Using a Cognitive Model for an In-Depth Analysis of the Tower of London

Rebecca Albrecht, Sven Brüssow, Christoph Kaller, and Marco Ragni
albrechr@informatik.uni-freiburg.de; sven@cognition.uni-freiburg.de
christoph.kaller@uniklinik-freiburg.de; ragni@cognition.uni-freiburg.de
Center for Cognitive Science, University of Freiburg
Friedrichstr. 50, D-79098 Freiburg, Germany

Abstract

The Tower of London (ToL) is a transformation task extensively used and well-established as a neuropsychological diagnostic tool for assessing human planning ability in clinical and research contexts. Behavioral experiments have recently shown that planning in the ToL is substantially influenced by structural task parameters. This work presents an ACT-R model of the ToL that explains structural influences by using different strategies, whereby, strategy selection depends on visually observable characteristics. Model evaluation was based on a problem selection that accounted for systematic variations of task demands. Based on comparisons with empirically observed planning latencies from previously published data, we argue that task-specific structural characteristics are necessary to explain human planning strategies.

Introduction

The Tower of London (ToL) task (see Fig. 1a) is a planning task originally proposed by Shallice as a neuropsychological tool to measure planning deficits in patients with frontal lobe damages (Shallice, 1982). Today it is widely used as a general assessment tool to evaluate executive and planning functions. In addition, the ToL has also been used in numerous studies within the domain of cognitive psychology (cf. Gilhooly, Phillips, Wynn, Logie, & Sala, 1999; Hodgson, Bajwa, Owen, & Kennard, 2000; Newman & Pittman, 2007; Phillips, Wynn, McPherson, & Gilhooly, 2001; Kaller, Unterrainer, Rahm, & Halsband, 2004; Kaller, Rahm, Bolkenius, & Unterrainer, 2009; Ward & Allport, 1997).

Typically participants receive a ToL problem (see Fig. 1a) as a start state (A) and a goal state (B). The task is to find a shortest sequence of moves transforming the start state into the goal state. A move consists of a colored bead, a start peg and a target peg. The constraints for executing a move are: (1) only one bead may be moved at a time and (2) only the top bead on any peg may be moved. ToL problems differ with respect to the number of moves, the number of beads and structural characteristics of problems to be solved (cf. Kaller, Rahm, Spreer, Weiller, & Unterrainer, 2011). The ToL task is in some respects similar to the Tower of Hanoi (ToH) task (see Fig. 1b) as it shares a similar environmental structure, the same constraints concerning the moves, and the kind of problem to be solved (cf. Kaller, Rahm, Köstering, & Unterrainer, 2011). The beads in the ToH task, however, are distinct by their size. Therefore, the task has the additional constraint that only smaller beads may be placed on larger beads.

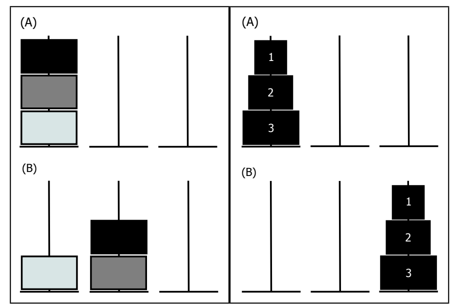

(a) Tower of London (b) Tower of Hanoi

Figure 1: The goal in both tasks is to transform the start state (A) into the goal state (B).

The present work aims at elucidating the influence of structural task properties on planning behavior. Analyses are focused on the ToL that - compared to the ToH - was applied in the vast majority of related publications in MEDLINE-listed journals (cf. Kaller, Rahm, Köstering, & Unterrainer, 2011). We provide an ACT-R model which uses a general heuristic capable of solving the ToL tasks assessed in the experimental study. This model uses different strategies which are selected based on the structural distribution of beads in the environment.

State of the Art

The Cognitive Architecture ACT-R. ACT-R is a modular cognitive architecture with an underlying production system operating on symbolic representations of declarative memory items – so-called chunks (Anderson & Lebiere, 1998; Anderson et al., 2004). The system consists of specific modules corresponding to certain aspects of human cognition. Hence, the system provides a module for processing visually presented information (vision module), a module for directing goal driven behavior (goal module), a module for inferring new information (imaginal module), and a module to store long-term memory items (declarative module). Each module has a dedicated interface (buffer) which can store one chunk at a time.

The functionality of the system is driven by production rules which represent the procedural memory component. A

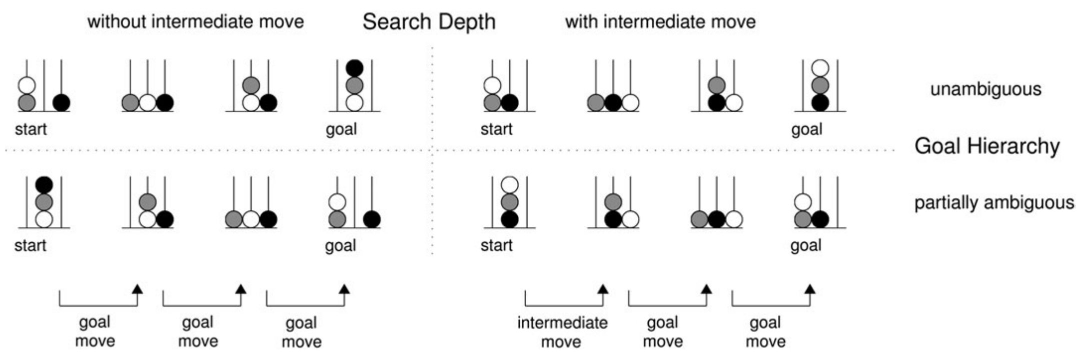

Figure 2: Problem classes as analyzed by Kaller, Rahm, Spreer, et al. (2011). The problem structure is unambiguous in the goal state if all beads are on one peg whereas it is ambiguous if beads are distributed over two pegs. An intermediate step in the planning process requires a move of a bead to another peg, not his goal peg (e.g. the white bead in the right/down corner). Figure taken from Kaller, Rahm, Spreer, et al. (2011), reprinted by permission of *Oxford University Press*.

production rule consists of a set of preconditions which are formulated in terms of buffer contents and a set of buffer manipulating instructions. If the preconditions of one production rule hold, it is selected and the buffer manipulations are executed, whereby replaced chunks are committed to declarative memory.

On a sub-symbolic level the accessibility of declarative memory items and the selection of production rules are influenced probabilistically. For long-term memory items an activation value is calculated which decays over time if the chunk is not used and boosted otherwise. This activation value determines the probability for correctly retrieving the requested chunk. If the preconditions of two or more production rules hold for the current buffer contents the selection depends on an utility value which is associated with each production. As a consequence, the utility value determines the probability for selecting a specific production rule.

ACT-R models of Tower of Hanoi. Anderson, Albert, and Fincham (2005) used Tower of Hanoi problems with an optimal solution of 28-moves to investigate brain regions which are active during the planning process in order to be able to provide a mapping to ACT-R components. Therefore, participants were trained with a general heuristic to enable them to solve the task accurately and in an optimal manner.

Altmann and Trafton (2002) assessed goal driven behavior of participants. They argued that task subgoals, which are built by goal decomposition, are ordinary memory objects and therefore have to be committed to memory if another subgoal is on focus. To revisit a subgoal it has to be retrieved from memory, whereby mental rehearsal is used for every constructed subgoal in order to resist decay. To assure the correct retrieval order cues presented by the environment are used.

Fum and Missier (2001) investigated the dependencies between strategy selection and contextual/task factors. For a specific problem they argued that for strategy selection achieving a good trade off between accuracy and cognitive effort is most important. The conducted experiment tested 34 participants on 4-bead ToH tasks. Participants were separated randomly into two groups, one instructed to solve the task in the shortest possible time the other one to solve the task with a minimum number of moves. To explain the differences regarding error rate and response time they modeled several known planning strategies for the ToH.

Empirical Findings

A study by Kaller, Rahm, Spreer, et al. (2011) analyzed Tower of London problems with an optimal solution of three moves from start state to goal state. 24 participants received 96 ToL problems in randomized order. In a first step, they had to plan all moves in advance. In a second step they were asked to execute the planned moves. This design has been chosen to separate the planning phase and the execution phase. Two task parameters search depth (with/ without an intermediate move) and goal hierarchy (unambiguous/ partially ambiguous) were experimentally manipulated resulting in a 2x2 factorial design:

1. *Goal hierarchy* describes the arrangement of beads on the target pegs.
   (a) Unambiguous goal hierarchy (U): All beads are placed on one target peg (see first row in Fig. 2).
   (b) Partially ambiguous (P): The beads in the goal state are arranged on two different target pegs (see second row in Fig. 2).
2. *Search depth* depends on the number of moves which have to be conducted for a certain bead before it is placed on its target position.
   (a) With intermediate step (U-I, P-I): A bead has to be moved out of the way before the task can be solved (see right side in Fig. 2).
   (b) Without intermediate step (U, P): Every bead can be moved to its target position directly. The solution for the

problem is therefore simply an ordering of moves which can be inferred directly from the environment (see left side in Fig. 2).

Cognitive Model

Problem Representation. As described above, a cognitive model in ACT-R is divided into declarative and procedural knowledge. In the reported experiment all declarative information is presented visually to the participants. Consequently, the cognitive model receives the information visually and must operate on an internal representation. Therefore, a special representation is used which enables the agent to interpret a visual object as a certain part of a ToL board including relations to other objects (e.g. a certain bead is located on a certain peg). Fig. 3 shows the information the model is able to infer from the environment. A chunk representing a peg captures additional information concerning the colors of the beads placed on it. Additionally, a bead holds information concerning the peg number and the slot number it is placed on.

Inferring Knowledge. Besides visual information provided by the environment additional declarative information is necessary to solve a given instance of the ToL. Therefore, we use subgoals which capture information of the start and the target position of a certain bead, whereby the bead is named by its color. For processing and linking subgoal information we included findings from Altmann and Trafton (2002) by treating these chunks as ordinary memory objects. Nevertheless, their solution had to be adjusted to capture the specific requirements of the ToL. First, the ordering of subgoals is not provided by cues in the environment for the reason that the beads itself are not ordered by size. Humans, however, are able to remember the last inspected bead. As a consequence, a subgoal holds additional information concerning the color of an old subgoal which could not be executed before. Nevertheless, as only the color is saved the old subgoal has to be retrieved from declarative memory. Additionally, as no ordering is provided by the environment we assume that the retrieval of old subgoals can fail. Therefore, we do not use a rehearsal mechanisms to artificially prevent subgoals from decay.

Furthermore, there is also a major difference in the experimental setting mostly used for the ToL compared to the experimental settings used to asses the ToH. Whereas in the ToH the environment may be changed if a subgoal is executable, in the ToL it must not. As a consequence, a mental representation is needed which can capture the current state during the solution process. Therefore, we use one chunk representing the current state in order to keep track of changes in the overall solving process. The representation of a current state is necessary for the reason that start and target position have to be updated to enable a correct solution process.

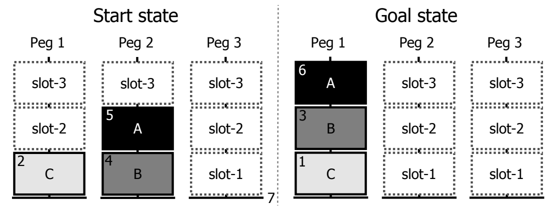

Figure 3: Environmental information encoded by the model.

Heuristic. Besides declarative information, procedural knowledge is used to model all activity necessary to solve a ToL problem. As our goal is to provide a generalized heuristic for solving tower tasks which is, however, capable of explaining planning behavior specific for the ToL, we integrated a heuristic developed by Anderson et al. (2005). To solve the ToH their model constructs a subgoal for the largest bead not on its target location (see Fig. 1b, bead 3). Given the constraint that only smaller beads may be placed on larger beads this subgoal has to be solved first. The largest bead (bead 3), however, is blocked by another bead (bead 2). Consequently, a subgoal for moving the other bead (bead 2) out of the way is constructed. This process goes on until a subgoal is finally executable. Eventually, the bead corresponding to the executable subgoal is moved and the last subgoal which was not executable is retrieved from long-term memory. For our ToL model we use a similar heuristic, which is shown in Fig. 5. In a first step, a bead of a certain color is selected (select-bead). For the selected bead, start and target location are determined and a subgoal is constructed based on the collected information (find-counterpart). Afterwards, the start and the target position are compared (infer-distance). If the corresponding bead is not already on the correct location, the constraints are checked in order to test if the subgoal can be solved directly (check-constraints). At this point, three different inferences are possible. First, if the subgoal is executable, the current working state is updated and the heuristic starts again for another subgoal (update-internal-representation, retrieve-IR). Second, if the subgoal cannot be executed given the constraints the heuristic starts over by constructing/retrieving a subgoal for the blocking disk. Third, it is possible that in a task no bead can be moved to its target location directly. In this case, a cycle between two subgoals is detected and an intermediate step for a bead is processed (select-peg-for-IS).

Example. Fig. 4 shows an example of how the ToL heuristic works on the given knowledge representation by solving the ToL instance shown in Fig. 3. The visual locations visited by the model are referenced in Fig. 3 using numbers in the upper left corner of a bead. The goal buffer stores information concerning a current subgoal. The imaginal buffer holds a chunk with one slot for each possible bead location in the initial state as this representation is changed during the planning process. The retrieval buffer contains a chunk which could be retrieved from long-term memory.

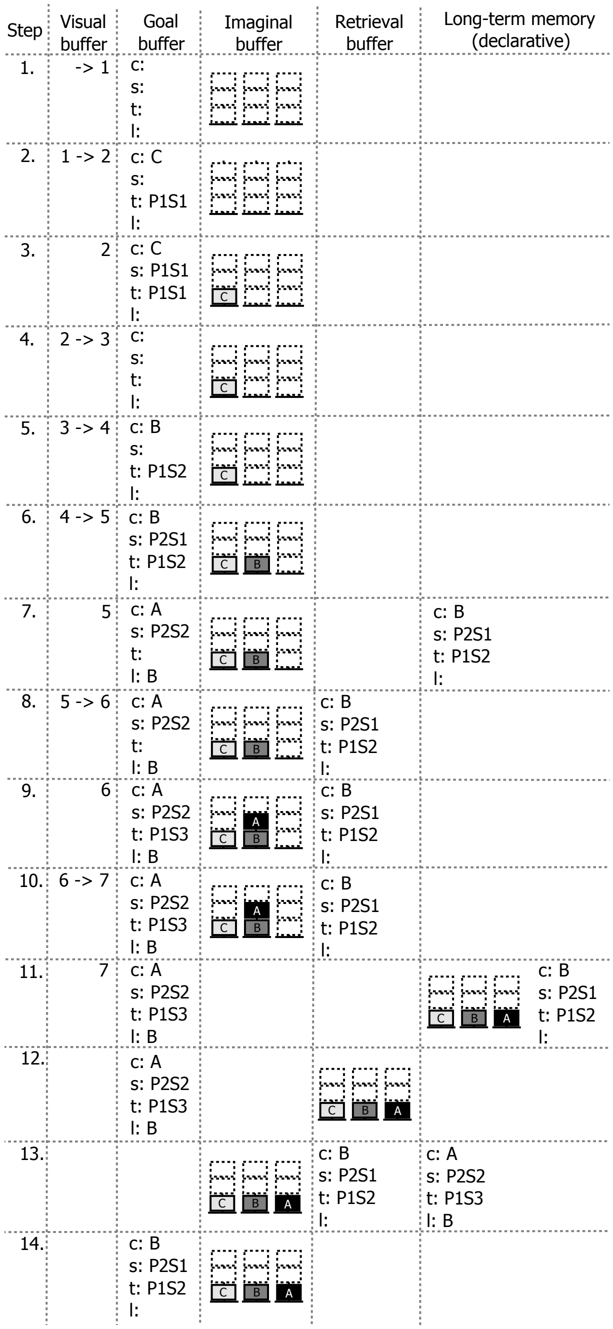

Figure 4: Overview: Model resolving an intermediate step as depicted in Fig 3. Columns denote buffer contents of active modules. Visual locations correspond to the visual locations marked with numbers from one to seven in Fig. 3. Goal buffer slots are: color of a bead (c), start location of a bead (s), target location of a bead (t) and last subgoal processed (l). The following steps are executed by the model (see Fig 5): 1. select-bead; 2. find-counterpart; 3. infer-distance; 4. select-bead; 5. find-counterpart; 6.infer-distance; 7. check-constraints; 8. find-counterpart; 9. infer-distance; 10. check-constraint; 11. select-peg-for-IS; 12. update-IR; 13. retrieve-internal-representation; 14. select-bead.

The heuristic works as follows: First bead C is selected (step-1) and a corresponding subgoal is constructed including information concerning the target location (step-2). The start location of bead C is determined and the subgoal is updated accordingly. As bead C is already on the correct location no further processing for the current subgoal is necessary (step-3). The heuristic starts over by selecting bead B (step-4). As for bead C a subgoal for bead B is constructed and checked (step-5, step-6). The check shows that bead B is not on the correct location. Therefore, it is determined if bead B can be moved given the constraints. The check results in a violation of a constraint, as bead B is blocked by bead A (step-7). A new subgoal for bead A is constructed (step-7), completed (step-8) and checked (step-9). The subgoal corresponding to bead B is committed to memory. Additionally, the color of bead B is saved to indicate the subgoal which structurally led to the new subgoal (step-7). In order to be able to infer structural dependencies between the positions of bead A and bead B, the model also tries to retrieve the subgoal for bead B (step-8). If the retrieval of the subgoal corresponding to bead A was successful, the model is able to infer that an intermediate step is necessary (step-10). Otherwise, the necessary information have to be collected again by visual search. For the processing of an intermediate step an intermediate position has to be selected. In this context, the intermediate position may not interfere with start and target locations of the involved subgoals. In this case, the model chooses peg 3 as an intermediate position for bead A (step-11). To complete this mental operation with capturing the costs the working state in the imaginal buffer is updated, committed to long-term memory (step-12), retrieved from long-term memory (step-13) and reconstructed in the imaginal buffer (step-14). Additionally, the subgoal corresponding to bead B is retrieved in order to check if it is executable now. At this point, all cycles are resolved. In the next two runs through the heuristic the moves for bead B and bead A can be planned and the problem is solved.

Strategies. Besides the general heuristic we assessed more sophisticated strategies with regard to structural task characteristics tested in the experimental study. These strategies allow for an improved heuristic if the correct inference can be drawn based on additional declarative information accessible in a step. These additional strategies, however, only extend the heuristic. The probability of selections depends on an associated utility value. All improved heuristics are shown in Fig. 5 as dotted lines. The labels of dotted lines comply with strategies represented in Table 1.

The first criteria assessed in this context is the goal hierarchy (U). We assume that certain constraints do not have to be checked if all beads are located on one peg in the goal state. The reason is, that the target peg for all disks remains the same and only the correct ordering has to be determined. After the first bead can be moved to the bottom position of the target peg it is not necessary to check for the second bead if the mid position of the target peg is reachable.

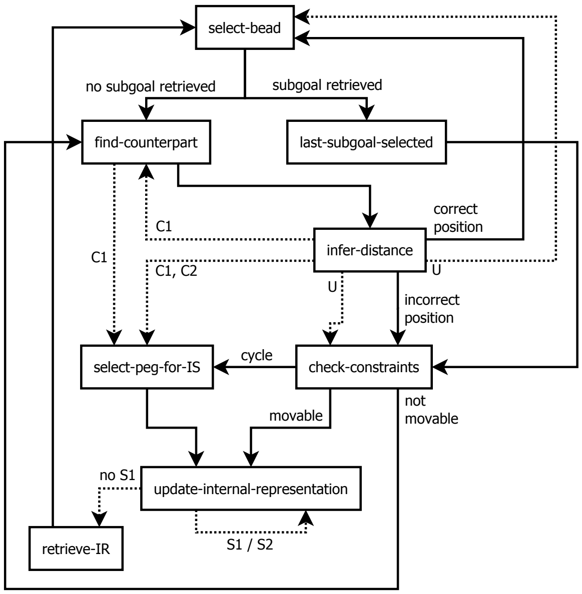

Figure 5: Overview of the Heuristic: Shows the heuristic used by the model to solve 3-move Tower of London problems. Black arrows relate to a specific decision based on current information. These arrows are labeled with the information used in a specific step. Dotted arrows show uncertain decisions based on task structure (for the labels, cf. Table 1). For the utilities please refer to Fig. 6.

For the second criteria concerning the search depth we provide additional mechanisms to infer that an intermediate step is necessary. In this context we assessed two task structures used in the experimental study. In the first structure a bead is already on the correct peg but not in the correct slot (C1). Therefore, it is necessary to move this bead out of the way first. This inference can be drawn based on a comparison of the start and a target position of a certain bead. In the second structure a tower consisting of two beads has to be moved to another peg without violating the order between the beads (C2). Fig. 3 shows an example for this structure. That an intermediate step is necessary to solve this structure can be inferred by comparing start and target locations of the corresponding subgoals. In both situations the necessary information are given before the constraints are checked for both subgoals. As a consequence, this step can be skipped if the inference is drawn by the model.

Additionally, we assessed the overall overview participants gained when the necessity of an intermediate step is inferred. This includes the question of how many subgoals can be solved in the phase where an intermediate step is processed. When an intermediate step is processed, the current goal is to enable a move for one bead to its target location. In this context, the goal of moving another bead out of the way is only a subgoal of the current goal. Therefore, we assume that the enabled goal move can be executed right after the other bead was moved out of the way. Additionally, it is possible that the goal move for a bead which was moved out of the way is also possible without an additional run through the heuristic.

Table 1: Overview of the strategies applied by the model to assess structural criteria. For an illustration of strategy application compare Fig. 5, labels on dotted lines.

U	Using structural information for unambiguous problems.
C1	Using the structural information that a bead is already on the correct peg but not on the correct slot.
C2	Using the structural information that two beads have to be moved with keeping the order between (see Fig. 3 bead A and bead B).
S1	Conclude an goal move for a blocked bead right after conducting an intermediate step
S2	Conclude a goal move for the bead moved in the intermediate step right after S1. (This is only possible after S1.)

Evaluation. To evaluate the cognitive model we assessed the performance of the general heuristic in combination with the additional strategies. For this reason, we varied the utility values of the additional strategies systematically. Adjusting utility values controls the probability for choosing specific strategies in contrast to executing the general heuristic. To exclude other structural influences we systematically varied the different start and goal states. The tested task structure, however, was preserved. Additionally, we used three possible starting strategies: backward planning (planning beads ordered from bottom to top on a peg in the goal state), forward planning (planning beads ordered from top to bottom on a peg in the start state) and a mixed strategy with forward planning for partially ambiguous problems and backward planning for unambiguous problems. The results of the evaluation are shown in Fig. 6, whereby the best fit is provided by the combination of probabilities shown in Fig. 6d.

Discussion

We developed a Tower of London ACT-R model able to replicate behavioral data reported in Kaller, Rahm, Spreer, et al. (2011). These problems consisted of four problem classes posing different degrees of difficulty. In a first step, we adapted and adjusted some heuristics from Tower of Hanoi (cf. Anderson & Douglass, 2001). In a next step, we evaluated these heuristics and extended them to account for task specific structural characteristics. The results indicate that the recognition of structural characteristics is an important factor for explaining human planning behavior in the context of the ToL. For problems without an intermediate step (U, I) the general heuristic is capable of explaining latencies in planning. Here, the correct order of moves has to be determined to solve the task and, therefore, all steps of the heuristic have to be processed. For problems with an intermediate step (U-I, P-I), however, certain steps of the heuristic may be skipped.

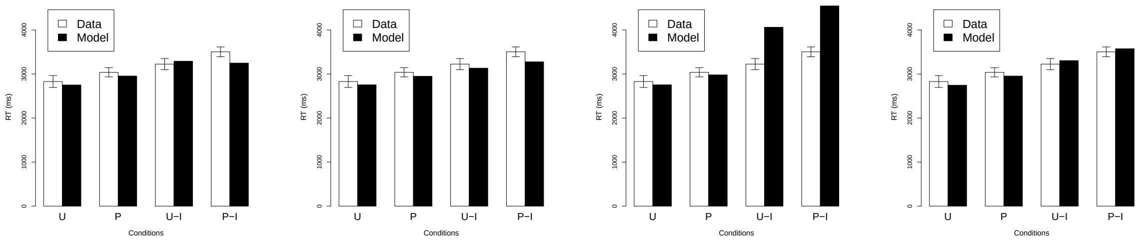

(a) C1, C2, S1: 50% (b) C1, C2: 91%; S1: 50% (c) C1, C2: 91%; S1: 0% (d) C1, C2: 91%; S1: 66%

Figure 6: Human data compared to model predictions illustrating the influence of the selective probability variation on strategy application. For the strategies used see Table 1. The probabilities for the different conditions are (a) C1, C2, U and S1: 50%; (b) C1 and C2: 91% S1 and U: 50%; (c) C1 and C2: 91%, U: 50%, S1: 0%; (d) C1 and C2: 91%, U: 50%, S1: 66%. The combination of probabilities in (d) provides the best fit.

The reason is, that structural information can be observed visually in the environment which trigger the inference that an intermediate step is necessary. Additionally, participants may gain enough overview of the overall task structure to solve the task directly after the intermediate step was inferred. A similar mechanisms was found for problems with an unambiguous goal hierarchy (U, U-I). In these cases it is not necessary for participants to check all constraints if the correct inference can be drawn based on the information provided by the environment. Taken together, the present systematic analyses allowed for an identification of specific strategies presumably applied by human planner. For future refinements, concurrent evaluation of eye-tracking data may provide more insights about the planning process (cf. Kaller et al., 2009).

Acknowledgements

This research was supported by the DFG (German National Research Foundation) in the Transregional Collaborative Research Center, SFB/TR 8 within project R8-[CSPACE] and a grant of the German Federal Ministry of Education and Research (BMBF-01GW0710).

References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39-83.

Anderson, J. R., Albert, M. V., & Fincham, J. M. (2005). Tracing problem solving in real time: fMRI analysis of the subject-paced Tower of Hanoi. *Journal of Cognitive Neuroscience*, *17*, 1261-1274.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036-1060.

Anderson, J. R., & Douglass, S. (2001). Tower of Hanoi: Evidence for the cost of goal retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1331-1346.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Fum, D., & Missier, F. D. (2001). *Adaptive selection of problem solving strategies.* Erlbaum.

Gilhooly, K. J., Phillips, L. H., Wynn, V. E., Logie, R. H., & Sala, S. D. (1999). Planning processes and age in the five-disc Tower of London task. *Thinking & Reasoning*, *5*, 339-361.

Hodgson, T. L., Bajwa, A., Owen, A. M., & Kennard, C. (2000). The strategic control of gaze direction in the Tower of London task. *Journal of Cognitive Neuroscience*, *12*, 894-907.

Kaller, C. P., Rahm, B., Bolkenius, K., & Unterrainer, J. M. (2009). Eye movements and visuospatial problem solving: Identifying separable phases of complex cognition. *Psychophysiology*, *46*, 818-830.

Kaller, C. P., Rahm, B., Köstering, L., & Unterrainer, J. M. (2011). Reviewing the impact of problem structure on planning: a software tool for analyzing tower tasks. *Behavioural Brain Research*, *216*, 1-8.

Kaller, C. P., Rahm, B., Spreer, J., Weiller, C., & Unterrainer, J. M. (2011). Dissociable contributions of left and right dorsolateral prefrontal cortex in planning. *Cerebral Cortex*, *21*, 307-317.

Kaller, C. P., Unterrainer, J. M., Rahm, B., & Halsband, U. (2004). The impact of problem structure on planning: insights from the Tower of London task. *Cognitive Brain Research*, *20*, 462-72.

Newman, S. D., & Pittman, G. (2007). The Tower of London: A study of the effect of problem structure on planning. *Journal of Clinical and Experimental Neuropsychology*, *29*, 333-342.

Phillips, L. H., Wynn, V. E., McPherson, S., & Gilhooly, K. J. (2001). Mental planning and the Tower of London. *Quarterly Journal of Experimental Psychology*, *54A*, 579-597.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London*, *B 298*, 199-209.

Ward, G., & Allport, A. (1997). Planning and problem-solving using the 5-disc Tower of London task. *Quarterly Journal of Experimental Psychology*, *50*, 49-78.