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, how-

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

Auszug aus: Künstliche Intelligenz, Heft 1/2008, ISSN 0933-1875, BöttcherIT Verlag, Bremen, **www.kuenstliche-intelligenz.de/order**