iPass: A Case Study of the Effectiveness of Automated Planning for Decision Support

Sachin Grover* Sailik Sengupta* Tathagata Chakraborti†

Aditya Prasad Mishra* Subbarao Kambhampati*

Arizona State Univeristy

ABSTRACT

Researchers in the automated task planning community have proposed decision support systems that can assist humans in their decision-making process. Although some of these works explain the intricate details of building these systems, but their effectiveness is not supported by any human factor studies. One of the major challenge in designing these user studies, has been getting access to domain experts to verify the usefulness of the decision support system. In this paper, we borrow some of the key features of automated planning for decision support and situate them in a domain (for constructing a "plan of study" at Arizona State University) that graduate students can relate to. This allows us to perform a comprehensive study of key elements of decision support techniques using automated planning with domain experts (students) for a challenging task, thus helping us validate key elements of the decision support paradigm. We analyze the data gathered from these experiments to determine to what extent automated task planning technologies proposed in the existing literature are effective as decision support systems for human-in-the-loop decision making.

The theory of decision support is built around the idea of enabling human decision makers make decisions faster and more accurately with the added commitment to *never take the decision making away from the decision maker*. The field of automated planning [1] – which aims to develop technologies that can compute a plan or a course of action given a problem description – seems to be a perfect fit for this endeavor. However, much of the effort in this field has focused on tackling the combinatorics of the problem without much consideration of the human in the loop who may be involved in the plan generation process itself and may even be the ultimate stakeholder of the plan. Our works on human-in-the-loop planning such as [2, 3] treads these fine lines between the human as the decision maker and the automated planner as the decision support. However, until now, effective synergy between the human planner and the automated planner as conceptualized in `RADAR` [2] has not been established through human factors studies that show demonstrable improvement in the collaborative planning process through the use of automated planning technologies. The purpose of this paper is to do a case study of two key components of `RADAR` – the ability to validate a given plan for correctness and the ability to suggest a completion to a partial plan and demonstrate to what extent these components are effective for collaborative planning.

`iPass` – System Overview

We begin with a brief description of the `iPass` interface and its decision support components.[1]

The iPOS Domain and Interface One of the major difficulties of designing user studies in the decision support paradigm is access to *domain experts* who can verify the practical usefulness of the system. Thus, earlier works that propose software to help the human in their decision making process [2, 3] are unable to provide any empirical evidence as to how effective they are in practice. Thus, we

*These authors are with the School of Computing, Informatics, and Decision Systems Engineering (CIDSE) at Arizona State University, USA. Email: {sgrover6, sailiks, amishra28, rao}@asu.edu

†The author works in IBM Research, Cambridge. Email: {tathagata.chakraborti1@ibm.com}

[1]A detailed description of the interface, domain and an instance of a corresponding planning problem could not be provided due to space constrains.

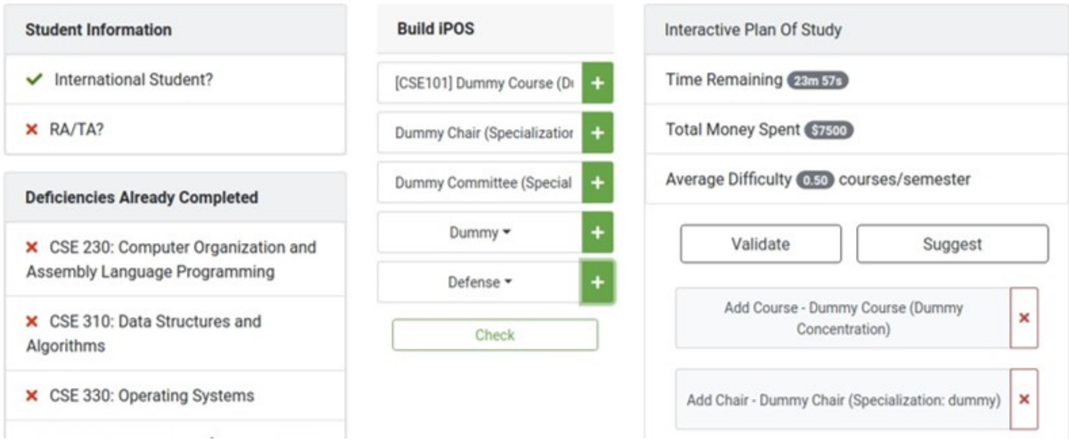

Figure 1: Illustration of the `iPass` interface.

have designed a study in the domain of constructing an *"interactive Plan of Study"* (iPOS) at Arizona State University. This has two implications. On one hand, this task is known to be challenging as per (1) evidence in existing literature [4], and (2) its use in the International Planning Competition [5] as a benchmark domain. On the other hand, graduate students are the experts in this domain, as they already maintain their iPOS as per university requirements and are easily accessible in the university. The interface (shown in Figure 1) has three panels – (1) the panel on the left shows the relevant information of the student (e.g. what deficiency courses they have, whether they are an international student, if they are research or teaching assistants, etc.); (2) the central panel provides the student with options to build the iPOS for the given student information. Actions in this panel can include adding course, specialization, committee members, etc.; (3) the panel on the right provides an interactive interface to work on the plan (such as rearranging or deletion of action) along with relevant information about the plan (e.g. difficulty or an average number of courses a semester, the total cost of tuition for the current plan, etc.). This panel also houses the decision support components that, let's the user ask for validation of the current plan or suggestions to complete it. We describe these next.

Decision Support Components We cast the iPOS design problem as a planning problem written in the popular Planning Domain Definition Language (PDDL) [6]. A planning problem consists of the initial state (which captures the student information), domain (which captures the constraints of the domain such as rules a student must follow) and the final goal (a complete plan of study). The solution to a planning problem is a sequence of actions or a *plan* (which in this case, is the plan of study).
Based on the formulation of this task as a planning problem, we use a collection of automated planning technologies to provide support during the planning process while user constructs their iPOS.
◇ *Plan Validation* – Plan validation checks the correctness of the iPOS. We use VAL [7] to check if a sub-plan (or plan) is executable in the planning domain and report validation error where needed. For example, if a user attempts to validate a plan in which they select a committee chair who is outside of the student's specialization area, VAL will catch it and provide an appropriate error message.
◇ *Plan Suggestion* – It is used to come up with a completion of a partially filled out iPOS. In order to do this, we use an existing compilation from [8], for a slightly different purpose (plan completion) than originally intended (plan recognition). The compiler takes in the plan already constructed by the student, turns them into observations that must be produced in a compiled version of the original planning problem, and then solves it to give a complete iPOS. For example, a student can choose their specialization and ask for suggestions that complete the rest of the course requirements and also select a possible committee chair that satisfies that specialization.
◇ *Plan Explanations* – When requested by the user, the planner is also capable of providing explanations for the suggested plan. It uses the technique of model reconciliation introduced in existing literature [9]. This is done by assuming an empty model of the user (i.e. a user who is not familiar with any of the constraints in the domain) and then providing a *minimal* subset of those constraints that support the suggested plan. For example, one supporting constraint for choosing a particular course may be that it is required by the selected specialization. Plan validation and explanations are complementary in nature, as explanations provide details of the domain that support a plan whereas validation points out constraints that invalidate a plan.

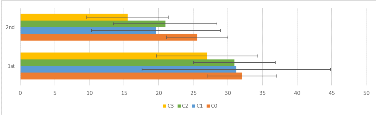

Figure 2: Average time taken (along with the standard deviation) by a participant to complete the two parts of the study for every condition.

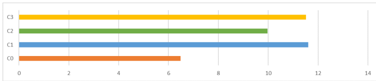

Figure 3: Time difference $\Delta T(C_i)$ between two tasks C_i^1 and C_i^2 of iPOS planning for every condition C_i.

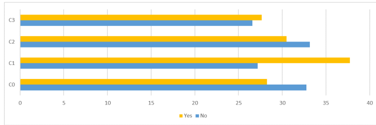

Figure 4: Time taken by experienced (in yellow) and non-experienced (in blue) users to make the first iPOS (C_i^1).

Aim of the Study & Results

In order to determine the individual as well as the cumulative impact of the two decision support components, we evaluated our interface in four conditions – [C_0] Both validation and suggestion capabilities are absent. The users do have to pass correctness before they can submit, [C_1] Only validation capability is enabled, [C_2] Only suggestion capability is enabled, [C_3] Both validation and suggestion options are available. Furthermore, each participant assigned to one of the study conditions C_i performed the iPOS planning task twice (with different, randomly generated, student information). We thus, have two sub-conditions (denoted using the super-script) C_i^1 and C_i^2 for every condition i.
The study was conducted on the university premises. Each subject was given \$15 for an hour of study when they used `iPass` software to make two iPOS. At the start of the study, participants were informed that they would be asked to explain each iPOS with the hope that it will help them be more invested in the task [10]. Then they were given a document explaining the planning domain and another document explaining the functionality of the elements in the interface. They were given 20 minutes to make each iPOS, after which they were presented with a feedback form. We had a total of 56 participants, out of which six were undergraduates and the rest were graduate students. A total of 18 participants had submitted an iPOS before. The participants were evenly distributed among the four study conditions. Given the setup, we now present the three hypothesis and the results for each one of them.

Hypothesis H1: Time to complete an iPOS and satisfaction about finalized plan follow a particular order. **H1a.** – Time to complete an iPOS follows the order $T(C_0) > T(C_1), T(C_2) > T(C_3)$, where $T(C_i)$ represents the time to complete the iPOS in condition i. We show the average time a participant took to complete the first and the second iPOS and submit their feedback[2] in Figure 2. The data shows a significant improvement in performance with regards to time as one goes from C_0 to C_3 ($p < 0.05$ for the first and $p < 0.01$ for the second iPOS) showing that the automated planning technologies in conjunction helped in improving the efficiency of the decision making process. Unfortunately, there was no significant improvement seen in performance from (1) C_0 to C_1 or C_2 and (2) C_1 or C_2 to C_3. Thus, *hypothesis H1a was found to be partially true*, thereby showing that all the planning technologies and not a subset of them were necessary to improve the planning performance for the user.
H1b. – User satisfaction for the constructed iPOS will follow the order $S(C_0) < S(C_1), S(C_2) < S(C_3)$, where $S(C_i)$ represents the degree of satisfaction as provided by the user for subjective questions. In Figure 6, we show the answers of the users to the subjective statement *Q3: I am happy with the final Plan of Study* on the Likert Scale for all the four conditions. In C_0, we noticed that the least number of users agreed (either agreed or strongly agreed) with the statement. This is not surprising because many users were not even able to come up with a valid plan of study without any planning support in C_0. For C_1, six participants said they were in unison with the statement Q3 and For C_2 and C_3, half of the participants were happier (i.e. either agreed or strongly agreed) with their plan of study, which is the highest across all the four conditions. But, in C_2 there was one participant who strongly disagreed with the statement, while for C_3 there were none. Thus, *the hypothesis H1b holds.*
H1c. – User satisfaction with the feedback from the interface will follow the same order as H1b. In Figure 7, we show the number of users who agreed with the ratings on the Likert Scale for the statement *Q2: The feedback from the interface helped the iPOS making process.* If we let n_{C_i} denote the number of participants who either agree or strongly agree with the statement, then the following relation holds, $n_{C_0} < n_{C_1}, n_{C_2} \leq n_{C_3}$. Although the equality holds n_{C_1} and n_{C_3}, the number of people who strongly agreed to the statement was, by far, the highest for C_3. Thus, we infer that *the hypothesis H1c holds.*

[2] Since feedback was part of all the conditions, this is indicative of, even though not the actual, planning time.

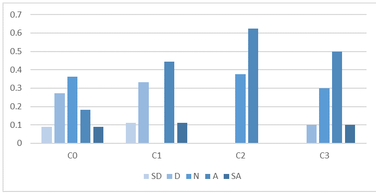

Figure 5: Feedback of non-experienced users about the statement 'Q1: The planning task was pretty simple for me' for each condition C_i^1.

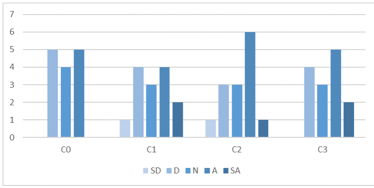

Figure 6: Average score for subjective 'Q3: I am happy with the final iPOS' for conditions C_i^1.

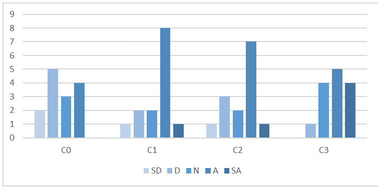

Figure 7: User agreement metrics for the statement 'Q2: The feedback from the interface helped the iPOS making process' for each condition C_i^1.

Hypothesis H2: Time to complete the plan will reduce in the second attempt We plot the average decrease in time in completing the second iPOS after doing the first iPOS with `iPass` for all the four study conditions in Figure 3. The lowest reduction in time for C_0 shows that feedback given to the user by the decision support system helps them learn more about the domain model, thereby improving their performance in making the second iPOS. We also saw that the highest reduction in time occurred for the conditions C_1 $(p < 0.1)$ and C_3 $(p < 0.01)$. We feel that the presence of plan validation in both these conditions informed the users about the reason behind each error they made while constructing the first iPOS that was effective in teaching the users about the actual domain. Due to a similar reason, we had also hypothesized that the presence of plan explanations in C_2 and C_3 will reduce the time significantly because these explanations will teach the user about the domain, thus reconciling the models. Unfortunately, this functionality was used very rarely $(0.14$ and 0.91 average number of times for C_2 and $C_3)$ and thus, improvement in performance was not observed. Hence, *H2 was only found to be partially true*, supporting the cause that use of automated planning C_3 for decision support improved the efficiency of the human thereby reduced the time for making the second iPOS.

Hypothesis H3: Less expert users benefit more from decision support components We noticed that the performance (time) was *not significantly better* for participants who had filled an iPOS before when compared to participants with no experience (Figure 4). Although the experienced participants did perform slightly better in C_0, C_1 and C_3, to our surprise, we noticed that for C_2, the users who had no prior experience performed better. This might be because the latter group had prior conceptions about the rules of making an iPOS and thus, spent time making plans that appeared valid *in their model*, but were invalid in the `iPass` domain. With the presence of 'validate' in C_1, they might have ended up having to correct their partial plans multiple times, resulting in a longer time and worse performance.
We plot the response of non-experienced users to the subjective question *Q1: The planning task was pretty simple for me* in Figure 5. Interestingly, the non-experienced users seemed to agree (or strongly agree) more with the statement in C_3 compared to C_0, indicating that support features have contributed to decrease in perceived difficulty of the task.

Conclusion

In summary, we found that two key decision support components – validation and suggestion – for human-in-the-loop planning tasks were, in general, useful in improving the performance and/or satisfaction of the decision-maker. From the written feedback, we noticed that 11 people asked for more feedback from the interface in C_0 (3 of whom mentioned suggestion feedback and 5 mentioned validation feedback) thus highlighting the role of the evaluated support components in the normative expectations of the user. These results provide partial validation about the effectiveness of using automated planning-based decision support systems in expert driven domains.

Acknowledgements

This research is supported in part by the ONR grants N00014-16-1-2892, N00014-18-1-2442, N00014-18-1-2840, the AFOSR grant FA9550-18-1-0067, and the NASA grant NNX17AD06G.

References

[1] Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated Planning: Theory and Practice*. Elsevier, 2004.

[2] Sailik Sengupta, Tathagata Chakraborti, Sarath Sreedharan, and Subbarao Kambhampati. RADAR - A Proactive Decision Support System for Human-in-the-Loop Planning. In *ICAPS Workshop on User Interfaces for Scheduling and Planning (UISP)*, 2017.

[3] Sailik Sengupta, Tathagata Chakraborti, and Subbarao Kambhampati. Ma-radar-a mixed-reality interface for collaborative decision making. *ICAPS UISP*, 2018.

[4] Omar Zia Khan, Pascal Poupart, and James P Black. Automatically Generated Explanations for Markov Decision Processes. In *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions*, pages 144–163. IGI Global, 2012.

[5] International Planning Competition (IPC) Probabilistic Track. Academic Advising Domain, 2018. https://ipc2018-probabilistic.bitbucket.io/#domains.

[6] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. PDDL –The Planning Domain Definition Language. 1998.

[7] R. Howey, D. Long, and M. Fox. VAL: Automatic Plan validation, Continuous Effects and Mixed Initiative Planning Using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2004)*, 2004.

[8] Miquel Ramírez and Hector Geffner. Plan recognition as planning. In *IJCAI*, 2009.

[9] Tathagata Chakraborti, Sarath Sreedharan, Yu Zhang, and Subbarao Kambhampati. Plan Explanations as Model Reconciliation: Moving Beyond Explanation as Soliloquy. In *IJCAI*, 2017.

[10] Hugo Mercier and Dan Sperber. Why do humans reason? arguments for an argumentative theory. *Behavioral and brain sciences*, 34(2):57–74, 2011.