Finding your Way Out: Planning Strategies in Human Maze-Solving Behavior

Florian Kadner[1] **Hannah Willkomm**[2] **Inga Ibs**[1] **Constantin A. Rothkopf**[1]

Centre for Cognitive Science & Institute of Psychology, TU Darmstadt

[1]{firstname.lastname}@tu-darmstadt.de

[2]{firstname.lastname}@stud.tu-darmstadt.de

Abstract

In many situations encountered in our daily lives where we have several options to choose from, we need to balance the amount of planning into the future with the number of alternatives we want to consider to achieve our long-term goals. A popular way to study behavior in these planning problems in controlled environments are maze-solving tasks since they can be precisely defined and controlled in terms of their topology. In our study, participants solved mazes that differed systematically in topological properties regulating the number of alternatives and depth of paths. Replicating previous results, we show the influence of these spatial features on performance and stopping times. Longer and more branched solution paths lead to more planning effort and longer solution times. Additionally, we measured subjects' eye movements to investigate their planning horizon. Our results suggest that people decrease their planning depth with increasing number of alternatives.

Keywords: planning; maze solving; eye movements

Introduction

In order to achieve our goals and navigate the complexities of daily life, we need to plan our actions. In doing so, tasks can turn out to be difficult, as they offer several alternative courses of action or because the consequences of our actions have to be considered over a longer period of time. One way to study this planning behaviour in controlled environments is via maze-solving tasks. In contrast to simple binary decision tasks, maze-solving involves sequential multi-step decision-making, which requires planning in order to achieve long term goals. Therefore, for analyzing decision-making and its underlying planning mechanisms, mazes have been of particular popularity in studying humans, animals and machines in various fields including Cognitive Science (Kryven, Kleiman-Weiner, Tenenbaum, & Yu, 2022; Wu et al., 2016; Buecher, Hölscher, & Wiener, 2009), Neuroscience (Alonso, van der Meij, Tse, & Genzel, 2020; Rosenberg, Zhang, Perona, & Meister, 2021), and Robotics (Dang, Song, & Guo, 2010; Aqel et al., 2017). Despite this broad interest in planning behavior in mazes, the analyses, modeling, and respective literature have been somewhat disconnected, reflecting the potentially relatively arbitrary conceptualizations ranging from sensorimotor planning to planning as higher cognition.

Real world naturalistic navigation tasks usually require the integration of internal and external cues, the execution of motor actions, and internal planning (Kessler, Frankenstein, & Rothkopf, 2022). However, one major advantage of mazes as experimental environments is that they can be clearly defined and generated in terms of their topology (Kim & Crawfis, 2018). Thus, we can use mazes to investigate planning mechanisms during navigation and are able to control the spatial structure of the environment. Elements of a maze's topology include for example, specific cell types (dead-ends, turns, crossings) and their overall distribution, but also the spatial arrangement of cells or the length of the solution path. In recent years, the possibilities for automated generation of mazes given various hyperparameters have been investigated and constantly developed further (Bellot et al., 2021; Kim & Crawfis, 2015). In human maze-solving, the underlying topological characteristics have been proposed to influence behaviour in terms of maze solving time: Solution path length and the number of turns along the solution path have been shown to render a maze more complex and therefore increase solving time (Crowe, Averbeck, Chafee, Anderson, & Georgopoulos, 2000). The number of alternatives influences the exploration behavior, where participants examine the task-relevant structure of the environment more thoroughly with an increasing amount of alternatives (Zhu, Lakshminarasimhan, Arfaei, & Angelaki, 2022).

Since planning and the underlying internal processes are not readily observable while decisions are made, eye movements have been used successfully as indication of ongoing cognitive processes (Spering, 2022; König et al., 2016; Hayhoe & Ballard, 2005), particularly in tasks in which spatial locations allow reducing uncertainty about task relevant quantities (Kaplan & Friston, 2018; Zhu et al., 2022). Furthermore, eye movements have been proposed to be closely related to the planning horizon since they can be understood as information sampling in the visual environment (Ma, Ma, & Gureckis, 2021) and have indeed been shown to be planned ahead (Hoppe & Rothkopf, 2019). Furthermore, the planning horizon and strategy can differ dynamically within a task and between subjects (Tsividis et al., 2021; Carton, Nitsch, Meinzer, & Wollherr, 2016) not the least by the simple fact that human scan paths are not independent of individual behavioral preferences in gaze selection (De Haas, Iakovidis, Schwarzkopf, & Gegenfurtner, 2019; Kadner, Thomas, Hoppe, & Rothkopf, 2023). Recent work suggests that humans balance depth and breadth searches (Vidal, Soto-Faraco, & Moreno-Bote, 2022) and prune decision trees related to their plans when encountering large losses (Huys et

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

al., 2012) in sequential decision making tasks different from mazes. However, whether and, if so, how humans potentially balance different strategies such as depth and breadth searches based on the availability of alternatives and the depth in search trees, is not known.

Eye movements have been investigated in previous studies involving mazes to gain insight into human maze-solving strategies, particularly related to planning. Previous studies suggest that gaze reflects a mental simulation process during maze-solving and is, therefore, reflective of the maze and its solution path structure (Li, Watters, Yingting, Sohn, & Jazayeri, 2022; Zhu et al., 2022; Crowe et al., 2000). (Zhao & Marquez, 2013) found that gaze patterns during maze-solving can be differentiated into those that subserve exploration and those that aid in motor guidance. Although these studies indicate that gaze patterns represent planning behaviour in mazes, the exact influence of topological features on human planning strategy and the adopted planning horizon remains unknown.

In this study, we parametrically generated different mazes by controlling topological parameters, which influence the number of alternative paths and the length of the solution path. We analyzed participants' behavior by converting mazes into equivalent decision trees, allowing us to compute principled features quantifying the topology. First, we confirm previous results showing that both the length of the solution path and the number of possible alternative routes impact performance, i.e. search time. Secondly, we look at the influence of the overall topology and the influence at the level of individual cells on solving time. Finally, by measuring subjects' eye movements, we can quantify participants' planning strategy by inferring depth and breadth features of their visual search. We find an effect of the number of alternate paths in a maze on the depth of their planning. Thus, with a larger number of alternatives, they plan less deeply, but keep the width of their planning constant, which hints at an adaptive planning strategy that adjusts the depth of planning with the number of paths that have to be considered. Such a strategy is computationally adequate, because memory resources for storing paths that have to be evaluated is limited.

Methods

Participants

16 subjects (9 female, 7 male; age M = 22, SD = 2.39) participated in the experiment. For seven subjects additional eye tracking data was recorded. All of the eye tracking subjects had normal or corrected to normal vision. All experimental procedures were approved by the ethics committee of the Technical University of Darmstadt and informed consent was obtained from all participants.

Apparatus

Mazes were presented on a 2560×1440 (559×335 mm) monitor. Participants were seated approximately 102 cm from the monitor using a chin rest. The mazes were presented centrally on an area of 1200×1200 pixels, such that the stimuli were displayed at a visual angle of $\sim 15° \times \sim 16°$. Eye movement data were collected using an Eyelink 1000 Plus eye tracker with a 35 mm lens, allowing online event parsing. We recorded the data from the participant's dominant eye determined before the start of the experiment. We performed a 9-point array calibration and validation procedure for each participant prior to the start of the experiment to ensure the accuracy of eye tracking data. The average validation accuracy was 0.3°, with all individual point measurements being under 0.94°.

Experimental Design

Maze Generation We used the search-based procedural content generation (SBPCG) approach (Kim & Crawfis, 2018) constrained to solution path length and the number of dead ends to generate mazes in the size of 25×25 cells varying in their difficulty in terms of number of alternative paths and depths. This was done to elicit different human exploration patterns. All the generated mazes were *perfect*, so there is exactly one correct solution path that leads from the starting point to the goal. We opted for a 3×3 experimental design, choosing three different solution path lengths (short, medium, long) and three different levels of number of junctions (low, medium, high). The solution path lengths and their classification into the three levels are motivated by previous

Figure 1: The nine mazes used in the experiment. The columns contain different generation algorithms, which correlate with the number of junctions in the maze. In the rows, the length of the solution path increases from top to bottom. The starting point is shown as a grey point, the yellow star visualises the goal, and green dots visualise the solution path.

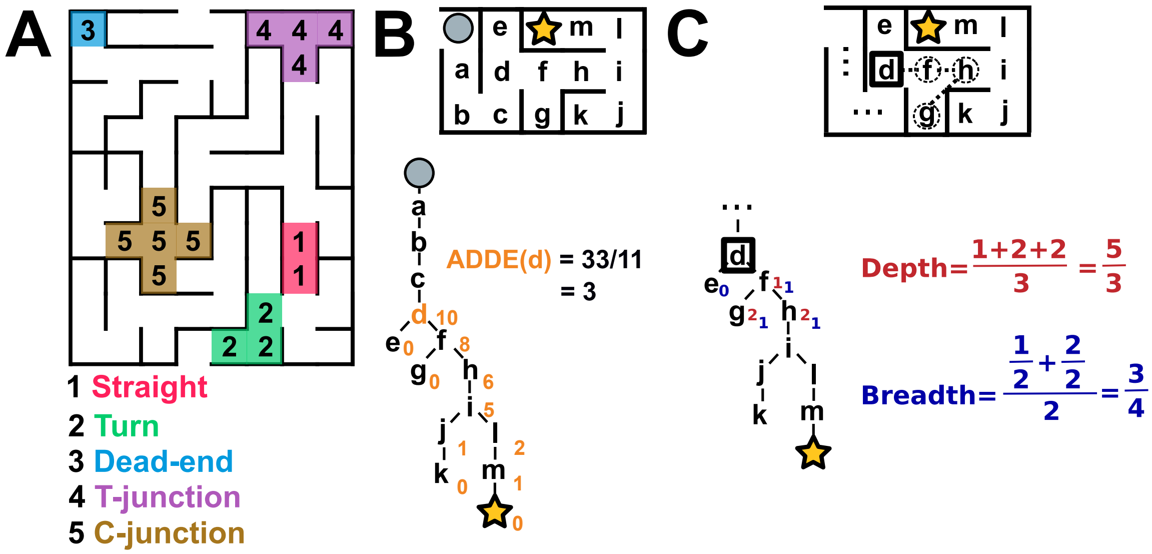

Figure 2: Topological features of mazes. (A) Visualisation of the cell types (1) straights (2) turns (3) dead-ends (4) T-junctions, and (5) C-junctions. (B) The internal representation of a maze in a tree structure. For node d, the calculation of the Average Distance to Dead-Ends is exemplarily shown. (C) Exemplary calculation of depth and breadth. The player is in cell d and then fixates on cells f, h and g.

studies that observed variation in solution times when manipulating these variables (Crowe et al., 2000). We ensured the different numbers of junctions by using different generation algorithms, which are known to generate different amounts of dead ends correlating with the number of junctions: Recursive Backtracking for low, Hunt-and-Kill for medium and Kruskal for a high amount of dead ends. Figure 1 shows the resulting nine mazes used and the respective subdivision given by solution path length and the number of junctions.

Procedure Subjects were asked to solve the mazes as quickly as possible. To navigate through the maze and reach the goal, they could move from cell to cell using the arrow keys on a regular keyboard. The player position was indicated via a blue dot. Before the first move, the player position was placed at the start position. The goal was indicated via a red cross. After the first movement in the maze, the start position disappeared. Cells already visited were not marked, so participants needed to keep their path in memory in order to be able to trace back their steps. In order to exclude potential biases simply stemming from the geometric orientation of the solution path, each maze was shown twice, once as shown in Figure 1 with the solution path aligned in the horizontal direction and, a second time, rotated 90 degrees to the left.

In addition to the nine mazes, an easy-to-solve test maze was generated using the Recursive Division algorithm (Reynolds, 2010). It was shown at the beginning of the experiment to familiarise participants with the control and task mechanics of the experiments. Afterward, the 9x2 experimental mazes were presented in random order. Finally, the rotated version of the test maze was shown. The test maze was excluded from further analyses.

Metrics of topology

Following (Kim & Crawfis, 2018), we define a set of metrics describing the global topology of a maze. First, this includes the different cell types that can occur within the maze, which are visualised in Figure 2A. The simplest cell types are dead-ends, which allow the player only to move backward, followed by straights and turns, which allow the player to move forward or backward. More complex cell types involve a higher number of alternatives, including so-called T-junctions for three alternatives and C-junctions for four, based on their shape. These cell types, together with their frequency and distribution especially along the solution path, build the basis for the topological properties of a labyrinth. Another important property to describe the complexity of the mazes is the length and branching of individual paths until they reach a dead end. In order to quantify this property, we introduce the Average Distance to Dead-Ends (ADDE) measurement, computed for all cells in a maze. To calculate this property, each maze is converted into its equivalent tree structure with its unique identity. Then, starting from one cell, the mean distance to all dead ends (without considering predecessor cells) is calculated. The value thus captures two important properties: First, it increases with the number of alternative paths given the current position since, for example, a C-junction can have one more path running into a dead end than a T-junction. Secondly, it increases with longer sub-trees after the given position, which makes it more difficult for the subjects to look at the entire set of paths, plan and remember the findings. The tree representation of a maze and the exemplary calculation of the ADDE score for a given cell can be seen in Figure 2B.

Measuring planning behavior

We use the subjects' eye movements to gain insight into their planning behaviour. To do this, we look at the eye movements while the player stands still at a location in the maze and the participant explores the maze with their gaze. For the planning process, we then compute two features of the search carried out by the participant's eyes, specifically a feature quantifying the depth of a search and a second feature quantifying the breath of a search. The depth of a search episode is calculated as the average distance of the fixed cells to the current player position. The breadth is calculated by averaging the proportion of covered cells per depth level up to the level of the deepest fixation. Figure 2C shows an exemplary calculation of these two measures. The depth measure for the player cell d is calculated by adding up the distance of the fixated cells f,g,h to d (1,2,2 respectively) and dividing by the number of fixated cells (here 3) resulting in a depth of $5/3$. For the player cell d the breadth measure is calculated as follows: On the first level only one of two possible cell (f) is covered by a fixation, which results in a coverage of $1/2$. On the next and last level both possible cells are fixated (g and h) resulting in a coverage of 1. This gives us an average coverage for the two depth levels up to the deepest fixations (g and h) of $3/4$ as our breadth measure.

Results

Ruling out Confounders

Comparing the average solution times of the two types of rotations showed no major effect (two-sample t-test, t=0.261, p=0.795). In addition, the subjects were presented with the mazes (except for the first and last maze) in a random order. We performed a Linear Regression where the position of the maze within the experiment was compared to the corresponding solution time. The order of presentation had no significant effect on subjects' performance ($F(1,16)=0.125$, $\beta=-0.49$, $p=0.072$).

Performance

In order to quantify the performance of individual subjects, we consider the total time they took to solve the mazes. The mean solving time for the maze was 89.39 seconds ($\sigma_{\text{mazes}} = 59.17$). The descriptive statistics and differences for all nine mazes are shown in Table 1. On average, the fastest subject needed 41.77 seconds to solve one maze and the slowest 154.75 seconds ($\sigma_{\text{subjects}} = 54.76$).

Table 1: Descriptive statistics for all nine mazes.

	Mean	Std	Min	Max
Maze 1	54.64	21.98	24.33	118.19
Maze 2	50.12	33.15	23.30	183.75
Maze 3	85.94	55.83	25.27	238.81
Maze 4	71.58	51.73	20.39	254.83
Maze 5	95.67	77.37	29.35	347.82
Maze 6	92.09	60.12	33.04	299.62
Maze 7	96.54	62.25	35.92	284.53
Maze 8	129.59	83.61	45.84	481.61
Maze 9	128.31	86.50	54.28	480.42

Why planning?

To investigate subjects' performance, given topological constraints at the global maze level, we used a Linear Mixed Effects Model. We chose the subjects as random effects, the solution path length and the number of junctions per path unit on the solution path as fixed effects. Increasing the solution path length should also increase the complexity of the maze since a larger planning horizon is needed to traverse down the path and rule alternatives out on the way. The number of junctions should also correlate with a higher complexity due to a higher number of alternatives that must be considered. The number of junctions is defined as alternative paths from the solution path, where C-junctions add two alternatives and T-junction one choice. We choose random slopes and intercepts for each subject to incorporate individual differences in personal performance. Equation 1 shows the resulting model for all the mazes.

$$\begin{aligned}\text{solving time} \sim (1|\text{subject}) + \text{length} + (\text{length} - 1|\text{subject})\\ + \text{junctions} + (\text{junctions} - 1|\text{subject}) \qquad (1)\end{aligned}$$

The relative effects are shown in Figure 3. Our results suggest that the solution path length and the density of junctions on the solution path make the mazes harder to solve for all participants and require increased planning effort.

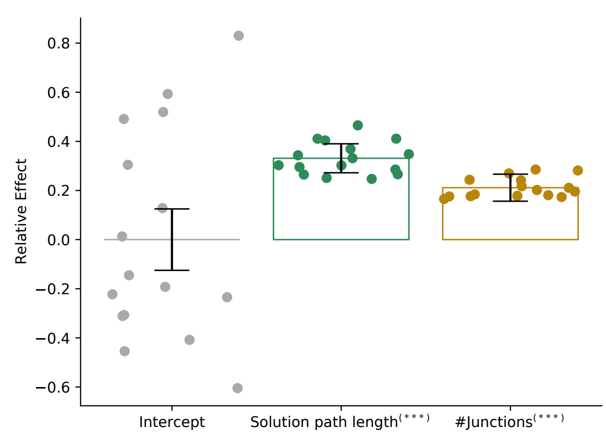

Figure 3: Relative effects for the Linear Mixed Effects Model. Participants were chosen as random effects with random intercepts and slopes. All variables were z-scaled. Errorbars correspond to the standard error of the mean.

Where to plan?

We differentiate between two different components in planning. Besides the regions a player looks at to find the solution path and the cognitive processes underlying this search, it is first important to know at which points in the maze this process started. Therefore, we first look for points at which the player stops moving and switches to exploring the maze with their gaze. For this purpose, we calculated a linear regression over all subjects, looking for cell properties favoring a prolonged stay and, thus, a prolonged planning time. In addition to the cell types (i.e. baseline straight/turn with only two alternative paths), we also considered their position on the solution path (i.e. the distance as number of cells to the target), their ADDE score, and the interaction effect of these. The fitted coefficients of the model ($F(5,1075)=157.8, p < 0.005$) are shown in Figure 4.

We found longer waiting times the more alternative paths the current cell had (note that here the 2-alternative cell types straights/turns were taken as the baseline for the categorical variables). We also found a significant influence of the current position on the solution path. The further away the player is from the goal point, the longer he stays. This is consistent with our hypothesis that more planning and exploration is necessary at the beginning to find a possible solution path. The ADDE score alone has no significant influence but the interaction between the ADDE and the distance to the target does. This can be explained by the fact that paths that are longer and have an increased branching have hardly any influence on planning shortly before reaching the goal since the solution path has already been found with high probability. However, this has a strong influence right at the beginning of solving a maze since possible solution path candidates have to be explored much longer during initial planning to see whether they are promising.

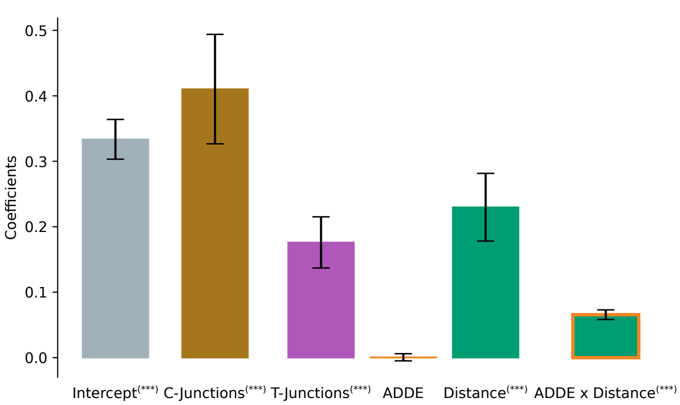

Figure 4: Coefficients for the linear regression to explain the times spent on cells given their specific properties. The reference variable for the cell type is turns/straights with only two alternative decisions. Errorbars correspond to the standard error of the mean.

How to plan?

To investigate people's planning strategy, we look at the breadth and depth values for their fixations paths as described and shown in Figure 2C. The calculated mean values for breadth and width for all participants in all nine mazes are

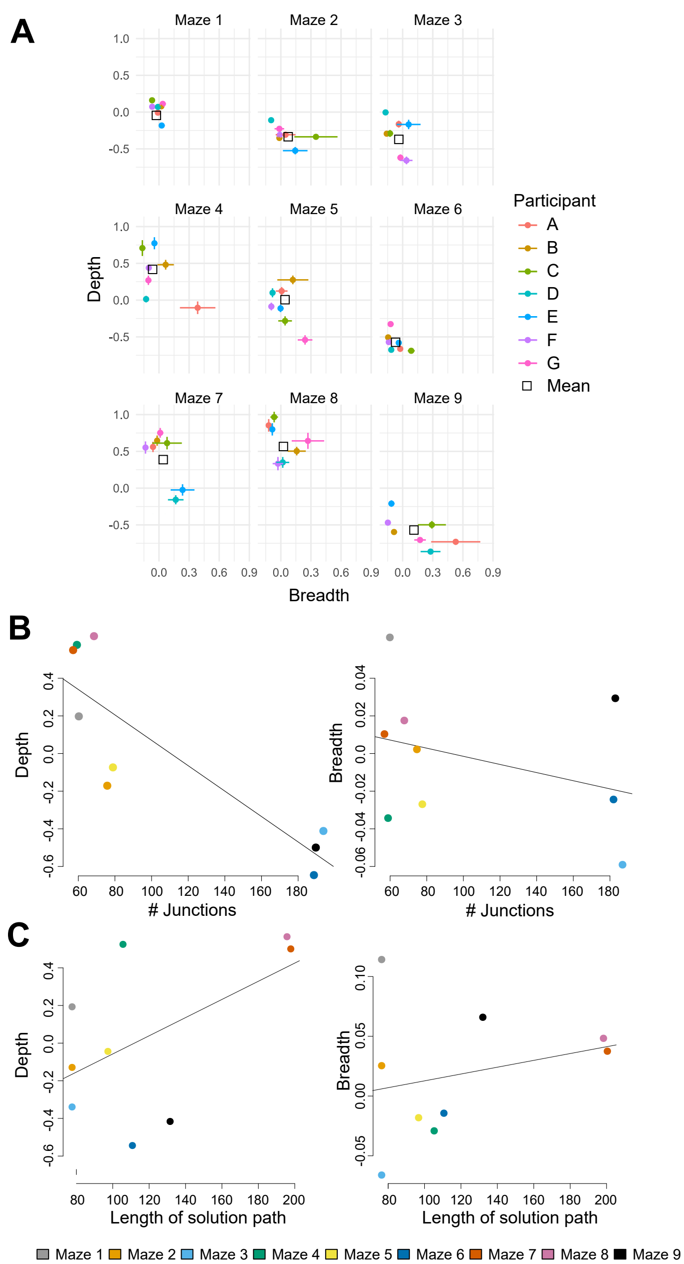

Figure 5: Estimated breadth and depth values (z-scaled). (A) Mean breadth and width values across all participants for the nine mazes. Errorbars correspond to the standard error of the mean. All values are z-scaled. (B) Relationship between the number of junctions in a maze and the breadth and depth search values of participants. (C) Relationship between the length of the solution path in a maze and the participants' depth and breadth search values.

displayed in Figure 5A. These plots show that subjects used consistent patterns of depth and breadth exploration across different maze types. Despite some variability across participants, their behavior generally followed the two manipulated topological features of the mazes (length of solution path and number of junctions) in the same way.

By construction, the mazes were generated with different characteristics in the dimensions *number of junctions* and *length of the solution path* (see Figure 1). To investigate the influences of these two features on the depth and breadth values, we calculated linear regressions. Figure 5B visualises the influence of the number of junctions on breadth and depth. The plot demonstrates a significant impact of the amount of junctions in the maze on the depth of the search (linear regression, $R^2 = 0.74, p = 0.003$) but not on breadth value (linear regression, $R^2 = 0.12, p = 0.36$), i.e. with increasing number of junctions, participants maintained the breadth of their search, but reduced its depth. For the influence of the length of the solution paths (Figure 5C) on breadth and depth, we see no influence, neither on the depth (linear regression, $R^2 = 0.28, p = 0.14$) nor the breadth values (linear regression, $R^2 = 0.06, p = 0.51$). Due to a longer solution path, the paths are less deep overall. People cannot anticipate these circumstances, though, because they cannot possibly know the length of the solution path in advance and thus don't know the depth of the alternatives. So, within the limits of their capacities, they keep both breadth and depth constant here.

Discussion & Outlook

In this study, we investigated human planning strategies in maze-solving tasks. We generated different mazes according to fixed topological parameters and looked at both the performance and the planning behaviour by means of exploration times and measured eye movements of the subjects. We were able to reproduce and extend previously known results that suggested that performance decreases as mazes become more complex (that is, a higher amount and deeper branches within). At the cell level, we systematically found those locations where subjects needed more time to explore and plan their next steps. These cells were easily detectable, because participants spent significantly more time at these points in the maze without moving while gaze was moving along alternative paths.

To investigate the underlying planning process, we transformed the mazes into their equivalent decision trees to quantify the number of available alternative decisions at each point with the introduced ADDE measurement. The results show a strong effect of junctions on the exploration time. However, the depth of the possible search tree mainly had an effect at the beginning of the solution path, where subjects had to explore more extensively than towards the end when a solution path had been found with high probability. Evaluating the participant's eye movements within the decision tree, we were able to assign their internal planning to breadth or depth seeking. Subjects tended towards breadth planning, which suggests that they may be more likely to adopt a strategy that allows them to explore multiple options before committing to a specific solution. However, they are able to balance their strategy and change to more depth planning with decreasing number of junctions in a maze, e.g., adapting their strategy when they encounter situations with limited options. One explanation for this result is that the more alternative paths present themselves at junctions, i.e. the larger the branching in the equivalent search tree, the larger the memory requirement for evaluating all the different paths. Thus, by exploring each path to shallower degree, the burden on memory is reduced. These findings are consistent with (Vidal et al., 2022), who designed the *BD apricot task*, an economic many-alternative task where subjects were asked to allocate finite search capacity to sample the reward of the alternatives with the goal to choose the best one. Their results suggested that participants preferred deeper searches in environments where good outcomes where more likely. This is comparable to the increasing search depth of participants observed at maze locations with less alternative ways, i.e. deeper exploration of one of the few paths could lead to better results.

While the results of the current study indicate the utility of topological features in analyzing human planning strategies in maze solving tasks, there are a number of limitations which should be addressed in future studies. The introduced depth and breadth measures give a good indication of the quality of human planning behaviour, however, computational modeling of complete gaze sequences could give a more detailed illustration of the planning behaviour deployed. Since the subjects' did not explore each alternative in complete depth, they must use heuristics for pruning alternatives. An example of such a heuristic could be based on the angular direction from the current position to the target location, which could bias towards specific paths and alternatives in the planning process. To investigate the computational efficiency of the deployed trade-off between deeper and broader planning horizons, constrained breadth-first and depth-search searches could be compared with the strategic behaviour of the participants. Finally, we conclude that, although not addressing all possible aspects of real-world navigation, the controlled visual search environment of algorithmically generated mazes allows to investigate how humans adapt their planning behavior based on topological features.

Acknowledgments

We would like to thank Sebastian Bock, Janis Fitzthum, Lilly Reisig and Jonah Weber for their help in programming and conducting the experiment. We acknowledge support by the Excellence Program of the Hessian Ministry of Higher Education, Science, Research and Art through the project WhiteBox in the priority program LOEWE and the cluster project 'The Adaptive Mind', as well as the BMBF project 'Plexplain'. C.R. additionally acknowledges support by the European Research Council (ERC; Consolidator Award 'ACTOR'-project number ERC-CoG-101045783).

References

Alonso, A., van der Meij, J., Tse, D., & Genzel, L. (2020). Naïve to expert: Considering the role of previous knowledge in memory. *Brain and neuroscience advances*, *4*.

Aqel, M. O., Issa, A., Khdair, M., ElHabbash, M., AbuBaker, M., & Massoud, M. (2017). Intelligent maze solving robot based on image processing and graph theory algorithms. In *2017 international conference on promising electronic technologies (icpet)* (pp. 48–53).

Bellot, V., Cautrès, M., Favreau, J.-M., Gonzalez-Thauvin, M., Lafourcade, P., Le Cornec, K., ... Rivière-Wekstein, S. (2021). How to generate perfect mazes? *Information Sciences*, *572*, 444-459. doi: https://doi.org/10.1016/j.ins.2021.03.022

Buecher, S. J., Hölscher, C., & Wiener, J. (2009). Search strategies and their success in a virtual maze. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 31).

Carton, D., Nitsch, V., Meinzer, D., & Wollherr, D. (2016). Towards assessing the human trajectory planning horizon. *PLOS ONE*, *11*(12), e0167021.

Crowe, D. A., Averbeck, B. B., Chafee, M. V., Anderson, J. H., & Georgopoulos, A. P. (2000). Mental maze solving. *Journal of cognitive neuroscience*, *12*(5), 813–827.

Dang, H., Song, J., & Guo, Q. (2010). An efficient algorithm for robot maze-solving. In *2010 second international conference on intelligent human-machine systems and cybernetics* (Vol. 2, pp. 79–82).

De Haas, B., Iakovidis, A. L., Schwarzkopf, D. S., & Gegenfurtner, K. R. (2019). Individual differences in visual salience vary along semantic dimensions. *Proceedings of the National Academy of Sciences*, *116*(24), 11687–11692.

Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in cognitive sciences*, *9*(4), 188–194.

Hoppe, D., & Rothkopf, C. A. (2019). Multi-step planning of eye movements in visual search. *Scientific reports*, *9*(1), 1–12.

Huys, Q. J. M., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012, 03). Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLOS Computational Biology*, *8*(3), 1-13. doi: 10.1371/journal.pcbi.1002410

Kadner, F., Thomas, T., Hoppe, D., & Rothkopf, C. A. (2023, 01). Improving saliency models' predictions of the next fixation with humans' intrinsic cost of gaze shifts. In *Proceedings of the ieee/cvf winter conference on applications of computer vision* (pp. 2104–2114).

Kaplan, R., & Friston, K. J. (2018). Planning and navigation as active inference. *Biological cybernetics*, *112*(4), 323–343.

Kessler, F., Frankenstein, J., & Rothkopf, C. A. (2022). A dynamic bayesian actor model explains endpoint variability in homing tasks. *bioRxiv*. doi: 10.1101/2022.11.09.515854

Kim, P. H., & Crawfis, R. (2015). The quest for the perfect perfect-maze. In *2015 computer games: Ai, animation, mobile, multimedia, educational and serious games (cgames)* (pp. 65–72).

Kim, P. H., & Crawfis, R. (2018). Intelligent maze generation based on topological constraints. In *2018 7th international congress on advanced applied informatics (iiai-aai)* (pp. 867–872).

König, P., Wilming, N., Kietzmann, T. C., Ossandón, J. P., Onat, S., Ehinger, B. V., ... Kaspar, K. (2016). Eye movements as a window to cognitive processes. *Journal of Eye Movement Research*, *9*(5), 1–16.

Kryven, M., Kleiman-Weiner, M., Tenenbaum, J., & Yu, S. (2022). Planning ahead in spatial search. *PsyarXiv preprint*. doi: 10.31234/osf.io/epmxf

Li, J., Watters, N., Yingting, W., Sohn, H., & Jazayeri, M. (2022). *Modeling human eye movements with neural networks in a maze-solving task.* arXiv.

Ma, I., Ma, W. J., & Gureckis, T. M. (2021). Information sampling for contingency planning. In *Proceedings of the 43rd annual conference of the cognitive science society.*

Reynolds, A. (2010). Maze-solving by chemotaxis. *Physical Review E*, *81*(6), 062901.

Rosenberg, M., Zhang, T., Perona, P., & Meister, M. (2021). Mice in a labyrinth show rapid learning, sudden insight, and efficient exploration. *Elife*, *10*, e66175.

Spering, M. (2022). Eye movements as a window into decision-making. *Annual Review of Vision Science*, *8*(1), 427-448. doi: 10.1146/annurev-vision-100720-125029

Tsividis, P. A., Loula, J., Burga, J., Foss, N., Campero, A., Pouncy, T., ... Tenenbaum, J. B. (2021). Human-level reinforcement learning through theory-based modeling, exploration, and planning. *arXiv preprint arXiv:2107.12544*.

Vidal, A., Soto-Faraco, S., & Moreno-Bote, R. (2022, 09). Balance between breadth and depth in human many-alternative decisions. *eLife*, *11*, e76985. doi: 10.7554/eLife.76985

Wu, Z., Zheng, N., Zhang, S., Zheng, X., Gao, L., & Su, L. (2016, 09 13). Maze learning by a hybrid brain-computer system. *Scientific Reports*, *6*(1), 31746.

Zhao, M., & Marquez, A. G. (2013). Understanding humans' strategies in maze solving. *arXiv preprint arXiv:1307.5713*.

Zhu, S., Lakshminarasimhan, K. J., Arfaei, N., & Angelaki, D. E. (2022). Eye movements reveal spatiotemporal dynamics of visually-informed planning in navigation. *Elife*, *11*, e73097.