in which up to 100 goals are to be connected [12], such a number would not be feasible whenever real navigation is involved. One major feature in which visual TSPs often systematically differed concerned the distribution of the goals within the overall field; they may be randomly distributed, grouped systematically, lined up in rows, spread out in patterns, etc. In navigational TSPs, as well, the distribution of goals stands out as a major factor of interest. Typically, this kind of variability is used as a crucial design feature to identify humans' planning strategies in the given task. Certain distributions of goals may be particularly suitable for triggering a specific strategy out of a range of options available to humans. Furthermore, distributions can be designed in a way as to disprove that particular strategies generally dominate users' behavior [20].

Furthermore, tasks differ crucially with regard to the presence of distracting places other than the goals. Also, if there are more places than actual targets to be visited, there may or may not be a given number or list of goals. In visual TSPs, usually all presented locations are target places for the current trial. Everyday navigation tasks such as shopping routes, in contrast, take place in complex spatial environments involving many places and objects. Only some of these need to be visited as part of a given task scenario, the remaining majority are distractors. Few experimental studies addressing TSP-like tasks have so far been designed to address humans' strategies in naturalistic scenarios (e.g., [6, 9]). Other studies take a step in this direction by purposively involving distractor objects [20, 16].

A further point concerns the nature and features of the targets. Natural target locations in the real world are associated with human activities and background knowledge of various kinds. In a shopping task, for instance, it is a matter of world knowledge to determine how the goal locations relate to each other [14]. There may also be various levels of importance or saliency that influence the planning process [9]. In general terms, it is likely that grouping processes have an impact on everyday navigation planning. The importance of such processes in more abstract TSP tasks was confirmed by our studies [20, 16], where the different colours and shapes of the symbols played a major role for the planning process. Such findings cannot be obtained by studies employing only identical targets, as in the visual TSPs. Additional effects, for example, of individual associations with objects or locations can be expected considering humans' particular choices when given freedom in deciding about the places to be visited. In this direction, first results involving a travel planning situation based on a map [15] again indicate effects of conceptual grouping and regionalization.

2.3 Setting

A visual TSP task, which is maximally abstract and remote from real navigation settings, does not necessarily involve the notion of navigation; participants may simply be told to connect the points by as short a line as possible. Alternatively they may be asked to find a route (invoking a mental navigation situation) that visits each place indicated by abstract symbols. This may take place either on a computer screen or by using paper and pencil; both of these are common in visual TSP tasks. In shopping and traveling situations, the people themselves move around in a large-scale environment, in corresponding experimental tasks, their trajectory is measured.

There are further possibilities, mediating between the highly restricted figural scale space (e.g., on a computer screen) and the (non-controllable) high complexity of naturalistic large-scale environments. On the one hand, a situation close to real navigation can be created in virtual reality, which provides considerable flexibility in experimental design [21]. On the other hand, participants may be asked to imagine traveling in a large scale area. In an innovative experimental paradigm [1], blindfolded humans can imagine any spatial scenario while walking on the spot. A different approach is to have participants imagine the planning of complex routes, for example, based on a map. This may involve the mental organization of multiple goals, as in organizing a busy day in a city [9] or making travel plans for a holiday in a distant country [15]. In such tasks, participants are asked to find a solution to a TSP-like problem without actually carrying out the navigation itself.

Finally, the nature of the representation that is measured and analyzed towards solving the TSP task may be varied. Usually, trajectories are measured in some way, by connected lines or by tracing the paths participants travel. However, further crucial insights can be gained from the analysis of humans' language about spatial strategies; this has been well investigated for wayfinding tasks between two locations, involving everyday human route descriptions [4]. We have started to explore this approach for TSP-like tasks [16]. Linguistic tasks may involve thinking-aloud protocols [5, 10], or the participants may be asked to describe the task as such, give instructions to someone else (e.g., route descriptions), or discuss it dialogically in various ways. Each version yields a different text type with systematic inherent features which need to be differentiated linguistically from the insights to be gained about the problem solving issue itself. The investigation of planning strategies solely on the basis of behavioral data (such as trajectories) has its limitations if a particular solution can be reached with multiple solution heuristics. This problem becomes particularly eminent if overall performance is close to optimal. If, for example, participants find the optimal solution for a given TSP, and different TSP-heuristics predict optimal solutions, it is impossible to determine which strategy participants applied. In such situations, linguistic data provide insights into underlying conscious or even subconscious heuristics that can be tested in follow-up studies.

2.4 Memory

Differences in the TSP setting can affect the demands posed on memory. Solving visual TSPs is often thought to primarily involve perceptual processes, while solving navigational TSPs usually requires different memory systems, such as a working and a long-term memory of target positions. Such factors can be systematically manipulated either by targeted experimental navigational TSPs or by introducing memory demands in visual

TSPs (see Section 3). Manipulations range from connecting all targets to marking the relevant targets in the field by the experimenter to providing the participant with a list of targets to be visited. In the latter case, further differences in memory load concern whether the goals already visited may or may not be marked (ticked off) by the participant, either in the field itself or in the list, or both. In visual TSPs, the path from the start to the current location is displayed as a polygonal line segment, and participants are allowed to correct their choices by undoing links between points. Both of the latter properties are not available during real world navigation. Displaying the chosen path up to the current position provides participants with (i.) visual information of the overall shape of the trajectory, and (ii.) it retains spatio-temporal information, i.e. the sequence in which targets have been visited. Both of these latter factors might influence the problem solving strategies applied. In all these respects, navigational TSPs will generally pose higher cognitive demands on the problem solver than visual TSPs.

In everyday path planning (e.g., [17]), navigators are faced with additional memory tasks that are absent in visual TSPs. For example, if no external representation of space (e.g., a street map) is available, the target locations have to be retrieved from the cognitive map. Additionally, once the targets have been localized in the cognitive map, they have to be held active in a working memory representation during the actual planning process. Here, it can be assumed that navigators have to deal with different memory related constraints, such as imprecise spatial knowledge or capacity limits of working memory [3].

3 Comparing the results

Table 1 sums up the major aspects discussed in the previous section, abstracting from the gradual differences that are conceivable and that have already been explored in the literature. While this list is far from being exhaustive, it demonstrates that solving visual TSPs and navigational TSPs cannot, a priori, be assumed to be based on the same cognitive principles. However, there are undeniable parallels especially with respect to the task itself. Therefore we briefly compare the main results from visual TSPs and on navigational TSPs, as far as they are available. To our knowledge, the only real navigation studies so far explicitly targeting the investigation of human performance and cognitive strategies underlying path planning in the TSP are [20, 16] for vista space, [6] for large-scale settings, and a few further studies that are either not directly concerned with solving the TSP task or that do not involve real navigation (e.g., [9, 15]).

For both visual and navigational TSPs, planning performance decreases with increasing TSP size. However, the overall performance level is clearly superior in visual TSPs. While this difference is unsurprising given the mental effort required in each case, it cannot easily be quantified due to the many factors that distinguish navigational and visual TSPs. Nevertheless, there are striking similarities regarding the proposed solution strategies. For visual TSPs, the hierarchical NN strategy [19] and the hierarchical pyramid algorithm [8] have been proposed (see Section 1.1). Both strategies essentially state that, based on perceptual grouping, abstractions of the original problem are generated and that first coarse tours are planned upon these abstractions. Then, more detailed tours incorporating all targets are generated based on the coarse path plans. For navigational TSPs, we have proposed a region-based planning heuristic and a cluster strategy that are based on similar principles (see Section 1.1). All of these hierarchical strategies are based on chunking or clustering of spatial information. Furthermore, the "convex hull" strategy [11] parallels some of our linguistic findings in [16] in which participants attend particularly to the shape of the trajectory. Also, the importance of proximity is represented clearly in our linguistic data, reminiscent of the earlier idea of a "Nearest Neighbor" (NN) strategy, which is however not sufficient to explain the data as such. In line with this, proximity is consistently mentioned along with other aspects in our data.

4 Conclusions and future work

According to the overview presented here, the investigation of TSPs in figural space is not sufficient to understand how the traveling salesman really solves his problem of finding the shortest route between a given number of places. In spite of the various systematic differences between visual and navigational TSPs that we have pointed out, however, the problem solving strategies allowing to handle complex and computationally expensive TSPs may at least be of a similar nature. It will be a challenging endeavor for future investigations to address systematically in how far generalizations and transfers can be drawn from existing computer-based studies to realistic task settings, and to spell out the differences and their specific reasons. It is a desirable long-term aim to tease apart the precise impact of each of the factors involved, including their relation to and interplay with each other. To develop an increased understanding of how the cognitive processes involved in solving visual and navigational TSPs relate to each other, future experiments in both domains need to be carried out that are matched more closely, and therefore more directly comparable.

References

[1] D. Vann Bugmann and K. R. Coventry. A method for studying representation of action and cognitive distance. In K.D. Forbus, D. Gentner, and T. Regier, editors, *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, pages 1381–1386. Lawrence Erlbaum Associates, Mahwah, NJ, 2004.

[2] N. Christenfeld. Choices from identical options. *Psychological Science*, 6:50–55, 1995.

[3] H. Couclelis. Verbal directions for way-finding: space, cognition, and language. In J. Portugali, editor, *The construction of cognitive maps*, pages 133–53. Kluwer Academic Publishers, Dordrecht, 1996.

[4] M. Denis. The description of routes: A cognitive approach to the production of spatial discourse. *Cahiers de Psychologie Cognitive*, 16(4):409–458, 1997.

[5] M. E. Fonteyn, B. Kuipers, and S. J. Grobe. A description of think aloud method and protocol analysis. *Qualitative Health Research*, 3(4):430–441, 1993.

[6] T. Gärling and E. Gärling. Distance minimization in downtown pedestrian shopping. *Environment and Planning A*, 20:547–554, 1988.

[7] B. Golden, L. Bodin, T. Doyle, and W. Stewart. Approximate traveling salesman algorithms. *Operations Research*, 28:694–711, 1980.

	navigation	setting	scale	number of goals	presence of distractors	choice of goals	features of goals	marking of visited goals	distribution of goals	setting of task
visual TSPs	no / imagined	computer / paper	small scale	up to 100	no	(no distractors)	identical	by a connecting line	random / clustered / grouped	laboratory
navigational TSPs	yes	symbolic / virtual reality	medium / large scale	up to 9	yes	by experimenter	systematic	no marking / marking of locations in the field	random / clustered / grouped	laboratory
everyday navigation	yes	natural	large scale	indefinite	yes	by participant	natural	no marking	relationships between goals accessible by world knowledge	real world

Table 1: Schematized similarities and differences between existing empirical approaches

[8] S. M. Graham, A. Joshi, and Z. Pizlo. The travelling salesman problem: A hierarchical model. *Memory & Cognition*, 28(7):1191–1204, 2000.

[9] B. Hayes-Roth and F. Hayes-Roth. A cognitive model of planning. *Cognitive Science*, 3:275 – 310, 1979.

[10] C. Hölscher, T. Meilinger, G. Vrachliotis, M. Broesamle, and M. Knauff. Finding the way inside: Linking architectural design analysis and cognitive processes. In C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, and T. Barkowsky, eds., *Spatial Cognition IV*, pages 1–23, Berlin, Heidelberg, 2005. Springer.

[11] J. N. MacGregor and T. C. Ormerod. Human performance on the traveling salesman problem. *Perception and Psychophysics*, 58:527–539, 1996.

[12] J. N. MacGregor, T. C. Ormerod, and E. Chronicle. A model of human performance on the travelling salesperson problem. *Memory & Cognition*, 28(7):1183–1190, 2000.

[13] D. R. Montello. Scale and multiple psychologies of space. In A. Frank and I. Campari, editors, *Spatial Information Theory: A theoretical basis for GIS*, pages 312–321. Springer, Berlin, Heidelberg, New York, 1993.

[14] S. Ratneshwar, L. Barsalou, C. Pechmann, and M. Moore. Goal-derived categories: The role of personal and situational goals in category representations. *Journal of Consumer Psychology*, 10:147–157, 2001.

[15] I. Seifert. Region-based model of tour planning applied to interactive tour generation. In *Proceedings of HCI International 2007, Beijing*, 2006.

[16] T. Tenbrink and J. Wiener. Wayfinding strategies in behavior and language: A symmetric and interdisciplinary approach to cognitive processes. In T. Barkowsky, M. Knauff, G. Ligozat, and D. Montello, editors, *Spatial Cognition V*. Springer, Berlin, Heidelberg, 2007.

[17] S. Timpf and W. Kuhn. Granularity transformations in wayfinding. In C. Freksa, W. Brauer, C. Habel, and K. F. Wender, editors, *Spatial Cognition III*, pages 77–88. Springer, Berlin, Heidelberg, 2003.

[18] I. Van Rooij, U. Stege, and A. Schactman. Convex hull and tour crossings in the Euclidean traveling salesperson problem: implications for human performance studies. *Memory & Cognition*, 31(2):215–220, Mar 2003.

[19] D. Vickers, M. D. Lee, M. Dry, and P. Hughes. The roles of the convex hull and the number of potential intersections in performance on visually presented traveling salesperson problems. *Memory & Cognition*, 31(7):1094–1104, Oct 2003.

[20] J. M. Wiener, N.N. Ehbauer, and H. A. Mallot. Planning paths to multiple targets: memory involvement and planning strategies. submitted.

[21] J. M. Wiener, A. Schnee, and H. A. Mallot. Use and interaction of navigation strategies in regionalized environments. *Journal of Environmental Psychology*, 24(4):475 – 493, 2004.

Contact

Dr. Jan M. Wiener
CNRS-Collège de France
11, place Marcelin Berthelot
75005 Paris/France
Tel.: +33 (0)1 44 27 14 21
Fax: +33 (0)1 44 27 13 82
Email: mail@jan-wiener.net

Dr. Thora Tenbrink
Universität Bremen
Enrique-Schmidt-Str. 5
Tel.: +49 (0)421 218-64-212
Fax: +49 (0)421 218-98-64212
Email: tenbrink@sfbtr8.uni-bremen.de

Bild

Jan M. Wiener is currently a research scientist at the CNRS in Paris. He studied Biology at the University of Würzburg and finished his dissertation in 2004 on human path planning behavior at the Max Planck Institute for Biological Cybernetics in Tübingen. His main research interest concerns the cognitive strategies and information processing underlying wayfinding and navigation behavior.

Bild

Thora Tenbrink pursues research on wayfinding and language in a joint project with Jan Wiener, financed by the Volkswagen Foundation. In addition she is one of the principal investigators of the project I5-[DiaSpace] as part of the SFB/TR 8 Spatial Cognition (Bremen / Freiburg, financed by the DFG). Her main interest lies in using discourse analytic methods to investigate linguistic representations of cognitive processes.