Climbing the Mazes: A cognitive model of spatial planning

D. Fum (fum@univ.trieste.it) F. Del Missier (delmisfa@univ.trieste.it) Department of Psychology, University of Trieste, via Sant'Anastasio 12 I - 34100 Trieste, Italy

Abstract

The paper presents an experiment on spatial planning in a 2D maze-like computerized environment. Participants were required to build on-line a plan to reach a target, and to execute it in the shortest possible time. We discuss the experimental results, and propose a detailed simulative model to provide an accurate account of the participants' step-by-step behavior. The simulation results support the hypothesis of the use by participants of an incremental hill-climbing strategy minimizing the number of turns to be performed in solving the mazes.

Introduction

Spatial planning can be considered as a particular kind of problem solving in which participants try to get at a definite location or to optimize a performance measure in spatially constrained environments. In comparison with more typical planning tasks, spatial planning requires a tighter interaction between central and peripheral processes: visual, attentive and motor issues play in fact a fundamental role, in addition to reasoning, in determining the final behavior.

Mazes and maze-like tasks are a common tool used in studying spatial planning processes. In a neuropsychological setting they have been employed since Porteus and Kepner (1944). Research on these tasks has dealt with various topics, ranging from individual differences (Krikorian & Bartok, 1998) to functional localization (Flitman et al. 1998). Studies have been carried out about motivational factors (Hickman, Stromme, & Lippman, 1998) and various social phenomena (Stephenson & Wicklund, 1984). Cognitively oriented research has tackled issues like learning (Crowe et al. 1999), and working memory capacity (Glassman, Leniek, & Haegerich, 1998).

In spite of the widespread use of mazes in psychological research, testing and clinical practice, there are no adequately specified and empirically grounded cognitive models for these tasks. The paper presents an experiment on spatial planning in a 2D maze-like computerized environment. Participants were required to build on-line a plan to reach a target, and to execute it in the shortest possible time. We discuss the experimental results and propose a detailed simulative model to provide an accurate account of the participants' step-by-step behavior. The simulation results support the hypothesis of the use by participans of an incremental hill-climbing strategy minimizing the number of turns to be performed in solving the mazes.

The Experiment

Participants

The participants were 24 undergraduate and graduate students, aged between 19 and 28. They were all right-handed and none of them was suffering from any perceptual, spatial or motor deficiency. The sample was balanced for gender, and all the participants had approximately the same degree of familiarity with computers.

Materials

For the experiment we built an interactive environment, called PathWorld, that is both challenging and amenable to complete experimental control. PathWorld allows to generate bi-dimensional grids containing randomly placed obstacles in a number specified by the parameters of a given probability distribution. A start and a target position are also placed on the grid after the length of the optimal (i.e. minimum-length) path connecting the two points has been specified. Placing the obstacles and fixing the start and end positions generates an instance of a maze-like task.

Mazes of 22x22 cells were used in the experiment. Each cell was represented by a light gray square with a side length of 24 pixels. The top and bottom rows and the first and last columns were considered as borders and completely filled with obstacles (black squares which block any movement). The target cell was identified by a red-and black "bull's eye" while a green circle containing a styled black "man" constituted the token to be moved. The participants had the complete view of the maze.

Design and Procedure

Two within-subject independent variables (number of obstacles and minimum path length to the solution) were manipulated during the experiment following a 4x3 repeated measures design. We randomly generated five mazes for each of the 12 experimental conditions, varying the number of obstacles placed onto the grid (100, 120, 140, 160, respectively), and the length of the minimum path (14, 24, and 34 steps). The complete specification of each maze was recorded and used for the analysis A maze for each condition was randomly selected to form the practice pool, while the remaining four entered into the test pool.

The dependent variables were: (a) the number of errors (i.e., the difference between the number of moves and the minimum path length), and (b) the total time to reach the goal.

Two other variables (i.e., the Euclidean and the Manhattan distance from the starting position to the target) were recorded to allow statistical control while several other variables had their values computed from the execution traces. Among these we mention: (a) time necessary to execute the first move, (b) mean inter-move latency (excluding the first move), (c) number of direction changes (turns) in the solution path, (d) number of moves in the path leading further from the goal.

For the experiment a PowerMacintosh 9500 computer was used. Participants interacted with the computer by using the arrow keys of the keyboard. They were instructed to use the right hand and to press the "up" and the "down" keys with the middle finger, the "forward" (i.e., right arrow) key with the ring finger and the "backward" (i.e., left arrow) with the index finger. Every key press caused a one-step movement of the token

on the grid. No "repeat" effect was associated with keeping a key down. All key presses and related times were recorded to produce a complete and detailed account of the participants' problem-solving behavior.

At the beginning of the experimental session the participants were asked to read the instructions which explained the experimental modalities and highlighted the task they had to perform: to reach the target cell in the shortest possible time by moving the token on the computer screen. The participants carried out a practice round with the twelve mazes of the practice pool. The mazes were presented in a different random order for each participant. After completing a maze, the participants were reminded (both by a message on the screen and by a sound) to press a key to start the next trial. Between the sixth and the seventh trial they took a one-minute break, timed by the computer. During the practice phase an experimenter was present and willing to answer possible questions. After the practice phase, the participants were left alone to solve the 48-mazes of the experiment. The mazes were delivered, utilizing block randomization, in four batches of 12 trials, with a pause of one minute between batches.

Results

We conducted an outlier analysis on both the participants and the mazes. While no outlier was found among the participants, the evaluation detected four particularly difficult mazes which caused exceedingly high values on errors or times. By using discriminant analysis and classification trees, we found that these mazes differed from the remaining ones by requiring the participants to perform a significantly higher number of turns to reach the goal. Considering their peculiar nature, we excluded them from further analyses.

Table 1 reports the mean number of errors and the time to reach the target for the mazes in the different experimental conditions. In order to analyze them, the data were transformed by taking the square root of the number of errors and the logarithm of the times.

| Condition | | Dependent | Variables |
|-----------|-------------|-----------|-------------|
| Obstacles | Path Length | Errors | Time (in s) |
| 100 | 14 | 1.30 | 4.88 |
| 100 | 24 | 4.10 | 9.69 |
| 100 | 34 | 3.32 | 11.68 |
| 120 | 14 | 2.52 | 5.76 |
| 120 | 24 | 5.28 | 10.40 |
| 120 | 34 | 6.10 | 14.80 |
| 140 | 14 | 1.42 | 5.00 |
| 140 | 24 | 4.32 | 9.30 |
| 140 | 34 | 5.75 | 14.80 |
| 160 | 14 | 2.00 | 5.60 |
| 160 | 24 | 2.95 | 8.40 |
| 160 | 34 | 6.90 | 14.80 |

Table1: Mean Times and Errors for Experimental Conditions

Errors

A two-way repeated measures ANOVA showed a significant interaction between path length and number of obstacles (F(6,138)=4.84, MSE = 0.25, p < .001). The Tukey HSD test highlighted many significant differences at p < .05 level.

In particular, to explain the interaction, the difference between the 160-34 and 160-24 conditions should be taken into account. The 24- and 34-steps path length conditions showed similar error rates for 100, 120 and 140 obstacles. With 160 obstacles, the condition with the longest path (160-34) resulted significantly more difficult. The 14-steps conditions were almost always significantly easier than those with a longer path and the same number of obstacles, the only not significant difference being that between the 160-14 and 160-24 conditions.

Also the main effects of the number of obstacles (F(3, 69) = 9.16, MSE=0.26, p < .001) and the path length (F(2, 46)=75.33, MSE=0.32, p < .001) were significant. Because the number of turns constitutes an important intervening variable that is strongly related to the number of errors and to the path length, but only weakly related to the number of obstacles, we used it as a changing covariate in a two-way repeated measures ANCOVA. The ANCOVA showed only the main effect of the path length (F(2,44)=7.19, MSE = 0.13, p < .01).

Time

A two-way repeated measures ANOVA showed a significant interaction between path length and number of obstacles (F(6, 138) = 15.24, MSE = 0.01, p < .001).

The Tukey HSD test highlighted many significant differences at the p < .05 level. For each obstacle level, the time required to solve the maze significantly increased with the path length. By taking the path length into account, it was found that in case of the longest (34 steps) path, the 100 obstacle mazes resulted easier than those in the other conditions (which did not differ from each other). Significant differences, of smaller magnitude, were also found for the other levels of path length. The main effects of the number of obstacles (*F* (3, 69)= 33.70, *MSE* = 0.01, *p* <.001) and the path length (*F* (2, 46)= 3209.44, *MSE* = 0.01, *p* <.001) were also significant.

Discussion

By considering the general pattern of the results, we can draw the following conclusions. The independent variables of path length and obstacle number had a significant effect on the dependent variables: time to reach the solutions, and errors (i.e. moves in excess of the optimum path). Unsurprisingly, mazes with short paths are quicker and simpler to solve; mazes with longer paths and a higher number of obstacles appear more difficult, requiring a longer time and yielding a greater number of errors. The nature of the interactions we have found suggests, however, that the impact of the structural features of the mazes is not direct but mediated by a complex problem-solving behavior. Moreover, some intervening variables seem to play a critical role. For example, as showed by the ANCOVA, taking into account the effect of the number of turns performed in reaching the goal completely rules out the main effect of the obstacles, and the interaction between obstacles and path length. For these reasons, we found necessary to get a more accurate picture of the factors contributing to the participants' performance.

The Simulation Model

To give a detailed account of the participants' step-by-step behavior we developed a simulative model that rests on the assumptions that, in devising a solution path for the maze, the participants: (a) plan incrementally according to a greedy, locally optimized approach similar to the "hill climbing" strategy, and (b) try to minimize the number of turns in the path.

The first assumption means that the participants do not create the complete plan at once but follow an incremental subgoaling strategy choosing, at every step, the next point (goal) to be reached. The goals coincide with the turning points and, in selecting them, the participants try to minimize the distance that separates the choosen goal from the final target. A goal capable of reducing such a length is considered to be a "privileged" goal and, being situated along one of the natural directions, will be reached without performing any "contrary" move. Should no privileged goal be found, the participants will choose as the next point to be reached an "alternate" goal, i.e. a point that will increase the distance from the target of the shortest possible amount. Every alternate goal could be reached by performing a single contrary move.

According to the second assumption, the participants will fix their goals along a direction being as free as possible from obstacles in order to make the greatest number of identical moves, i.e. to maximize the distance which is possible to cover without changing direction. Forced to choose among equivalent goals, the participants will obey an inertia principle and will continue to follow the direction they were moving along.

These assumptions have been translated into a Common Lisp program that simulates the planning processes involved in devising a solution path for the mazes. The core of the model is constituted by a set of procedures that incrementally find a goal until the target has been reached. Whenever a goal has been fixed, the model prescribes the moves necessary to reach it. The planning process is performed by trying sequentially the following procedures, until one succeeds: found-privileged-goal, found-alternate-goal, or backtrack.

To find a privileged goal, the following steps are performed:

1. Find every possible goal along the natural directions and choose the best, i.e. the one that minimizes the distance from the target.

2. If no goal has been identified, return with failure.

3. If a unique privileged goal has been identified, execute the steps necessary to reach it.

4. If several equivalent goals have been found, choose randomly among them.

In case no privileged goal can be found (because no such a goal exists or the point representing the goal has already been visited) an alternate goal is tried.

To find an alternate goal:

1. Find every possible alternate direction, i.e. any move that leads away from the target.

2. Unless such a direction has been identified, backtrack to a previous choice point.

3. If a unique alternate direction exists, make a step along that direction.

4. If more alternate directions exist and the current one is among them, make a step in the current direction, otherwise choose randomly a direction and make a step along it.

To backtrack:

1. Pop the current goal.

2. Mark it as visited.

3. Plan again from the previous goal (avoiding to choose an already visited one).

The Model Tracing Procedure

We compared the performance of the model with the behavior of all the participants on half (i.e., 24) randomly selected test mazes (two for each experimental condition). We selected the model-tracing methodology (Anderson, 1993) as the best method to get a fair comparison for our sequential tasks. This method requires to compare at each step the prediction of the model with the behavior of one of the participants, and to possibly reset (if a mismatch has been found) the former to the latter.

In particular, at every step, we identified the goal chosen (or pursued, whenever one already existed) by the participant, and compared it with that predicted by the model. We recorded the moves carried out by the participant and by the model and, whenever it was needed, we reset the position occupied by the model token to the position reached by the participant. Since the model fixes its subgoals only at turning points directly reacheable from the current position, in some cases of position resetting the model was forced to set an intermediate subgoal at a turning point to be able to reach a participant's goal situated beyond it. In these special cases the mismatch was recorded and reported in the results.

To establish, from the execution traces, the goals set by the participants we looked at the points where a significant execution slow down occurred. These are the places where the cognitive load does not allow both online planning and execution to take place without additional cost (i.e., more processing time). To identify these points we computed, in a preliminary experiment with 10 participants solving simpler mazes that did not require any planning or subgoaling, independent estimates of the execution times, i.e., the time required to execute the different moves. Three different patterns were distinguished: (a) one for performing straight moves, (i.e. a move in the same direction of the preceding one), (b) one for performing an up-right turn, and (c) one for performing the other kind of direction changes. We computed these estimates obtaining the following means (and 95% confidence intervals, henceforth: c.i.):

- 1. straight moves 158 ms. (c.i. 148-168);
- 2. up-right turns: 268 ms. (c.i. 199-338);
- 3. other turns: 282 ms. (c.i. 246-317).

We looked then at execution trace selecting the inter-move latencies with a value higher than the upper confidence limit for the specific kind of move (i.e., 168 for straight moves, 338 for up-rights, and 317 for other turns). These points were identified as the goals choosen by the participants, and utilized in the model-tracing procedure.

The tracing procedure took into account not only the goals and the moves needed to reach them, but also the time spent to to solve the problem. To make the model predict the time for every single move, we used the execution times identified in the preliminary experiment (i.e. 158 ms. for straight moves, 268 for up-right turns and 282 for other turns).

To estimate the time needed to choose a given goal, we utilized the mazes not used for model tracing (two for each experimental condition). For each maze and each subject, we identified the points with a significant processing load (see above) and computed their processing time (i.e. the difference between the inter-move latency and the time needed to execute the move). We averaged over subjects and the mazes, and the result (172 ms.) was used as an estimate for the goal selection time.

Finally, we computed the mean time (1047 ms) necessary to execute the first move, i.e, the time necessary to visually code the maze.

Simulation Results

For clarity's sake, we summarize the results averaging over the test mazes. All confidence intervals are at the 95% level.

Moves

The mean agreement between the moves of the participants and the model was .82 (c.i.: .63 and .93). Given the complexity of the task and the large number of potential choice points, this agreement is quite high. In the 81% of cases the subjects followed a "hill-climbing" strategy and executed a move bringing them nearer to the goal, giving support to the first model assumption.

Goals

The use of model-tracing methodology allows an evaluation of the agreement between the goals chosen by the model and by the participants. The location was exactly the same for 25% of the goals (c.i: .12 - .45). In the 66% of the cases in which the location differed (c.i.: .47 - .81), the mean Manhattan distance between the two goals was very low (2.2 steps, c.i.: 1.9 - 2.4) and remained constant for all the experimental conditions. It should be remarked that in two thirds of the cases of choice mismatch, the model's and the participant's goals were situated along the same direction. Finally, in 8% of cases (c.i: .02 - .25) the model was forced to set an intermediate subgoal at a turning point in order to be able to reach the goal set by the participant.

These results show a good overall agreement between the times in which the model and the participants set a goal, and between the spatial positions of the goals. The small constant difference between the positions may be related to motor and monitoring factors.

As an independent support for our second assumption, we run on the test mazes a modified version of the A* algorithm to perform a local minimization in the number of turns. The modified algorithm is able to to find the shortest path to the goal mantaining, whenever possible, the direction already taken. We ran the algorithm for 1000 cycles and computed the mean number of turns for each maze. The comparison with the number of turns performed by the participants yielded no significant difference (t(23)= -.30, p =.77). A linear regression showed a good fit between observed and predicted turns (R^2_{adj} =.743, F(1,22)=67.63, p.<.001). These data provide good evidence for the local mimimization assumption showing, at the same time, the effectiveness of the participants' performance.

Times

The mean correlation between the participants' and model's inter-move latencies is 0.74 (N = 24, p < .001). This is a promising result, considering two main simplifications embedded in the model that lower the correlation. First, our model does not separate the cost of replanning in extremely critical situations (such as long dead-ends, sometimes present in our data) from the cost of ordinary subgoaling. Second, the time for the initial encoding of each maze is kept constant and does not depend on the specific features of the perceptual input (e.g., start and target locations, number and position of the obstacles).

Conclusions

Two important conclusions can be drawn from the results of the experiment and from the simulation. The first concerns the evidence for an incremental subgoaling strategy. Even if

the participants could always see the whole maze, they did not plan a complete path but seemed to follow a step by step approach guided by hill-climbing. This strategy shares many aspects with other well-known heuristics used in problem solving (Anderson, 1993), and with some features of a recently developed model of animal route finding (Reid & Staddon, 1998). The strategy allows the participants to find, almost always, a good and quick solution, and that is exactly what they were required to do. They seem therefore to adopt a rational behavior that can be explained in the context of the adaptive theories of cognition (Anderson, 1990; Payne, Bettman & Johnson, 1993). The second important result is the empirical support for the role of the operators cost (O'Hara & Payne, 1998) on strategy selection: participants tend to locally minimize the number of turns probably because turning moves have a high temporal cost.

Our future work will take two directions: (a) ameliorating the simulative model, and (b) making further experimental research. The main improvement to the model will consist in the inclusion of replanning process in critical situations where hill-climbing fails. Following our theoretical interests, we are also going to translate the model into the ACT-R/PM framework (Anderson & Lebiere, 1998). Further direct experimentation will be carried out for testing specific predictions of the model (e.g., manipulating the minimum number of turns allowed by a maze, and providing external aids).

References

Anderson, J.R. (1990). The adaptive character of thought. Hillsdale, NJ: Erlbaum.

- Anderson, J.R. (1993). Rules of the mind. Hillsdale, NJ: Erlbaum.
- Anderson, J.R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Crowe, S.M., Barclay, L., Brennan, S., Farkas, L., Gould, E., Katchmarsky, S., & Vayda, S. (1999). The cognitive determinants of performance on the Austin maze. *Journal* of the International Neuropsychological Society, 5, 1-9.
- Flitman, S., O'Grady, J., Cooper, V., & Grafman, J. (1997). PET imaging of maze processing. *Neuropsychologia*, *35*, 409-420.
- Glassman, R.B., Leniek, K.M., & Haegerich, T.M. (1998). Human working memory capacity is 7 ± 2 in a radial maze with distracting interruption. *Brain Research Bulletin*, 47, 249-256.
- Hickman, K.L., Stromme, C., & Lippman, L.G. (1998). Learned industriousness: Replication in principle. *Journal of General Psychology*, *125*, 213-217.
- Krikorian, R., & Bartok, J.A. (1998). Developmental data for the Porteus maze test.. *Clinical Neuropsychologist, 12,* 305-310.
- O'Hara, K.P., & Payne, S.P. (1998). The effects of operator implementation cost on planflulness of problem solving and learning. *Cognitive Psychology*, *35*, 34-70.
- Payne, J.W., Bettman J.R., & Johnson, E.J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- Porteus, S.D. & Kepner, R.D. (1944). Mental changes after bilateral prefrontal lobotomy. *Genetic Psychology Monographs*, 29, 3-115.
- Reid, A.K., & Staddon, J.E.R. (1998). A dynamic route finder for the cognitive map. *Psychological Review*, *105*, 585-601.
- Stephenson, B., & Wicklund, R.A. (1984). The contagion of self-focus within a dyad. *Journal of Personality and Social Psychology*, 46, 163-168.