

Path planning under spatial uncertainty

JAN M. WIENER, MATTHIEU LAFON, AND ALAIN BERTHOZ
LPPA, Collège de France–CNRS, Paris, France

In this article, we present experiments studying path planning under spatial uncertainties. In the main experiment, the participants' task was to navigate the shortest possible path to find an object hidden in one of four places and to bring it to the final destination. The probability of finding the object (*probability matrix*) was different for each of the four places and varied between conditions. Given such uncertainties about the object's location, planning a single path is not sufficient. Participants had to generate multiple consecutive plans (*metaplans*)—for example: If the object is found in A, proceed to the destination; if the object is not found, proceed to B; and so on. The optimal solution depends on the specific probability matrix. In each condition, participants learned a different probability matrix and were then asked to report the optimal metaplan. Results demonstrate effective integration of the probabilistic information about the object's location during planning. We present a hierarchical planning scheme that could account for participants' behavior, as well as for systematic errors and differences between conditions.

Path Planning Under Spatial Uncertainty

The selection between path alternatives, and the planning of novel paths, are essential and frequent tasks for human navigators. Several studies have addressed the underlying cognitive strategies and mechanisms (e.g., Christenfeld, 1995; Gärling & Gärling, 1988; Golledge, 1995; Hölscher, Meilinger, Vrachliotis, Brösamle, & Knauff, 2006; Wiener & Mallot, 2003; Wiener, Schnee, & Mallot, 2004). All prior studies on path selection and path planning behavior assumed or stated that all required spatial information was available. In real life, however, navigators have to deal with incomplete or imprecise spatial knowledge resulting in spatial uncertainties. They might, for example, be confronted with situations in which the exact location of a goal or goal place is unknown, or can be judged with limited probability only. In this article, we intend to investigate the strategies and heuristics employed by navigators during path planning under such spatial uncertainties.

The main experimental task was analogous to the following scenario:

On your way home from work you realize that you are missing your keys. You know, however, that your roommate, who has another set of keys, has an appointment for dinner in a nearby restaurant after work. Obviously, the probability to find your roommate at work or in one of the restaurants depends on the time of day. Given different timings, what is the best path (i.e., in which order should you visit your roommate's workplace and the nearby restaurants) to find him/her, get the keys, and get home as quickly as possible?

What is being described here is a path planning task with a destination, and with an intermediate target whose exact whereabouts are uncertain but which can be de-

scribed by a probability matrix over multiple locations. This probability matrix might change with different timings: If you leave work early, the probability of finding your roommate still at work is rather high. As time goes by, however, the probability of finding your roommate at work decreases, whereas the probability of finding him or her in a nearby restaurant increases.

Decision Making and Uncertainties

In the scenario described above, decisions have to be made about the order in which different locations are visited, although the exact whereabouts of the target can be described with only a certain limited probability. Probabilistic decision making has been extensively studied in nonspatial contexts in probability learning experiments, in which participants are faced with response alternatives that differ in terms of payoff (e.g., Shanks, Tunney, & McCarthy, 2002; West & Stanovich, 2003). Alternative A, for example, is rewarded in 70% of the choices, whereas Alternative B is rewarded in 30% of the choices. Rational choice theory requires that after learning these probabilities, participants should eventually allocate all their responses to the alternative with the higher payoff (Alternative A). However, many studies demonstrated that participants match their response probabilities to the payoff probabilities; that is, participants choose Alternative A in ~70% of the trials and Alternative B in ~30% of the trials (for an overview, see Vulkan, 2000). In contrast to rational choice, such probability matching behavior results in suboptimal overall payoff. In context of the spatial planning scenario introduced above, in which one faces a single trial only, both rational choice and probability matching predict that the different locations are visited in an order resembling their probability of containing the

J. M. Wiener, mail@jan-wiener.net

desired target. Note, however, that the spatial layout is not considered. In spatial terms, trajectories resulting from such nonspatial decision-making theories could, therefore, be rather long and inefficient.

Strategies and Heuristics Involved in Path Planning

Even without spatial uncertainties, planning reasonably short paths for visiting several locations is a complex and computationally expensive task. Several planning strategies and heuristics have been described that allow for a reduction of computational effort: A particularly simple strategy to solve path planning tasks with multiple target locations is the *nearest neighbor* (NN) strategy (e.g., Bureš, Burešová, & Nerad, 1992; Gärling & Gärling, 1988; Gärling, Säisä, Böök, & Lindberg, 1986; Hirtle & Gärling, 1992). It states that the closest target location is chosen repeatedly until all target locations have been visited. The NN strategy is a purely local strategy that is universally applicable if distances between locations are available and that is known to result in reasonably short paths in many situations (Golden, Bodin, Doyle, & Stewart, 1980).

A number of further planning strategies can be subsumed under the term *hierarchical planning schemes*. Essentially, hierarchical planning schemes such as the *cluster* strategy (Gallistel & Cramer, 1996; Wiener et al., 2004), region-based planning strategies (Wiener & Mallot, 2003), or the *hierarchical nearest neighbor* strategy (hNN; see Vickers, Bovet, Lee, & Hughes, 2003), reduce mental effort during planning by first generating approximate solutions using abstractions from the actual planning task. These approximate or coarse solutions (path plans) are refined in later planning steps. The cluster strategy, for example, assumes that neighboring target locations collapse into clusters and that paths are planned so that the largest cluster is visited first (cf. Gallistel & Cramer, 1996; Wiener et al., 2004). By collapsing multiple targets into one cluster that becomes a target in the first planning step, the problem size is considerably reduced, many suboptimal solutions are eliminated, and a simple yet coarse path plan is generated that is easily remembered and that can be refined during navigation by planning a detailed path within each cluster. Similarly, region-based planning strategies assume that target locations are assigned to different environmental regions and that paths are planned toward the region(s) rather than toward the single target places (e.g., Hölscher et al., 2006; Wiener & Mallot, 2003; Wiener et al., 2004). During planning, detailed spatial information is made available for the close surroundings only, whereas distant locations are represented at superordinate region levels. Such region-based planning will result in a detailed path plan for the close surroundings only, allowing for immediate movement decisions. During navigation, the plan then has to be refined as the navigator enters new regions.

Path planning and optimization behavior has also been investigated by means of visual versions of the *traveling salesman problem* (TSP; see, e.g., Graham, Joshi, & Pizlo, 2000; MacGregor, Ormerod, & Chronicle, 1999; Vickers,

Lee, Dry, & Hughes, 2003; Vickers, Lee, Dry, Hughes, & McMahan, 2006), in which participants are confronted with a number of locations (usually dots) on a computer monitor. Their task is to connect these locations such that the resulting path (tour) is optimal with respect to overall length. There is an ongoing debate on the strategies participants apply in these experiments. MacGregor and Ormerod (1996) proposed the convex hull method, assuming that participants use the convex hull as part of their strategy (see also MacGregor, Chronicle, & Ormerod, 2004; MacGregor, Ormerod, & Chronicle, 2000). They argue that the convex hull strategy is supported by the fact that a tour following the convex hull method is, by definition, free of crossings, and that humans tend to avoid crossings. Van Rooij, Stege, and Schactman (2003), however, argue that participants know that crossings will result in suboptimal solutions. They therefore proposed the *crossing avoidance hypothesis*, stating that—rather than following the convex hull method—humans avoid crossings when solving TSPs. Whereas the above-mentioned strategies are specific to the TSP, requiring the planning of closed loops, the hNN method is more general (Vickers, Bovet, et al., 2003). The hNN method assumes that participants first establish clusters of several dots based on NN distances, which they then sequentially link into a tour, using some variant of the NN algorithm.

To the authors' knowledge, path planning has not yet been investigated in situations in which participants were confronted with spatial uncertainties. All planning strategies described above are based solely on spatial information; hence, they do not allow predictions of how navigators might cope with spatial uncertainties. However, if not stated otherwise, the reported strategies can be applied to the spatial aspects of the present experimental scenario.

Motivation and Synopsis

When asked to plan a path to find an object that can reside in different places, one faces the problem that the exact target location can be determined with a certain limited probability, at best. It has not yet been determined how human navigators integrate such probabilistic information during path planning, so Experiment 1 was designed to address this question. In Experiment 2, participants were asked to plan paths of similar complexity as in Experiment 1, but without spatial uncertainties. Comparison of planning performance between Experiment 1 and Experiment 2 allowed the impact of spatial uncertainties during path planning to be quantified. Efficient planning requires comparing path alternatives and selecting between them according to criteria such as overall path length. For the planning tasks in Experiment 1, several alternative solutions existed that differed only marginally with respect to overall length. Experiment 3 quantified how well participants could perceive differences in the lengths of alternative paths of similar complexity to the paths navigated in Experiment 1. Results from all experiments are then discussed and related to predictions from different path planning strategies, as well as from probabilistic choice-making theories.

EXPERIMENT 1

Experiment 1 aimed at investigating whether human navigators can integrate probabilistic information about the exact target locations in their overall path planning scheme. Three general hypotheses that were tested here could be distinguished. (1) Path planning takes into account spatial (i.e., metric) information only. This hypothesis predicts that participants' planning and navigation behavior is not influenced by the probabilistic target information but can be described by spatial planning strategies alone. (2) Planning behavior depends only on the probabilistic information about the whereabouts of the target object. In this case, spatial information is ignored and the task resembles a nonspatial probabilistic decision-making task in which navigation behavior can be described by probability matching or rational choice alone. (3) Both spatial and probabilistic information about target locations are taken into account. In this case, navigation behavior changes when the probabilistic information about the exact target location changes and can only be explained by a combination of spatial path planning strategies and nonspatial probabilistic decision-making strategies.

Method

Participants

Fourteen participants (7 women) took part in the experiment. They were mostly postgraduate students from the lab and were unaware of the hypotheses at the time of testing.

Experimental Setup

Six positions within the experimental room (6×10 m) were marked with squares (30×30 cm; see Figure 1). One position was marked as the starting place; another as the destination. The remaining four positions were intermediate target places, one of which was marked as the origin (O). The others were marked by blank squares (Positions A, B, and C in Figure 1).

Experimental Task

The general experimental task was to plan and navigate the shortest possible path from the starting place to find an object hidden in




Figure 1. Experimental setup: Six positions within the experimental room were marked by squares.

one of the four intermediate target places and to bring the object to the destination. Participants were told that the object was at the origin (Place O) at the beginning of the trial and that it would jump to one of the remaining three intermediate target places (A, B, C) with a certain probability (*jump probability*). The likelihood that the object had jumped to Place A, Place B, or Place C was identical in each case. A single trial always started at the starting place and ended at the destination. During navigation, participants visited the intermediate target places by stepping onto the corresponding squares. If participants stepped onto the square containing the hidden (virtual) target, they received verbal feedback from the experimenter and proceeded directly to the destination.

Probability Matrix

The likelihood that the object jumped from the origin (Place O) to one of the three remaining intermediate target places (A, B, or C) resulted in a probability matrix describing the whereabouts of the object across the four intermediate target places. For example, a 50% jump probability resulted in a 50% chance of finding the object still at the origin (Place O) and a ~16.7% chance of finding it in any of the three remaining intermediate target places. In order to successfully solve the experimental task, participants had to learn this probability matrix.

Metapanning

It is important to understand that, in principle, the object could reside in any of the four intermediate target places. Given such uncertainties about the exact whereabouts of the object, planning a single path is not sufficient to solve the task. Participants had to generate what we refer to as *metaplans*. A metaplan consists of multiple consecutive path plans; for example, if one leaves the start place to search for the object in Place B and finds it, one proceeds directly to the destination. If, however, the object is not found in Place B, one has to navigate to the next place to search for it (e.g., Place C). Here again, two possibilities exist: The object is found and one proceeds to the destination, or the target is not found and one navigates to the next intermediate target place (see Figure 2, left panel). The metaplan is therefore not identical with the actual trajectory, but describes the order in which participants plan to visit the four intermediate target places, taking into account that they proceed directly to the destination once the target object is found, and that their task is to minimize overall path length.

Calculating the Optimal Metaplan

The optimal metaplan was different for each probability matrix. To calculate it, all possible orders for visiting the four intermediate target places were calculated. For each of these 24 sequences, the path lengths were calculated, given that the object resided in the four different intermediate target places (O, A, B, C). For example, if the metaplan was Start–B–C–O–A–Destination, and the object resided in Place B, the resulting path would be Start–B–Destination (5.24 m). If the object resided in Place C, the resulting path would be Start–B–C–Destination (5.91 m); if the object resided in Place O, the resulting path would be Start–B–C–O–Destination (7.32 m); and if the object resided in Place A, the resulting path would be Start–B–C–O–A–Destination (9.48 m). The path lengths for these alternatives were multiplied with the particular probability of finding the object in the respective locations (see Figure 2). By summing over the resulting four values, the average length of a given metaplan for a particular probability matrix was calculated. The metaplan that resulted in the shortest average path length is referred to as the *optimal* metaplan for the corresponding probability matrix. The optimal metaplans for the three different probability matrices are displayed in Figure 3.

Chance level performance was calculated for each probability matrix by averaging over the *percentage above optimal* (PAO) values of all 24 possible metaplans (see Table 1). It was 56.35 PAO for Probability Matrix I, 30.58 PAO for Probability Matrix II, and 18.75 PAO for Probability Matrix III.

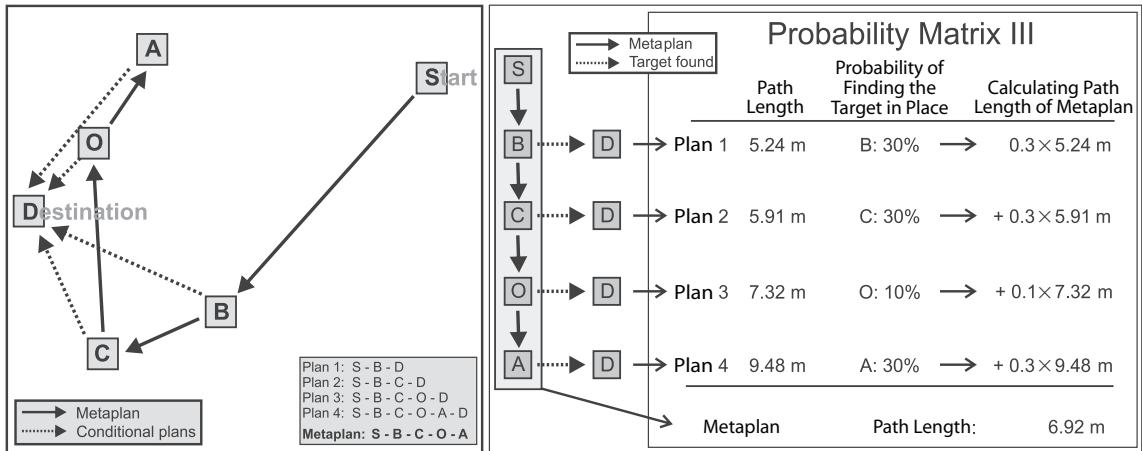


Figure 2. Left: Metaplan. Right: Exemplary calculation of the path length of a particular metaplan for the 90% jump probability condition (Probability Matrix III).

Experimental Procedure

Each participant took part in three experimental conditions. In each condition, the target's jump probability—that is, the probability that the object moved from the origin (Place O) to one of the remaining intermediate target places—varied, resulting in three different probability matrices. For Condition 1, the jump probability was 10% (Probability Matrix I); for Condition 2, it was 50% (Probability Matrix II); and for Condition 3, it was 90% (Probability Matrix III).

Figure 3 displays the resulting probabilities of containing the target object for all four intermediate target places. The sequence in which participants were exposed to the different conditions was randomized. Each experimental condition consisted of a training phase and a test phase.

Training phase. Participants were naive with respect to the specific probability matrix. To learn it, they received 30 training trials in each experimental condition. For example, for Probability Matrix II

Table 1
All Possible Metaplans for Visiting All Intermediate Target Places

Paths With Intermediate Target Places	Condition 1 Probability Matrix I			Condition 2 Probability Matrix II			Condition 3 Probability Matrix III		
	$p(0) = 90\%$	$p(A,B,C) \sim 3.3\%$		$p(0) = 50\%$	$p(A,B,C) \sim 16.7\%$		$p(0) = 10\%$	$p(A,B,C) = 30\%$	
	Length	PAO	Frequency	Length	PAO	Frequency	Length	PAO	Frequency
Start O A B C Dest.	475.59	0	9	645.97	7.72	3	816.36	17.86	
Start O A C B Dest.	478.00	0.51	3	658.03	9.73	2	838.05	20.99	
Start O C A B Dest.	492.47	3.55		730.39	21.79		968.31	39.80	
Start C O A B Dest.	740.53	55.71		816.94	36.23		893.36	28.98	
Start C O B A Dest.	747.12	57.09		849.89	41.72		952.67	37.54	
Start O C B A Dest.	480.60	1.05	1	671.05	11.90	1	861.49	24.38	
Start O B C A Dest.	481.95	1.34	1	677.78	13.02		873.61	26.13	
Start O B A C Dest.	491.41	3.33		725.07	20.91		958.73	38.42	
Start B O A C Dest.	629.84	32.43		710.89	18.54		791.94	14.34	
Start B O C A Dest.	636.34	33.80		743.37	23.96		850.40	22.78	
Start B C O A Dest.	727.75	53.02		710.20	18.43	1	692.64	0	5
Start C B O A Dest.	842.33	77.11		824.40	37.47		806.46	16.43	
Start C B A O Dest.	990.94	108.36		899.47	49.99		808.01	16.66	
Start B C A O Dest.	911.43	91.64		810.32	35.12		709.21	2.39	6
Start B A C O Dest.	1,136.26	138.92		977.26	62.96		818.27	18.14	
Start B A O C Dest.	777.20	63.42		779.70	30.02		782.21	12.93	
Start A B O C Dest.	837.83	76.17		813.98	35.73		790.12	14.07	
Start A B C O Dest.	938.48	97.33		826.96	37.90		715.43	3.29	1
Start A C B O Dest.	974.70	104.95		857.79	43.04		740.88	6.96	
Start C A B O Dest.	1,251.99	163.25		1,097.25	82.97		942.51	36.07	
Start C A O B Dest.	925.45	94.59		923.33	53.97		921.20	33.00	
Start A C O B Dest.	871.52	83.25		843.40	40.64		815.29	17.71	
Start A O C B Dest.	503.22	5.81		599.70	0	3	696.17	0.51	2
Start A O B C Dest.	503.32	5.83		600.17	0.08	4	697.02	0.63	
Chance level:		56.35			30.58			18.75	

Note—Average path lengths and percentages above optimal (PAOs) are displayed for each metaplan as well as for each probability matrix (condition). Frequencies with which participants have reported and navigated the corresponding metaplans are shown. Near optimal metaplans that resulted in less than 1.7 PAO are shown in boldface. Results from Experiment 3 suggest that the metaplans in boldface are so close in path length that, for participants, their length differences are essentially indistinguishable.

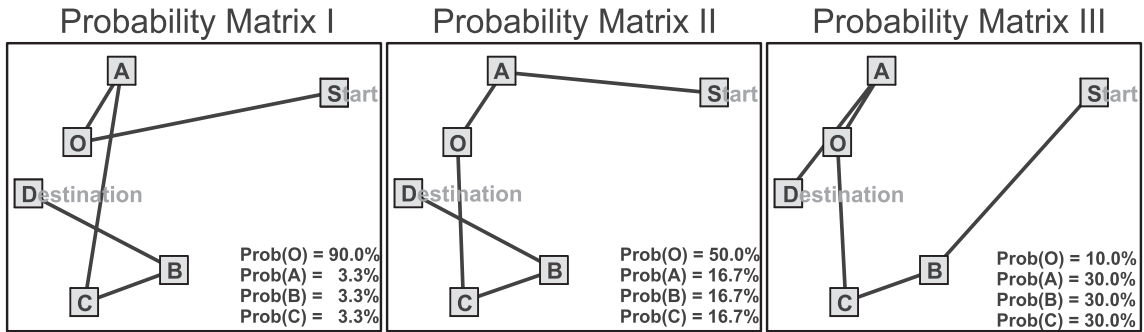


Figure 3. Optimal metaplans for the three different probability matrices.

(50% jump probability), participants would find the object in the origin (O) 15 times and 5 times in any of the three remaining intermediate target places (A, B, and C), adding up to a total of 30 trials. The order in which participants visited the intermediate target places was recorded for each trial.

Test phase. In the subsequent test phase, participants were asked to (1) judge the jump probability (i.e., the probability that the object had moved from its original position to one of the remaining intermediate target places) and (2) report and navigate the optimal metaplan. To clarify the latter task, participants were told that they should provide a friend unaware of the specific probability matrix with a sequence in which to visit the different target places in search of the objects.

Analysis

In order to evaluate path planning performance, the order in which participants visited the different places was recorded and compared with the optimal solution (*optimal metaplan*). Path planning performance is described by the measure PAO obtained by dividing the length of the traveled path by the length of the shortest possible path, subtracting 1, and multiplying the result by 100. A path with a PAO value of 10 is therefore 10% longer than the shortest possible path.

Results and Discussion

Jump Probability

On average, participants produced paths during the training phase with 9.20 PAO. Performance was significantly influenced by the specific probability matrix [$F(2,26) = 13.67, p < .001$]. Whereas performance was best (low PAO) for Probability Matrix I (PAO, 3.37), it was comparable for Probability Matrix II (PAO, 13.62) and Probability Matrix III (PAO, 10.61).

Learning

In order to assess the learning of the specific probability matrix during the training phase, the 30 trials in each condition were subdivided into three blocks (Trials 1–10; Trials 11–20; and Trials 21–30). An ANOVA revealed a significant main effect of experimental block [$F(2,26) = 4.46, p = .02$]. Specifically, PAO significantly decreased between Block 1 (PAO, 11.40) and Block 2 (PAO, 8.25), and no further decrease was observed for Block 3 (PAO, 7.95).

Trajectories

Participants' trajectories during the training phase are displayed in Figure 4. Clearly, participants preferred different paths when confronted with different probability matrices.

Test Phase

Reported jump probability. On average, participants recognized and reported the actual jump probabilities very precisely. For Probability Matrix I (10% jump probability), they reported 9.4%; for Probability Matrix II (50% jump probability), they reported 47.1%; and for Probability Matrix III (90% jump probability), they reported 82.8%. In none of these cases were significant differences between reported jump probabilities and actual values found [t test 10%, $t(13) = 0.39, p = .70$; t test 50%, $t(13) = 0.93, p = .37$; t test 90%, $t(13) = 1.66, p = .13$].

Reported metaplans. On average, the reported metaplans resulted in 2.28 PAO. For Probability Matrix I, the reported metaplans generated 0.28 PAO; for Probability Matrix II, 5.23 PAO; and for Probability Matrix III, 1.33 PAO (see Table 1). An ANOVA revealed a significant main effect of the probability matrix [$F(2,26) = 8.06, p < .01$]. Specifically, PAO was highest for Probability Matrix II. In all three conditions, participants' planning performance clearly exceeded chance level performance (see Table 1) [Probability Matrix I, t test, 0.28 PAO vs. 56.35 PAO, $t(13) = 471.09, p < .001$; Probability Matrix II, t test, 5.23 PAO vs. 30.58 PAO, $t(13) = 15.85, p < .001$; Probability Matrix III, t test, 1.33 PAO vs. 18.75 PAO, $t(13) = 51.51, p < .001$].

To statistically assess differences in trajectories between experimental conditions (Probability Matrices I, II, and III), we analyzed participants' first movement decisions (i.e., which of the four intermediate target places was visited first). A chi-square test [$\chi^2(4, N = 42) = 40.7, p < .001$] revealed highly significant differences between conditions (see Figure 4). Table 1 explicitly lists which metaplans were reported and how often. For Probability Matrix I, 12 out of 14 participants reported the optimal metaplan, or a near optimal metaplan, that produced only 0.51 PAO (detour length < 3 cm). No participants reported metaplans producing more than 1.5 PAO, and all reported visiting Place O (origin) first; their first movement decision therefore followed the optimal metaplan. For Probability Matrix II, 7 out of 14 participants reported the optimal or a near optimal metaplan that produced only 0.08 PAO (detour length < 1 cm). The remaining 7 participants, however, reported a metaplan with over 7.7 PAO (detour length > 44 cm). Specifically, 6 out of these 7 participants

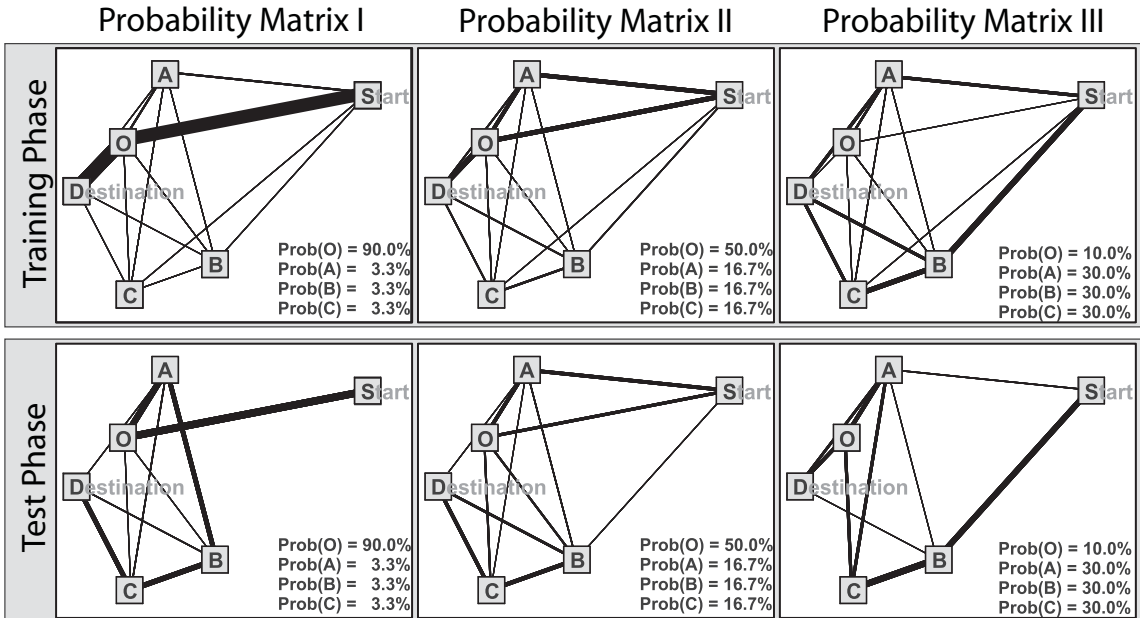


Figure 4. Participants' superimposed trajectories during the training phase (upper row) and during the test phase (lower row) for all three probability matrices. The thickness of the lines reflects the frequency with which the corresponding connections were traveled.

chose a metaplan that first visited Place O (origin) rather than Place A, as suggested by the optimal metaplan. For Probability Matrix III, 7 out of 14 participants reported the optimal or near optimal metaplan (0.51 PAO, detour length <4 cm). The remaining 7 participants reported metaplans that produced 2.39 PAO or 3.39 PAO. Interestingly, one very good metaplan that produced only 0.63 PAO (detour length <5 cm) has never been reported. In contrast to the optimal metaplan first visiting Place B, the latter first visited Place A.

Results from both the training and test phases clearly demonstrate that participants adjusted path planning and navigation behavior according to the different probability matrices, thus preferring different paths in each condition (see Figure 4 and Table 1). These results clearly contradict Hypothesis 1 (see above): that path planning under spatial uncertainties is not influenced by probabilistic information about the exact location of the target object but is based on spatial information only. Results also contradict Hypothesis 2: that path planning under uncertainty is based solely on the probabilistic information about the exact location of the target object. This hypothesis predicted that the intermediate target places would be visited in ascending order of their probability of containing the target object. Results for Probability Matrix II from both the training and test phases clearly show that in approximately 50% of the trials, participants first visited Place A, with only a ~17% probability of holding the target object, compared with Place O, with a 50% probability (see Figure 4). Furthermore, for Probability Matrix III, three intermediate target places with equal probabilities of holding the target existed (Places A, B, and C, each with a 30% probability). Hypothesis 2 pre-

dicted that these places would be visited in random order; clearly, that was not the case (see Figure 4).

Overall, path planning performance during the training and test phases clearly exceeded chance level performance for all probability matrices. Moreover, the result that the reported metaplans in the test phase were only ~2.3 PAO emphasizes participants' very good path planning performance. Together with the fact that the reported jump probabilities matched the actual values very well, these results suggest (1) that during the 30 trials of the training phase, participants learned the probability matrix to find the object in the four different places quite precisely; and (2) that participants very effectively integrated this probabilistic information into their path planning process, supporting Hypothesis 3.

EXPERIMENT 2

Experiment 1 demonstrated surprisingly good planning performance, even if the exact location of a target object could be determined with only a certain limited probability. Experiment 2 was designed to study planning performance in situations in which participants were not confronted with spatial uncertainties, but in which they had to navigate paths of similar complexity, as in Experiment 1. Comparison of planning performance between Experiments 1 and 2 will allow quantification of the impact of spatial uncertainties for path planning.

Method

Participants. Fourteen participants (2 women) took part in the control experiments. They were mostly postgraduate students from the lab who did not participate in Experiment 1.

Experimental setup. The experimental setup was identical to Experiment 1 but differed in the position of the start location. Altogether, three different start positions were used in Experiment 2, labeled in Figure 5 as Start Positions 1, 5, and 9.

Procedure. Participants were asked to navigate the shortest possible path from each of the three different start places to the destination (D), visiting all four intermediate target places (Places O, A, B, and C). For Start Position 1, the shortest path was Start–A–O–B–C–Destination; for Start Position 9, it was Start–B–C–A–O–Destination; and for Start Position 5, both alternatives were of equal length and no shorter path existed. The order in which participants started from the three start positions was randomized.

Results and Discussion

On average, the chosen paths in Experiment 2 resulted in a PAO value of 1.56; 1 participant was excluded from the final data set, whose PAO values overshoot average performance by more than two standard deviations. Eight of the 14 participants produced PAO values of below 0.01 on average (i.e., they found the optimal solution in most cases). The results demonstrate that most participants were extremely good at planning the shortest path from start to destination, with four intermediate target places. However, on average, planning performance in Experiment 2 did not exceed planning performance in Experiment 1 with spatial uncertainties [1.56 PAO vs. 2.28 PAO, *t* test: $t(25) = 0.64, p = .5$]. This null effect could be due to several reasons, and should therefore not be overinterpreted. However, as the existence of spatial uncertainties results in substantial additional planning costs, the lack of pronounced differences in performance between Experiment 1 and Experiment 2 at least emphasizes the fact that the strategies employed by participants in Experiment 1 very efficiently integrated the probabilistic information about the exact location of the target into the general planning scheme.

EXPERIMENT 3

Efficient path planning requires comparing path alternatives and selecting between them according to such criteria as overall path length. Table 1 demonstrates that for solving the path planning task in Experiment 1, several optimal or near optimal solutions existed that differed only marginally with respect to overall length. However, whereas some of these alternatives were chosen often, others were neglected. For the interpretation of these results, and for the further understanding of interactions between spatial and probabilistic information during path planning, it is important to develop a basic understanding of how well participants could distinguish between purely metric properties of path alternatives. Hence, Experiment 3 was designed to quantify how well participants could perceive differences in the lengths of alternative complex paths consisting of a start place, four intermediate places, and a destination place.

Method

Participants. The same 14 participants who took part in Experiment 2 also participated in Experiment 3.

Experimental setup. The experimental setup was identical to that in Experiment 1 but differed in the position of the start place.




Figure 5. Experimental setup of Experiments 2 and 3: In Experiment 2, Start Positions 1, 5, and 9 were used; in Experiment 3, all nine start positions were used. The numbers depicted in either the white or the gray part of the five target places (O, A, B, C, D) defined the two alternative paths for Experiment 3 (white path, A, O, B, C, D; gray path, B, C, A, O, D).

Altogether, 9 different start positions were used in Experiment 3 (see Start Positions 1–9 in Figure 5). We predefined two alternative paths, Start–A–O–B–C–Destination and Start–B–C–A–O–Destination, by marking the respective locations with numbers describing the order in which these places are visited along the paths (see Figure 5). One path was marked by blue digits, the other by black digits.

Procedure. Participants were positioned at the nine different start positions in random order and were asked to judge which of the two path alternatives was shorter from their current position. From Start Positions 1–4, Start–B–C–A–O–Destination was shorter; from Start Positions 6–9, Start–A–O–B–C–Destination was shorter; and from Start Position 5, both paths were of equal length (see Figure 5). When standing at Start Position 4 or 6, the two alternative paths differed about 15 cm in overall length, corresponding to a difference in overall path length of ~1.7%.

Results and Discussion

Figure 6 displays results from Experiment 3. From Start Position 5, both path alternatives were of equal length, and participants chose randomly between the alternatives. From Start Position 4, the shorter path had an overall length of 8.81 m and the longer path had a length of 8.96 m; the latter was thus ~1.7% longer. From Start Position 6, the shorter path was 8.86 m long; the longer path was 9.01 m and, thus, ~1.7% longer as well. From Start Positions 4 and 6, participants chose the shorter alternative in 82.14% of the cases, which is just above the commonly used threshold of 75%–80% when determining just noticeable differences (JNDs). Although Experiment 3 was not designed to perform an in-depth analysis of the JND for path-length differences, these results demonstrate that in the present experimental setup participants reliably detected differences in path length of 1.7% (1.7 PAO) and above.

How do these results relate to Experiment 1? The metaplans generated in Experiment 1 have a more complex structure than the paths in Experiment 3: Metaplans consider conditional paths, taking into account that the target

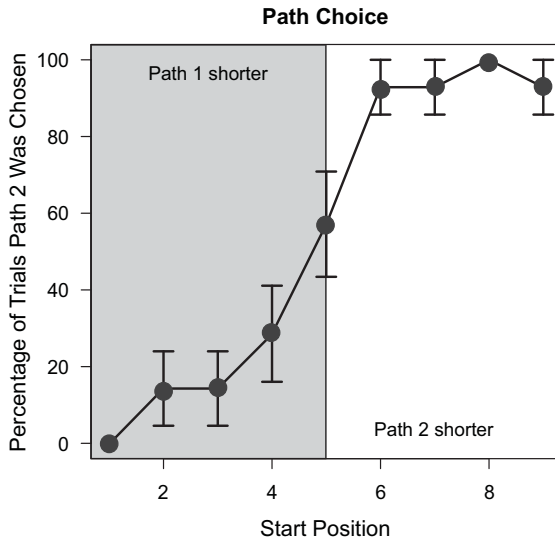


Figure 6. Results of Experiment 3: Percentage of trials in which Path 1 or Path 2 was chosen.

object can reside in different places. One might therefore assume that the threshold of 1.7% detected in Experiment 3 overestimates how well two alternative metaplans in Experiment 1 can be distinguished by length. On the other hand, one can be certain that metaplans that differ less than 1.7% in overall length could not be distinguished by participants. In Table 1, all metaplans that differed less than 1.7% in overall length—that is, that were indistinguishable according to their length—are presented in bold type.

GENERAL DISCUSSION

In this study, we investigated path planning under spatial uncertainties. For this, we developed an experimental paradigm in which participants had to find a hidden (virtual) object that could, with a certain probability, jump from its original location to one of the three remaining intermediate target places (Experiment 1). After training, the whereabouts of the object could be described by a probability matrix over four places. When one faces such uncertainties, planning a single path is not sufficient, since the exact object location is unknown. Instead, participants had to generate what we refer to as metaplans, considering multiple potential object locations. These metaplans describe the order in which the potential object locations will be visited, taking into account that one directly proceeds to the destination once the object is found and that the task is to minimize overall path length. Participants were confronted with three different probability matrices. For each probability matrix, a different metaplan was optimal. The chosen paths in Experiment 1 could be explained neither by assuming that during planning only spatial (metric) information was taken into account nor by assuming that participants visited the intermediate target places in the order that resembled their probability of

holding the target object. Results could be explained only by assuming that probabilistic information about the exact location of the target object and spatial (metric) information were integrated during path planning.

Overall, planning performance in Experiment 1 was very good; the reported metaplans in the test phase were only 2.3 PAO. The excellent planning performance under spatial uncertainties is further emphasized by a comparison with Experiment 2. In Experiment 2, participants solved planning tasks that required navigating paths of similar complexity to those in Experiment 1 but that did not require taking uncertainties into account. Nevertheless, planning performance in Experiment 1 and Experiment 2 was roughly comparable.

It is, of course, of great interest to understand how participants reached such good performance levels when planning paths under spatial uncertainty. One possibility is that they actually computed and compared all possible solutions, but that judgments of path length and/or distances between places were noisy and the sources of errors and variance in the data. There are, however, arguments against this interpretation. (1) We found systematic performance differences for the different probability matrices that cannot be explained by assuming nonspecific noise in the data. (2) The path planning task was computationally rather complex; the number of path alternatives that had to be considered clearly exceeded working memory capacities. (3) Results from Experiment 3 suggest that for paths with length and complexity similar to those of the metaplans of Experiment 1, participants could reliably discriminate paths that differed in 1.7% of overall length. In Experiment 1, however, participants often chose metaplans with PAOs far above 1.7, and systematically neglected metaplans with PAOs below 1.7 (see Table 1). Taken together, these considerations strongly suggest that participants applied planning strategies and heuristics, rather than actually computing the optimal solution.

Which planning strategies could account for the navigation data in Experiment 1? As discussed above, strategies based solely on spatial information, or solely on probabilistic target information, are not sufficient to explain the recorded data. We therefore assume a combination of spatial planning strategies and probabilistic decision-making strategies. As pointed out in the introduction, not all spatial planning strategies described in the literature allow deriving predictions in the present scenario. Hence, we concentrate on the following planning strategies: First, the NN strategy, predicting that the closest target place is repeatedly chosen until all targets have been visited (Bureš et al., 1992; Gärling & Gärling, 1988; Gärling et al., 1986); and second, hierarchical planning strategies that can, in principle, be applied to the present scenario—the cluster strategy (Gallistel & Cramer, 1996; Wiener et al., 2004), the hNN method (Vickers, Bovet, et al., 2003), and the region-based planning strategy (Wiener & Mallot, 2003). All three hierarchical strategies have this in common: that the environment is structured into the different clusters or regions taken into account during planning. In the present scenario, the

four intermediate target places (O, A, B, and C) would thus form two clusters ([O A] and [B C]), since distances within these clusters are smaller than any distance between clusters. This clustering is backed by informal interviews with participants after the experiments. The cluster strategy, stating that the larger cluster is visited first, can be translated to the present scenario quite easily: Whereas both clusters were of equal size, each containing two locations, the probability that they actually held the target object differed for all three probability matrices. In the present scenario, the cluster strategy would thus state that participants first visit the cluster with the higher summed probability to hold the target object. The cluster strategy is inherently spatial, based as it is on the distribution of places within the environment. However, in its proposed form, taking probabilistic information into account, it already resembles a simple combination of a spatial planning strategy and a probabilistic decision-making strategy.

Finally, we consider theories from (nonspatial) probabilistic decision making: rational choice and probability matching (e.g., Shanks et al., 2002). When participants are faced with a single trial only, both theories predict that the order in which alternatives are selected resembles their payoff probabilities. This can be directly transferred to the present scenario: Probabilistic choice making states that participants visit the target places in an order reflecting the probability that those places hold the target object. We refer to this strategy as the *rich-target-first* strategy.

We do find empirical evidence for all three proposed planning strategies: the cluster strategy, the NN strategy, and the rich-target-first strategy. This evidence is discussed in the following and a hierarchical planning scheme is presented allowing one to predict the empirical data of Experiment 1. First, a coarse path is planned on the level of the target clusters, describing in which order the two clusters are visited (cf. Gallistel & Cramer, 1996; Wiener et al., 2004): The cluster strategy (as proposed above) predicts that cluster [O A] is visited first in the 10% and the 50% jump probability condition and that cluster [B C] is visited first in the 90% jump probability condition (cf. Figure 3). In fact, participants' behavior was consistent with these predictions in all three conditions: Before visiting the other cluster, the great majority first visited the cluster with the higher probability of holding the target object. This is particularly interesting; for the 90% jump probability (Probability Matrix III), two metaplans existed that first visited the cluster with the lower probability (in this case [O A], with a summed probability of 40%) but that resulted in PAO values of only 0.5 and 0.6 (see Table 1). According to results from Experiments 2 and 3, these solutions were indistinguishable from the optimal solution but were chosen by only 2 out of the 14 participants. In contrast, a metaplan, resulting in 2.4 PAO, that first visits the cluster with the higher probability of holding the target object was reported by 6 of the 14 participants. Together, these results suggest that participants avoided first visiting the cluster with the lower probability, which in turn supports the cluster strategy.

Once a decision on the level of the clusters is made, the order in which the places within the clusters are visited is planned. Here, two strategies, the NN strategy (cf. Gärling et al., 1986) and the rich-target-first strategy, come into play. The interaction of these strategies is best demonstrated by analyzing the first step in the 50% jump probability condition. Here, the strategies made differential predictions: Given the decision to first visit cluster [O A], the NN strategy predicted that Place A (the closest target place) would be visited first, and the rich-target-first strategy predicted that Place O (with the highest probability of holding the target object) would be. Despite the fact that visiting Place O first resulted in clearly suboptimal solutions, with PAO values of at least 7 (see Table 1), participants chose to visit Place A first and Place O roughly equally often. Navigating an optimal or near optimal solution requires visiting Place A first. Apparently, in the 50% jump probability condition, the NN strategy and the rich-target-first strategy both influenced navigation behavior. A similar situation also occurred for the 90% jump probability: Here, the great majority of participants first visited Places B and C. With their next step, they could then either visit the closer Place O, which has a relatively low probability of holding the target object (10%), or the more distant Place A, which has a higher probability of holding the target object (30%). Again, participants chose between these alternatives roughly equally often (see Figure 4). Furthermore, results suggest that the influence of the NN and rich-target-first strategies was weighted according to the specific parameters: In the 10% jump probability condition, for example, the predictions for the first step of the NN strategy and the rich-target-first strategy were identical to those for the 50% jump probability condition. Nevertheless, all participants chose to visit Place O first. This difference can be explained by considering that in the 10% jump probability condition, Place O had a 90% chance of holding the target, whereas Place A had only a 3.3% chance (in contrast to 50% for Place O and 16.7% for Place A in the 50% jump probability condition). Apparently, the rich-target-first strategy had a stronger influence in the 10% than in the 50% jump probability condition, since the difference in probabilities to hold the target object was more pronounced between alternatives. It will be a challenging task for future studies to investigate the exact nature of these interactions in more detail.

Overall, we have presented a new experimental paradigm to investigate human path planning behavior under spatial uncertainties. We demonstrated that participants very successfully integrated probabilistic information about the whereabouts of a target object during path planning. Finally, we have presented a hierarchical planning scheme combining theories from spatial planning and probabilistic decision making that could account for participants' behavior, as well as for systematic errors and differences between conditions.

AUTHOR NOTE

We gratefully acknowledge funding from the Volkswagen Foundation and the EU Grant "Wayfinding" (6th FP NEST). M.L. benefits

from a grant from the French government and EDF for his doctoral research. Correspondence concerning this article should be addressed to J. M. Wiener, Laboratoire de Physiologie de la Perception et de l'Action, CNRS, 11, place Marcelin Berthelot, 75005 Paris, France (e-mail: mail@jan-wiener.net).

REFERENCES

- BUREŠ, J., BUREŠOVÁ, O., & NERAD, L. (1992). Can rats solve a simple version of the traveling salesman problem? *Behavioural Brain Research*, **52**, 133-142.
- CHRISTENFELD, N. (1995). Choices from identical options. *Psychological Science*, **6**, 50-55.
- GALLISTEL, C. R., & CRAMER, A. E. (1996). Computations on metric maps in mammals: Getting oriented and choosing a multi-destination route. *Journal of Experimental Biology*, **199**, 211-217.
- GÄRLING, T., & GÄRLING, E. (1988). Distance minimization in downtown pedestrian shopping. *Environment & Planning A*, **20**, 547-554.
- GÄRLING, T., SÄISÄ, J., BÖÖK, J., & LINDBERG, E. (1986). The spatiotemporal sequencing of everyday activities in the large-scale environment. *Journal of Environmental Psychology*, **6**, 261-280.
- GOLDEN, B., BODIN, L., DOYLE, T., & STEWART, W. (1980). Approximate traveling salesman algorithms. *Operations Research*, **28**, 694-711.
- GOLLEDGE, R. (1995). Path selection and route preference in human navigation: A progress report. In A. U. Frank & W. Kuhn (Eds.), *Spatial information theory: A theoretical basis for GIS (COSIT '95)* (Lecture Notes in Computer Science, No. 988, pp. 207-222). Berlin: Springer.
- GRAHAM, S. M., JOSHI, A., & PIZLO, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & Cognition*, **28**, 1191-1204.
- HIRTLE, S. C., & GÄRLING, T. (1992). Heuristic rules for sequential spatial decisions. *Geoforum*, **23**, 227-238.
- HÖLSCHER, C., MEILINGER, T., VRACHLIOTIS, G., BRÖSAMLE, M., & KNAUFF, M. (2006). Up the down staircase: Wayfinding strategies in multi-level buildings. *Journal of Environmental Psychology*, **26**, 284-299.
- 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. C. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics*, **58**, 527-539.
- MACGREGOR, J. N., ORMEROD, T. C., & CHRONICLE, E. P. (1999). Spatial and contextual factors in human performance on the traveling salesperson problem. *Perception*, **28**, 1417-1427.
- MACGREGOR, J. N., ORMEROD, T. C., & CHRONICLE, E. P. (2000). A model of human performance on the traveling salesperson problem. *Memory & Cognition*, **28**, 1183-1190.
- SHANKS, D., TUNNEY, R., & MCCARTHY, J. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, **15**, 233-250.
- VAN ROOIJ, I., STEGE, U., & SCHAFFMAN, 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., BOVET, P., LEE, M. D., & HUGHES, P. (2003). The perception of minimal structures: Performance on open and closed versions of visually presented Euclidean traveling salesperson problems. *Perception*, **32**, 871-886.
- VICKERS, D., LEE, M. D., DRY, M., & HUGHES, P. (2003). The roles of the convex hull and the number of potential intersections in performance on visually presented traveling salesperson problems. *Memory & Cognition*, **31**, 1094-1104.
- VICKERS, D., LEE, M. D., DRY, M., HUGHES, P., & MCMAHON, J. A. (2006). The aesthetic appeal of minimal structures: Judging the attractiveness of solutions to traveling salesperson problems. *Perception & Psychophysics*, **68**, 32-42.
- VULKAN, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, **14**, 101-118.
- WEST, R., & STANOVICH, K. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, **31**, 243-251.
- WIENER, J. M., & MALLOT, H. A. (2003). "Fine-to-coarse" route planning and navigation in regionalized environments. *Spatial Cognition & Computation*, **3**, 331-358.
- WIENER, J. M., SCHNEE, A., & MALLOT, H. A. (2004). Use and interaction of navigation strategies in regionalized environments. *Journal of Environmental Psychology*, **24**, 475-493.

(Manuscript received March 27, 2007;
revision accepted for publication October 12, 2007.)