Towards a Companion System Incorporating Human Planning Behavior

A Qualitative Analysis of Human Strategies

B. Leichtmann[1], P. Bercher[2], D. Höller[2], G. Behnke[2], S. Biundo[2], V. Nitsch[1] und M. Baumann[3]

[1] Institute of Industrial Engineering and Ergonomics, RWTH Aachen University
b.leichtmann@iaw.rwth-aachen.de, v.nitsch@iaw.rwth-aachen.de

[2] Institute of Artificial Intelligence, Ulm University
pascal.bercher@uni-ulm.de, daniel.hoeller@uni-ulm.de, gregor.behnke@uni-ulm.de, susanne.biundo@uni-ulm.de

[3] Human Factors Department, Ulm University
martin.baumann@uni-ulm.de

Abstract

User-friendly *Companion Systems* require Artificial Intelligence planning to take into account human planning behavior. We conducted a qualitative exploratory study of human planning in a knowledge rich, real-world scenario. Participants were tasked with setting up a home theater. The effect of strategy knowledge on problem solving was investigated by comparing the performance of two groups: one group (n = 23) with strategy instructions for problemsolving and a control group without such instructions (n = 16). We inductively identify behavioral patterns for human strategy use through Markov matrices. Based on the results, we derive implications for the design of planning-based assistance systems.

Kurzzusammenfassung

Zur Berücksichtigung menschlichen Planungsverhaltens in Companion Systemen – eine qualitative Analyse menschlicher Planungsstrategien

In nutzerfreundlichen *Companion Systemen* muss die Künstliche Intelligenz in ihrer Planung auch menschliches Verhalten berücksichtigen. Daher untersuchten wir in einer qualitativen Studie explorativ menschliches Planungsverhalten in einer wissensreichen, realistischen Anwendungsdomäne. In einem Experiment sollten die Teilnehmenden eine Heimkinoanlage aufbauen. Dabei wurde der Effekt des Strategiewissens auf Problemlöseverhalten über zwei Gruppen untersucht: Eine Gruppe mit Instruktionen einer Problemlösestrategie (n = 23) und eine Gruppe ohne (n = 16). Durch Induktion wurden aus Markovnetzen Muster im menschlichen Planungsverhalten identifiziert und entsprechend Implikationen für die Gestaltung von planbasierten Assistenzsystemen abgeleitet.

Keywords

ill-defined problem solving, human-computer interaction, plan linearization, planning

1 Assisting Humans in Action Planning

Digital personal assistants become more and more present in today's society. There is a wide range of such systems, e.g. personal assistants running on smartphones or control engines for smart homes. But most everyday technical devices give us support that still needs to be adopted or transferred to the current situation at hand. *Companion Technology* aims at utilizing cognitive technical systems that assist its users in an individualized and user- and situation-adaptive way [Biu16b]. At the heart of many *Companion Systems* is Artificial Intelligence (AI) planning [Biu16a]. It allows to reason about complex courses of action, i.e. what actions need to be applied in which order to achieve a certain goal.

When assisting humans in action planning based on AI planning technology, several challenges arise: Which planning decisions should be made by the user and which ones by the AI planning component, and in what order should alternative options be presented? The knowledge about human planning behavior might thus help the system in circumventing misunderstandings with human users by establishing a common understanding of the problem and its solution. After a plan has been generated, i.e., a set of actions that achieves the user's goals, there are usually several orderings (also called *linearizations*) of the actions that can be applied. These are generated explicitly when using Partial-Order Causal-Link (POCL) planners; when using state-based planners, they can be easily obtained via post-processing. Some of the linearizations may be more intuitive for humans than others; and a most suitable order in which they are presented has to be found. While there has been done ample research on human planning behavior (see e.g. Morris and Ward [Mor05], the integration of psychological findings with AI planning for automatic decision support has not received much attention. In order to address the research gap, we have conducted an exploratory study of human planning behavior in a knowledge-rich, ill-defined real-world scenario as a first attempt. The aim of our research was to show a difference in strategy use across available instructions as a manipulation of information availability or existing knowledge. We investigated behavioral patterns of planning strategies which might serve as cues indicating the reasoner's mental representation of the plan through inductive qualitative analysis of transition diagrams. With our findings we aim to provide support for the decision-making processes of *Companion Systems* with psychological insights and, thereby, to improve their overall functionality and usefulness.

In the following section, we provide an overview of related psychological perspectives as well as AI concepts, before describing the experimental study and its results. We close with a discussion of our findings with respect to psychological planning research and implications for AI planning technology.

2 Planning in Psychology and AI

2.1 Human Planning Behavior

According to Mumford et al. [Mum01] planning is a mental simulation of single actions in a dynamic environment, which is a goal-driven and resource-intensive activity. In order to reduce costs, humans use heuristics to overcome limitations in working memory capacity or the lack of knowledge in long-term memory. This is especially the case when the problem is ill-defined with an unknown problem space and a dynamic and uncertain environment as found in real-world settings.

In studies with well-defined or well-structured puzzle problems [Dav05], where the initial state, the goal state, and operators (= rules how moves can be made) are given, models have been proposed suggesting that problem-solving takes place in a rational, systematic, top-down manner. For well-defined problem spaces, heuristics such as hill-climbing or means-end analysis are used including a systematic solution search. With such strategies, the difference between the current state and a goal state is assessed and operators are used that minimize that difference.

In ill-defined real-world tasks, parameters such as the final goal state or some operators may not be known and planning behavior is thus different from the systematic approach to well-defined problems. In many real-world, knowledge-rich tasks, human planning was found to be mostly of opportunistic nature [Hay79] or in other words bottom-up. That means that plans are continuously changed and evolve as environmental cues trigger opportunities for plan refinement. Human planning occurs in partial order and non-hierarchically online during task implementation (= concurrent planning). Planning steps are continuously formulated as the solution evolves [Dav05]. In ill-defined problems, only little is known about the problem space, which makes the above mentioned process of difference evaluation such as means-end analysis difficult [Orm05]. Ormerod [Orm05] gives an overview of several human strategies. The usage depends highly on the task type, the task complexity, and the level of expertise of the problem solver.

Most strategies share an approach similar to means-end analysis, where a goal is decomposed into smaller subgoals until a known operator is applicable, but not in a purely systematic manner. Humans formulate partial plans, but environmental cues might trigger a reevaluation of the plan resulting in changing the plan opportunistically. The planning behavior is thus rather non-hierarchical, especially when confronted with ill-defined problems.

Therefore, planning can include structured or unstructured (= opportunistic) aspects. For ``creative'' problem tasks with no single best outcome, Ormerod [Orm05] proposes a mixed strategy. An example of such a ``creative'' problem may be the design of a novel educational task [Orm05] or planning a dinner [Beh17]: Several courses have to be matched and constraints such as allergies have to be considered, but there is no single best outcome.

2.2 Mixed-Initiative AI Planning

AI planning is concerned with finding so-called *plans* to solve a given *planning problem*, i.e., to find a (maybe partially ordered) sequence of actions that transforms an initial world state into a state that satisfies all goals, which are specified in the problem description. These goals do not need to describe a complete world state. Consider a problem where a user wants to set up a home theater consisting of various HiFi devices by means of several cables and adapters. Here, the goals do not need to mention which cables are plugged into which devices; instead, the planning goals only specify the signals required by the respective devices. In fact, there are various possibilities which cables will be used and where - and the planner will find one solution on its own. Usually, planning systems do this without a user taking part in this selection process. In contrast, in mixed-initiative planning (MIP), the planning process is done collaboratively with the user and certain decisions are taken by her or him [All96, Bur96].

There are various challenges yet unsolved in MIP [Smi12, Beh17]. One of them is deciding in which situations the planner should take initiative (i.e., solve the problem, or parts of it, fully autonomously), and in which ones the user should do so by her- or himself.

After a planning problem has been solved, i.e., the steps that need to be carried out to reach the user's goals are known, the respective solution is usually communicated to the user step by step to be carried out by her or him [Ber17].

In previous work [Höl14], we proposed techniques to find user-friendly plan linearizations, i.e., we try to find the most appropriate order in which the plan's actions are communicated to the user. Our strategies are based on a POCL plan - a partially ordered plan incorporating causal links, i.e., annotations that explicitly state which action fulfills the preconditions (the conditions necessary to execute the action) of another action; the former is called the *producer*, the later the *consumer*. This plan linearization has nothing to do with *solving* the task, as *any* linearization of the solution's actions adhering the present constraints solves the problem. When using POCL planners, the result of the planning process is already in that representation. When using state-based planners, it can easily be obtained via post-processing in polynomial time. Some of the linearizations might be more intuitive than others to human users. We have developed various domain-independent strategies to prioritize the different possibilities in a way that deemed plausible [Höl14, Ber17], but that had not yet been evaluated in a study: Two of our strategies are based on information

present in any POCL plan: The first one is based on the structure of causal links. It finds a linearization ensuring that the producer and the consumer of a causal link are close together. Consider our home theater domain: When the action of *open the case* of a DVD fulfills a precondition of *taking* the DVD, they should be close to each other in the plan (though from a technical view, one could reassemble the whole home theater in the meantime). The second strategy exploits the fact that most planning models are given in a lifted fashion (though planning is often done grounded). As a consequence, a plan has parametrized actions, e.g., *⟨openCase(dvd1), take(dvd1), putInto (dvd1, player1)⟩*. We use this representation to group actions that include the same constants. \label{linConstants} Our third strategy is applicable when planning in a hierarchical way, i.e., decomposing an abstract overall task into subtasks until an executable plan is found. Here, actions that are introduced by the same decomposition step are grouped together since they contribute to the achievement of the same task at a more abstract layer.
Closely related is also the task of generating so-called plan explanations [See12, Ber17]. Instead of only presenting a plan's actions to the user, additional explanations about the respective steps can also convey the purpose of that step. That is, it is not just shown *what* needs to be done by the user, but *why* this should be done. So far, these explanations are only provided upon user request. Providing such explanations pro-actively might be one way to establish a common ground.

3 Qualitative Pilot Study

In our exploratory study we took the home theater setup task proposed by Bercher et al. [Ber14, Ber15] as a real-world scenario and analyzed human planning strategies. We used a between-subject design with two groups varying in strategical background knowledge. One group received further instructions about a useful strategy and one group got no such instructions. Differences in planning behavior were analyzed as dependent variable qualitatively through observations of the order of actions. We expected to find more strategy use in the instruction group compared to the no-instruction group in all scenarios. Therefore we expected to find cues of opportunistic behavior rather than any use of the proposed strategies as systematic behavior in the non-instruction group.

3.1 Methods

3.1.1 Research Design

In our experiment, experience was manipulated as independent variable with two groups (between-factor), with one group receiving instructions on a successful problem solving strategy (experimental group; description below) and the other group receiving no such instructions (control group). To increase reliability the task was presented to each participant in three different scenarios. Problem solving success was assessed as dependent variable. Furthermore, domain knowledge, working memory capacity, need for cognition, and technology commitment were assessed as control variables.

3.1.2 Task: Setting Up a Home Theater

We have chosen the task of setting up a home theater as proposed by Bercher et al. [Ber14, Ber15] because of the knowledge-rich and realistic nature and because an AI planning-based assistant was available for that task. For this task, different technical devices have to be connected so that every HiFi component receives the required audio/video signals. For example, the television has to receive the video signals of a blu-ray player and a satellite receiver, and the audio/video receiver has to receive the respective audio signals. For the connections, several cables and adapters are available with different characteristics. For instance, not all of the cables transfer the same signal types or can be used for every device. The advantage of choosing this task compared to others is the possibility of systematically varying the difficulty by varying the availability of different cables.

The task was carried out in a virtual desktop environment where devices and cables could be moved, connected, and disconnected via mouse input. An example of the task is pictured in Fig. 1.

3.1.3 Measurements

Because the described setup task is a knowledge-rich task, domain knowledge of participants was assessed. For this purpose, a performance test had been developed and pre-tested in a pilot run with 29 participants.
One such question was *What signal(s) is/are transmitted by a HDMI cable?* with possible answers *a. video only*, *b. audio only*, *c. both audio and video*, and d. *there are different HDMI types with different signal transmissions*.
In addition, two other methods measuring domain experience other than domain-specific knowledge were used, namely a self-report with 4 items
(e.g., *How would you estimate your practical skills in setting up a home theater?*);
and a retrospective questionnaire, also with 4 items (e.g., *How often have you set up a sound system in a room?*).
Furthermore, other cognitive and personality constructs were assessed (working memory capacity, need for cognition and technology commitment), but analyses and results are not reported here. All items were phrased in German and can be obtained from the authors upon request.

3.1.4 Participants

In total N = 39 German-speaking psychology, biology, and medical students of Ulm University participated in the experiment. 26 identified themselves as female, 13 as male, and no person as any other gender. The age ranged from 18 to 49 years with an average of M = 25.1 (SD = 5.3). The participants were split into groups with 23 being in the experimental group with strategy instructions and 16 being in the control group with no instructions.

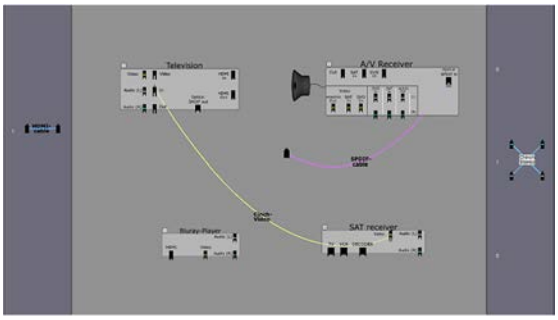

Figure 1: Task “Setting Up A Home Theater“ in a virtual desktop environment showing scenario C of the experiment. The participant already made connections

3.1.5 Procedure

In the experiment, both groups started with a session of filling out the questionnaires and performing the cognitive tests mentioned above. Afterwards, both groups received written instructions about the task and overview sheets about the cables and devices. Participants had access to these instructions and overviews throughout the experiment. Also, after reading the material, participants were able to ask questions in case of misunderstandings, followed by test scenario as a training. Subsequently, the experimental group received further information on a useful strategy. They were instructed to follow the signal flow beginning from the source devices (bluray player, satellite receiver) to the output devices (television, audio/video receiver). For example following this instruction a participant can transfer the video signal by plugging in a video cinch cable at the respective port of the satellite receiver first and then at the television afterwards. The control group received no further strategy instructions and continued directly with the main task. It was carried out by both groups and consisted of three different scenarios (*A, B, C*) varying in cable availability. Fig. 1 shows scenario C. In the other scenarios signals had to be transferred with other sets of cables or adapters and thus other solutions, but the devices and signal types stayed the same.

3.2 Results

For the statistical data analysis, the statistic software R [Dev12] was used. For the exploratory analysis of planning behavior, sequence analyses were performed using the “markovchain“ package [Spe17].

3.2.1 Descriptive Results

In Table 1 the percentage of users solving the scenarios can be found for every scenario of the task, separately for each group or in sum. It shows that participants in the instruction group seem to perform better in every task, although because of the small sample size and thus low statistical power, statistical inference does not seem legit and this observation needs to be interpreted with caution. It can also be seen that scenarios seem to vary in difficulty.
The newly developed knowledge test with 6 items in total had a mean difficulty of $M = .45$ ($SD = .29$). 4 of these items were close to the chance level ($\leq .29$) and 2 items showed a low difficulty above $\geq .79$. Participants were found to recognize an HDMI cable and knew its functions, but had no or very low knowledge about other cables (S/PDIF) or devices (blu-ray player). Also 64% of the participants rated their ability to set up a home theater as “very low“, “low“, and “rather low“.
In the retrospective questionnaire 59% of the participants reported that they had never set up a home theater.
In sum data suggests that participants are rather novices than experts in the domain of setting up a home theater.

Table 1: Percentage of correctly solved tasks for each group and scenario

Group	Scenario		
	A	*B*	*C*
Instruction Group	.61	.48	.78
Control Group	.31	.25	.50
Sum	.49	.38	.67

3.2.2 Strategy Usage

The actual planning behavior and the usage of potential strategies were investigated through qualitative log file analysis. All actions (connections and disconnections of cables) were automatically saved by the computer while the participants were solving the problem in the virtual environment. For qualitative sequence analysis, Markov models were calculated indicating the probability with which one action followed another during the tasks. Markov chains are examples of stochastic processes referring to a sequence of random variables $X_0, X_1, X_2\ldots$ evolving over time (X_n with a discrete state at time n) and the assumption that the future state is solely conditional on the current state and independent of the history of past events. The transition probability from state *i* to state *j* is q_{ij} with ($P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1} = i_{n-2},\ldots, X_0 = i_0) = P(X_{n+1} = j \mid X_n = i) = q_{ij}$. To analyze which actions were taken after another, transition probabilities were calculated for each experimental group separately and for each scenario *A*, *B*, and *C*. The transition probabilities were calculated for every action (what cable was connected with which device?), for the devices (which device was used after another?), and cables (which cable was used after another?).
The main results are reported in the following.

For the devices, it was found that participants in the experimental group started with devices defined as source devices (blu-ray player or satellite receiver) with a higher probability ($q = .74$) than output devices (television and audio/video receiver; $q = .26$) in scenario *A*. This was also true for scenario *B* ($q = .87$ for starting with a source device compared to $q = .13$ for starting with an output device). In scenario *C* the percentage was still higher for using source devices first, but was closer to the chance level ($q = .50$) and the distribution looks more balanced than in the other scenarios ($q = .61$ for using a source device first compared to $q = .39$ for using an output device first). Participants in the control group, however, did not start the problem solving with a particular source device - the probabilities are

more balanced (cf. Table 2).
The transition probabilities for cable use show higher values for reusing the same cable again than using another one for most cables. E.g., the state diagram of the experimental group of scenario *C* is shown in Fig. 2. The probabilities for cable reuse are about $q \approx .5$ - even for the cinch stereo cable because "Cinch Stereo Red" and "Cinch Stereo White" only describe the differently colored cable parts of the same cable and thus the transitions between these parts $q = .45$ and $q = .67$ describe the reuse of the same cable as well.

Table 2: Transition probabilities for participants using a certain device as a first action. Probabilities q_{ij} for $X_n = i \equiv$ start state and $X_{n+1} = j \equiv$ device usage; in cells transition probabilities for the experimental group/control group are shown in comparison

Device	Scenario		
	A	*B*	*C*
Blue-ray player	.44/.38	.35/.31	.22/.19
Satellite rec.	.30/.06	.52/.19	.39/.12
Television	.04/.31	.04/.38	.17/.38
A/V receiver	.22/.25	.09/.13	.22/.31

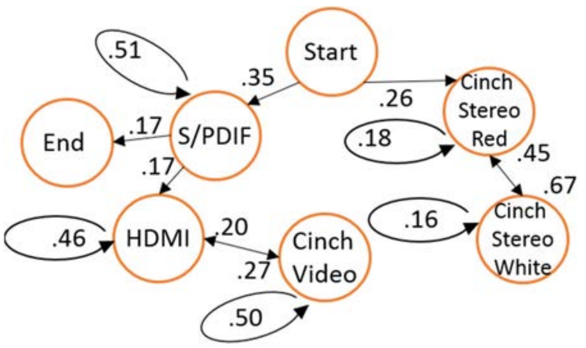

Figure 2: State diagram of cable usage in scenario *C*. The cables are depicted as states and the transitions as arrows with the probabilities attached. Only transition probabilities $q \geq .16$ are shown

Furthermore, the cable that the participants used most after the start state was the HDMI and Cinch Video Cable in scenario *A* and *B* for both groups (except in scenario *B*, where for the experimental group the Cinch Stereo was used most with $q < .22$) with a transition probability higher than $q < .25$ besides others. These were the same cables the participants already knew from the training scenario or even from existing domain knowledge (as 85% of the participants already knew the HDMI cable according to the knowledge test). Other (unknown) cables had not been used after the start state with a probability higher than $q < .15$. This is not true for scenario *C*, where in both groups the S/PDIF cable was used by $q = .35$ in the experimental group and by $q = .44$ in the control group. But this cable was known from scenario *A*. Furthermore, this cable was very easy to handle because it fitted only in two slots in the entire task (and thus only one connection is possible) whereas other cables could be used at different slots (and therefore different connections are imaginable). According to the knowledge test, only 23% of the participants knew the S/PDIF cable before the experiment and thus it was not used directly after the start state (only $q = .19$ in the control and only $q = .13$ in the experimental group) in scenario *A*, where it was introduced the first time

4 Discussion

In our study, we analyzed human planning behavior in an exploratory manner in order to derive implications for an AI planning system and to specify further research need.

4.1 Experimental Findings and Limitations

The descriptive results suggest that the participants in our experiments were rather novices and had limited knowledge of the task domain. The knowledge test showed high difficulty. The self-report and the retrospective questionnaire indicated that conclusion, too.

In the analysis of the transition probabilities it was found that people tend to reuse a cable with a higher probability than using another cable. They seem to decompose problems thus cable-wise instead of device-wise because the reuse of the same device was not found to be more probable than the usage of another device. Thus it was possible to identify a favored parameter of decomposition above others in this specific task. Furthermore, we found indicators of the influence of instruction on problem-solving behavior. First of all, a difference between the two groups could be observed. The group with the strategy instructions tended to start at source devices first compared to the control group with no strategy instructions, where device usage was more balanced. Another indicator of the influence of domain-specific knowledge was the observation that humans tended to use already known cables first before the usage of new cables. For example, in the first test scenario, the HDMI cable was used first in many cases, which was already known by most participants as the knowledge test showed. In the last scenario, the S/PDIF cable was used in many cases at the beginning, which most participants did not know before the experiment, but got familiar with through another scenario they solved before. This finding points to the opportunistic manner of human problem solving [Hay79] in knowledge-rich problem domains.

These results are certainly not applicable to other types of problems such as "creative" tasks [Orm05].

As in every experiment, there are limitations of this study, which should be taken into account in the interpretation of the results. First, the sample was homogeneous as all the participants were students and technological laypersons. Therefore, the knowledge test was also found to be too difficult and had low selectivity. Second, the sample size was small resulting in low robustness of results. Hence, the results have to be interpreted with caution. Third, the scenarios were presented in fixed order. So it is not clear if differences in human behavior or outputs between scenarios were due to differences in the scenario characteristics such as complexity, or due to order effects. Fourth, because of the exploratory nature of this study, the results are data-driven and consequently of inductive nature with post-hoc explanations of behavior rather than a strictly deductive test of a priori hypotheses. Fifth, we decided to study problem-solving in ill-defined domains and choose a specific task with a specific domain, thus generalization of these results to other domains is limited. The abstract principles need to be shown across other tasks and domains and integrated in a model of planning for deeper understanding of mechanisms.

4.2 Implications for AI Planning

In our experiment, we have observed that the participants tended to use known cables before trying different solutions. This planning behavior can be exploited in several ways. When a MIP system assists the user in solving a subproblem by presenting a set of possible solutions for it and letting the user choose, our finding can be exploited by ranking the different options, by grouping them together, or even by excluding certain options. Such strategies are especially important in case there are many options to be presented to the user, because, up to now, there are no empirically justified strategies that address the issue of presenting the available options to a user.In case the planner solves a subproblem on its own, it seems reasonable to prefer those solutions that are known to the user. That is, the planner regards solutions containing subsolutions known to the user of higher quality compared to solutions without that property. This implies that a different quality metric is optimized compared to standard planning metrics such as number of actions. Finally, in situations where the planner chooses a solution to a subproblem that is *not* known to the user although a known one might seem applicable, the system could explain why the known solutions have not been selected or are not possible. Consider the case where the planner first connects two devices with a cable known to the user, say HDMI, and then, for connecting the next two devices, it does *not* choose the other HDMI cable although this had technically been possible. Then the system could offer an explanation stating why a different cable was chosen instead of another one of the same sort. A possible reason might be that there is only one HDMI cable remaining that

has to be used somewhere else due to missing alternatives there.
The causal dependencies in the given problem allow for a wide range of possible plan linearizations, i.e. even when committing to one specific set of actions to solve the problem, there are many possibilities of how to order them. Based on the planning behavior observed in our study, we can first judge whether one of our linearization strategies given in Section 2.2 can actually be used to generate linearizations that are intuitive for human users, i.e. that adhere their expectations.
Especially Strategy 2.2 (based on *constant similarity*) seems promising: It leads to linearizations where actions including similar constants are ordered close to each other (where constants are the formal representation of objects such as cables). This enables a cable-wise completion of the overall task and thus reproduces the strategy observed in the experiments, i.e.\ it is possible to mime the human planning behavior. This enables an expectation-conforming system behavior and thus improves the system's general user-friendliness.

5 Outlook

In our experiment, we focused on planning behavior of non-creative nature. However, planning seems to be different in other domains with other characteristics [Orm05]. For example, planning behavior in more creative tasks such as planning a dinner [Beh17] could be investigated and compared. In our study, planning was analyzed without the interaction with an AI planning system. Therefore, AI systems need to be implemented and the implications evaluated. In addition, further research should be conducted to investigate under which circumstances and in which form explanations by an AI planner are needed or requested by a human when collaboratively solving a problem with a companion system. Furthermore, our approach was restricted and contextual aspects as well as multi-user situations had not been considered. These factors make the whole task more complex and dynamic. In a next step these factors should also be addressed in order to be even more realistic.

Acknowledgements
We acknowledge the support of the Transregional Collaborative Research Center SFB/TRR 62 “A Companion-Technology for Cognitive Technical Systems” funded by the German Research Foundation (DFG) and the support by the Bavarian Research Foundation.

Literature

[All96] J. F. Allen, G. Ferguson and L. K. Schubert: Planning in Complex Worlds via Mixed-Initiative Interaction. In: Advanced Planning Technology: Technological Achievements of the ARPA/Rome Laboratory Planning Initiative. AAAI Press, pp. 53-60, 1996.

[Beh17] G. Behnke, B. Leichtmann, P. Bercher, D. Höller, V. Nitsch, M. Baumann and S. Biundo: Help me make a dinner! Challenges when assisting humans in action planning. In: Proc. of the 2nd Int. Conf. on Companion Technology (ICCT). IEEE, 2017.

[Ber14] P. Bercher, S. Biundo, T. Geier, T. Hörnle, F. Nothdurft, F. Richter and B. Schattenberg: Plan, Repair, Execute, Explain – How Planning Helps to Assemble your Home Theater”. In: Proc. of ICAPS. AAAI Press, pp. 386-394, 2014.

[Ber17] P. Bercher, D. Höller, G. Behnke and S. Biundo: User-Centered Planning”. In: Companion Technology – A Paradigm Shift in HumanTechnology Interaction. Ed. by Susanne Biundo and Andreas Wendemuth. Cognitive Technologies. Springer, Chap. 5, pp. 79-100. DOI: 10.1007/978-3-319-43665-4_5, 2017.

[Ber15] P. Bercher, F. Richter, T. Hörnle, T. Geier, D. Höller, G. Behnke, F. Nothdurft, F. Honold, W. Minker, M. Weber and S. Biundo: A Planning-based Assistance System for Setting Up a Home Theater. In: Proc. of the 29th AAAI Conf. on AI (AAAI). AAAI Press, pp. 4264-4265, 2015.

[Biu16a] S. Biundo, D. Höller, B. Schattenberg and P. Bercher: “Companion-Technology: An Overview. In: Künstliche Intelligenz 30.1, pp. 11-20, DOI: 10.1007/s13218-015-0419-3, 2016.

[Biu16b] S. Biundo and A. Wendemuth: Companion-Technology for Cognitive Technical Systems. In: Künstliche Intelligenz 30.1, pp. 71-75. DOI: 10 . 1007/s13218-015-0414-8, 2016.

[Bur96] M. H. Burstein, B. Beranek, N. Inc and D. V. McDermott: Issues in the development of human-computer mixed-initiative planning. In: Cognitive Technology. Elsevier, pp. 285-303, 1996.

[Dav05] S. P. Davies: Planning and problem solving in well-defined domains. In: The cognitive psychology of planning. Ed. by Robin Morris and Geoff Ward. Psychology Press, Chap. 2, pp. 35-51, 2005.

[Dev12] Development Core Team of R. R: A Language and Environment for Statistical Computing. Vienna, Austria, 2012.

[Hay79] B. Hayes-Roth and F. Hayes-Roth: A cognitive model of planning. In: Cognitive Science 3.4, pp. 275-310, 1979.

[Höl14] D. Höller, P. Bercher, F. Richter, M. Schiller, T. Geier and S. Biundo: Finding User-friendly Linearizations of Partially Ordered Plans. In: 28th PuK Workshop ”Planen, Scheduling und Konfigurieren, Entwerfen”, 2014.

[Mor05] R. Morris and G. Ward (eds.): The cognitive psychology of planning. Psychology Press, 2005.

[Mum01] M. D. Mumford, R. A. Schultz and J. R. van Doorn: Performance in planning: Processes, requirements, and errors. In: Review of General Psychology 5.3, pp. 213-240, 2001.

[Orm05] T. C. Ormerod: Planning and ill-defined problems. In: The cognitive psychology of planning. Ed. by Robin Morris and Geoff Ward. Psychology Press, Chap. 3, pp. 53-70, 2005.

[See12] B. Seegebarth, F. Müller, B. Schattenberg and S. Biundo: Making Hybrid Plans More Clear to Human Users – A Formal Approach for Generating Sound Explanations”. In: Proc. of ICAPS. AAAI Press, pp. 225-233, 2012.

[Smi12] D. E. Smith: Planning as an Iterative Process. In: Proc. of AAAI. AAAI Press, pp. 2180-2185, 2012.

[Spe17] G. A. Spedicato: Discrete Time Markov Chains with R. In: The R Journal. R package version 0.6.9.7. URL: https://journal.r-project.org/archive/2017/RJ-2017-036/index.html, 2017.