

Key Issues in Interactive Problem Solving: An Empirical Investigation on Users Attitude*

Gabriella Cortellessa, Vittoria Giuliani,
Massimiliano Scopelliti, and Amedeo Cesta

ISTC-CNR, Institute for Cognitive Science and Technology,
Italian National Research Council, I-00137 Rome, Italy
{name.surname}@istc.cnr.it

Abstract. This paper explores the interaction between human and artificial problem solvers when interacting with an Intelligent Scheduling System. An experimental study is presented aimed at investigating the users' attitude towards two alternative strategies for solving scheduling problems: automated and interactive. According to an automated strategy the responsibility of solving the problem is delegated to the artificial solver, while according to an interactive strategy human and automated solvers cooperate to achieve a problem solution.

Previous observations of end-users' reactions to problem solving systems have shown that users are often skeptical toward artificial solver performance and prefer to keep the control of the problem solving process. The current study aims at understanding the role played by both the users' expertise and the difficulty of the problem in choosing one of the two strategies. Results show that user expertise and task difficulty interact in influencing this choice.

A second aspect explored in the paper concerns the context in which the end-users rely on explanations to understand the solving process. Explanations are in fact expected to play an important role when artificial systems are used for cooperative and interactive problem solving. Results support the hypothesis that explanation services are more often called into play in case of problem solving failures.

1 Introduction

The introduction of intelligent systems for solving complex problems in real world domain has been characterized by the raising consciousness that in most cases a completely automated approach is neither applicable nor suitable for a successful deployment of solving technologies. Traditionally, the main concern of scholars in the problem solving field has been to develop efficient and powerful algorithms for finding solutions to complex problems. This is true for both Artificial Intelligence and Operations Research approaches. A usually neglected issue has been the lack of effective front end design through which an end user can benefit from the potentialities of the artificial tools, taking, at the same time, an active role in the resolution process. In fact, on the

* This research is partially supported by MIUR (Italian Ministry of Education, University and Research) under project ROBOCARE (A Multi-Agent System with Intelligent Fixed and Mobile Robotic Components).

one hand the development of automated algorithms to solve difficult problems can relieve a user from complex or tedious tasks, on the other it needs to be combined with the design of effective and usable instruments for the users to maintain a level of control over the solving phase.

An integrated system *(human, artificial solver)* would create a powerful and enhanced problem solver applicable to the resolution of difficult real world problems. In this light, a new solving paradigm has been proposed in the literature, which promotes a collaboration between human and artificial solvers during the problem solving process. This emerging paradigm is known as mixed-initiative approach [2,4] and fosters the human-computer cooperation during the resolution of complex problems. The approach is based on the idea that experienced users and automated technologies bring complementary problem-solving strengths to the table, and the goal is to synergistically blend these combined strengths. Hence, successful technology application requires effective integration of decision-making and representation models of the user and the system. However, several problems, stemming from the composite nature of the system user-machine, need to be addressed.

This work investigates the interaction between human beings and computer-based decision support systems for the resolution of difficult problems. In particular we consider the case of an intelligent scheduling system that is endowed with different levels of interactivity with the user.

The design of mixed-initiative planning and scheduling systems has mainly to do with the general problem of interfacing the user with the automated system and placing him/her into the problem-solving loop. One basic task is to show system outcomes in user-comprehensible terms to save users the cognitive burden of understanding the internal representation usually inherent to specialized solving techniques. An additional feature, very relevant for human problem solving, is the ability to provide explanatory support in different problem solving phases. Unfortunately, most current systems provide no guidance in such cases, and force the user to diagnose the problem using the system's internal model. The ability of the system to provide user-level rationalizations of generated solutions may promote user acceptance [21]. Current proposals for mixed-initiative systems (see [2] for example) are very often presented as system description, whereas less work has been devoted to understanding how it is possible to evaluate the utility of the whole approach and the different features. This paper offers a methodological contribution in this direction, applying an experimental approach to the problem of understanding users' attitude toward interactive problem solving features.

Two main issues will be considered, concerning (a) users' preferences in selecting the solving strategy and (b) their willingness to rely on explanation. It is worth emphasizing that while several works on the mixed-initiative paradigm claim that end-users of automated systems prefer to maintain control over the problem solving, thus appreciating mixed-initiative systems, nonetheless, no empirical evidence is given to support this statement.

Plan of the Paper. The paper is organized as follows. The reference problems are introduced and the aims of the study are spelled out in Section 2. Section 3 describes the methodology used in the experiment. The results of the research are illustrated in Section 4. A general discussion and some final conclusions end the paper.

2 Preface to the Study

Before describing the methodology and results of the study we first introduce some definitions and notions that prepare the ground to better understand the context of the work. In particular this section introduces the type of problems we will consider in the study, then briefly describes an Artificial Problem Solver by which this work has been inspired and outline the main aims of the research.

2.1 Planning and Scheduling Problems

For the purpose of this paper we shall first give a definition of the reference problems which will be used in the study. The problems we are referring to, are known in literature as *planning* and *scheduling* problems. Planning and scheduling problems are present in many different applicative contexts, and are the object of many applicative computer-based systems. Considering that in most cases a clear distinction between the two categories is very difficult, definitions aiming at highlighting some crucial differences are given. While *planning* is concerned with the problem of deciding *what* to do – which actions are to be executed– to reach some *goals*, without any particular attention to the resources needed to execute the actions, the main concern of *scheduling* is to determine when to execute the actions and which resources are to be used to complete the whole set of actions. In other words planning consists of synthesizing a sequence of actions that enable us to reach a desired final state of the world, starting from a given initial state. The sequence of actions is called *plan* and it allows to reach a goal. Scheduling specializes the plan adding information about *time* and resources needed to execute the actions. A scheduling problem, see for example [1], consists of a set of activities, a set of resources and a set of temporal and resource constraints the activities should satisfy in order to be correctly executed. A solution of such problem is a *schedule*, that is an assignment to the start time of the activities such that all the specified constraints are satisfied.

2.2 COMIREM: A Reference Planning/Scheduling System

The mixed-initiative system we have used as a reference for our work is COMIREM [19], a web-based mixed-initiative problem solver devoted to the resolution of planning/scheduling problems.

The system is based on a CSP paradigm [20] and promotes a problem solving process that combines the actions of the automated solver and the human planner. COMIREM is composed of two main modules, named respectively *Automated Solver* and *Interaction Module*. The first is devoted to modeling domain entities through a CSP representation and provides the *automated* problem solving strategy. The Interaction Module directly interacts with the user, and allows him/her to take part in the process of finding a solution via advanced interactive facilities. It represents the communication channel between the user and the automated solver and a means to exploit various features of the automated system. Due to its interactive nature, the system can exploit human-planner knowledge and decision making, and in fact promotes a mixed-initiative problem solving process. Through the Interaction Module a user can choose

to either generate a solution automatically or iteratively build one through a *step by step* mixed-initiative procedure that interleaves human choices with system calculation of consequences. In the first case a user decides to completely delegate the system, in the second case he/she takes an active role in the problem solving. In accordance with the mixed-initiative theory, the ambitious idea behind COMIREM is to capture the different skills that a user and an automated system can apply to the resolution process, by providing both powerful automatic algorithms to efficiently solve problems and interactive facilities to keep the human solvers in the loop. For the purpose of our study we developed a simplified and simulated version of COMIREM, which is devoted to solve scheduling problem instances in a TV broadcasting station domain.

2.3 Objectives and Hypothesis

The aim of this study is twofold, addressing both user strategy preferences and use of explanations.

Empirical Investigation on Users' Solving Strategy Selection. Previous research on human problem solving procedures has mainly addressed either the influence of the type of problem [14] or the users' assessment of different strategies (e. g., how successful different strategies are or how costly they are to apply), focusing on how success/failure resulting from previous trials influence solving processes [13]. This work aims at investigating the influence of both human solvers' expertise and problem difficulty on the selection of interactive vs automated strategy. In addition, contrary to traditional research in the field of problem solving, where the object of study has usually been mathematical or logic problems, in this case "real world" problems will be considered. In our research, the user is presented an alternative between a completely automated procedure and an interactive approach. By choosing the first alternative, the user will delegate each action to the artificial solver, thus keeping no control over the problem solving process, whereas in the second case the system and the human solver will actively cooperate to produce a solution to the problem.

There is some evidence that humans do not always adopt an optimal strategy in getting help from artificial tools, ignoring advices or solutions proposed by the system [10]. A possible explanation for this behavior is provided by some research in the field of human-computers interaction, reporting that humans tend to attribute a certain degree of anthropomorphism to computers, assigning to them human traits and characteristics. In [12,16] a series of experimental studies are reviewed, showing that individuals mindlessly apply social rules and expectations to computers. It is plausible to hypothesize that human problem solvers show the same tendency toward artificial solvers, and refuse to delegate the solution of the problem, for many reasons. For instance, they could mistrust the automated agent's ability to solve the problem or they could enter in competition with it. However, we have no data on possible differences in the behavior of users with different levels of expertise. Experts are people with some knowledge of the design of artificial solvers and they are aware of the limits and merits of the system. We assume they would adopt a more pragmatic strategy, thus delegating the machine to solve the problem in order not to waste time. On the other hand they may be interested in understanding the procedure applied by the system. Hence, when facing

difficult tasks, they might be motivated to test themselves and actively take part in the process. Conversely, non-experts do not know the mechanisms behind the automated algorithms and thus might have a different degree of trust. Nonetheless the greater the difficulty of the problems, the more likely the choice to commit the problem solving to the machine. For these reasons we believe that some differences might exist between experts and non-experts while interacting with an artificial problem solver. In particular we formulate the following hypotheses:

Hypothesis 1. *Solving strategy selection (automated vs interactive) depends upon user expertise. In particular it is expected that scheduling expert users use the automated procedure more than non-experts. Conversely, non-expert users are expected to use the mixed-initiative approach more than experts.*

Hypothesis 1b. *In addition it is expected that when solving easy problems, inexperienced users prefer the mixed-initiative approach, while expert users have a preference for the automated strategy. Conversely, for solving difficult problems inexperienced users may prefer the automated strategy while expert users have a tendency to choose the mixed- initiative approach.*

Explanation Recourse. The term explanation usually means “to make something plain or understandable”, to “give the reason for some event, be it expected or not”. In AI, explanation has been investigated from a variety of different perspectives.

The traditional notion of explanation in expert systems is based on the early work on the MYCIN system [5] developed at Stanford in the 1970s to diagnose and recommend treatment in the medical field. In that case, explanation was devoted to answering “how” and “why” questions, based on a trace of the system’s reasoning. In knowledge-based system explanations facilities have been designed to teach and clarify systems’ intentions or to convince the users of the reliability of system results [8].

Explanations, by virtue of making the performance of a system transparent to its users, are influential for user acceptance of intelligent systems and for improving users’ trust in the advice provided [9]. In systems where explanation is provided as a part of problem solving, it is important to consider how it relates to the problem solving process. In fact it is possible that explanations can either be based on problem solving, in which case they elucidate reasoning and problem solving methods, or be used as a dialogue with the user to allow him/her to influence the problem solving progress. This second alternative is appropriate in those contexts in which the user and the system are collaborating to solve a problem. In these cases explanations are seen as an effective means to cooperate in the problem solving process [11]. In previous research [3] expectation of failures and perceived anomalies have been identified as an occasion for accessing explanations [6,17]. According to these results we formulate the following hypotheses:

Hypothesis 2. *The access to explanation is more frequent in case of failure than in case of success.*

Hypothesis 2b. *The access to explanation is positively associated with the number of failures and negatively associated with the number of successes.*

In the context of knowledge-based systems, the role of explanations in cooperative problem solving has been investigated [7] and results show that participants in cooperative problem solving conditions, made a greater use of explanations. In accordance with the Mixed-Initiative Theory we hypothesize that the human solver, actively participating in the problem solving, possesses a higher level of control in the problem solving, thus showing a lower need to access the explanation. In particular we formulate the following hypothesis:

Hypothesis 3. *Access to explanation is related to the solving strategy selection. In particular participants who choose the automated solving strategy access more frequently the explanation than subjects who use the mixed-initiative approach.*

In addition to the hypotheses expressed above, we want to investigate the relationships between the user's level of expertise and the use of explanation. Evidence has been provided that experts are more likely to use explanations for resolving perceived anomalies [15] or because they are surprised by conclusions [21]. In accordance with cognitive learning theory, we expect that non experts will use explanation more frequently than experts when solving problems with automated systems. In particular we formulate the following hypothesis.

Hypothesis 4. *During problem solving non experts access explanations more frequently than experts.*

Finally we also expect that the difficulty of problems will affect the recourse to explanation. In particular we hypothesize that:

Hypothesis 5. *Access to explanation is more frequent in case of difficult problems.*

Based on the taxonomy of explanation types provided in [3] the explanation we will use in our experiment will have the following features: it aims at explaining problem solvers' choices, it is expressed in textual form and has a user-invoked provision mechanism.

3 Method

The general experimental design of this research aims at investigating the influence of the variables *expertise* and *problem difficulty* on the solving strategy selection and access to explanation. The variable *expertise* is a *between* factor with two levels, expert or non-expert, while the *problem difficulty* represents a *within* factor with two levels, low and high¹. A further independent variable is represented by *failure* during the problem

¹ The two levels of this variable have been determined considering the problems dimension in terms of number of activities to be scheduled and alternative resources available for each activity.

solving. This last variable has two levels, present or absent. As general measures, the choice of the solving strategy and the frequency of access to explanation have been considered. With respect to the solving strategy, two general scores were computed (*choice_auto* and *choice_mixed*). They measure the overall frequency of choice of each strategy in the experiment.

As regards the access to explanation the following indexes were calculated:

- *access_failure* which represents the frequency of access to explanation in case of failure during problem solving;
- *access_success* which measures the frequency of access to explanation in case of correct decision during problem solving;
- *access_low_difficulty* indicating the frequency of access to explanation in case of problems of low difficulty;
- *access_high_difficulty* indicating the frequency of access to explanation in case of problems of high difficulty.

3.1 Tools

A web-based software has been developed, inspired by the software COMIREM described in 2.2. The simulated system allows users to solve instances of scheduling problems by means of two alternative procedures, automated and mixed-initiative. The system is accessible through a web browser and is organized as follows:

- *Presentation*: A general description of the study and the list of software requirements.
- *User data input form*: Data collected through this input form were registered in a data base implemented in MySQL Language. For each participant the following data were registered: identifier, profession, education, sex, age, language, expertise in planning & scheduling and participant's problem solving pattern.
- *Instructions*: A list of instructions to be followed during the experiment.
- *Training session*: This session was implemented through a sequence of animated web pages showing the actions necessary to use the system. The layout of the screen has been subdivided into two parts. On the left part the list of instructions was presented, which described the interface of the system and called upon the users to actively use the system. The right part of the screen was devoted to presenting the Problem Solver and its behavior consequently to user actions. The training session also allowed users to use and practice the system.
- *Session 1*: It was implemented through a sequence of web pages showing an instance of a scheduling problem to be solved. A textual description of the problem was shown, followed by a graphical presentation. Consequently to the user's actions, the system showed updated results.
- *Questionnaire 1*: an 11-item questionnaire was presented at the end of the first session. The questionnaire was subdivided into three sections:
 1. the first section was devoted to the *manipulation check* of the variable *difficulty*;
 2. the second section was devoted to verifying how clear the two description modalities (textual and graphic) were;

3. the last section aimed at investigating users' strategy selections and the reasons for their choices.

The first two sections included 6 items on a 5-step Likert type response scale (from "not at all" to "very much"). For the remaining items, related to reasons for the strategy selection, participants were asked to choose among different options. Participants were given the possibility to indicate possible suggestions or general comments.

- *Session 2*: It was implemented through a sequence of web pages showing the instance of a scheduling problem to be solved.
- *Questionnaire 2*: The first three sections were the same as for questionnaire 1. In addition a fourth session was added designed for investigating the access to explanations during the experiment and their perceived utilities. Questions related to explanations were evaluated on a 5-step item Likert scale.

3.2 Participants and Procedure

A group of 46 subjects was contacted, aged from 23 to 58 years (Mean 33,3). The sample was balanced with respect to expertise in planning and scheduling (23 experts and 23 non experts) and with respect to gender, education, age and profession.

All subjects participated in the experiment by connecting from their own computer to the experiment web site².

At the beginning of the experiment, an animated tutorial provided subjects with instructions on how to use the software, and showed which type of problems were to be solved. Then, it solved an example of scheduling problems by using both the automated and the mixed-initiative procedure. Participants could repeat the tutorial session until they felt confident with the use of the system. Then a problem was presented to the subjects and they were asked to choose between one of the two available solving strategies. During the problem solving, participants could either access explanations through the *explanation* button or go to the next step. User's interactions with the system were registered in the data base. At the end of the first session subjects were asked to fill in Questionnaire 1. The same procedure was followed for session 2. In order to avoid effects due to the order of the presentation, the two sessions (which corresponded to different degrees of difficulty³) were randomly presented to the users.

3.3 Stimuli

Four scheduling problems were defined in the field of a broadcast TV station resources management. Two solvable problems (1 low difficulty and 1 high difficulty) were presented during the first and the second session to all subjects, and two unsolvable problems (1 low difficulty and 1 high difficulty) were presented only to subjects who chose the automated procedure. The reason for adding these further problems in case of automated selection is twofold:

² <http://pst2.istc.cnr.it/experiment/>

³ Results of a preliminary analysis showed that our classification of difficult and easy problems coincides with the one perceived by the participants.

- the mixed-initiative selection entailed more time to solve problems. In this way all subjects had a comparable workload in term of time spent in solving problems.
- the mixed-initiative selection entailed that almost all participants encountered some failures during the problem solving, thus introducing unsolvable instances (failure) which were also necessary to the automated procedure.

4 Results

4.1 Solving Strategy Selection

The first analysis investigated the influence of *expertise* on the solving strategy selection. A between subjects ANOVA was performed to test Hypothesis 1, separately for the two different strategies. The dependent variables used in this analysis were the indexes *choice_auto* and *choice_mixed* respectively (see Section 3). With respect to the strategy selection no significant difference was found ($F(1,22)=1.94$, n. s.).

To test Hypothesis 1b a χ^2 test was performed, separately for low and high level of difficulty. In case of low difficulty problems, a significant effect was found ($\chi^2=5.58$, $df=1$, $p < .05$). In particular the analysis of standardized residual shows that when solving easy problems, experts prefer the automated strategy, while non-experts prefer the mixed-initiative approach (see Fig. 1).

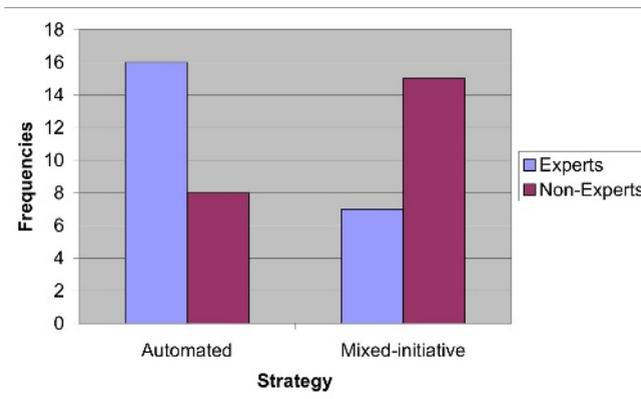


Fig. 1. Strategy selection preferences: easy problems

The analysis shows no significant difference between the two groups in case of difficult problems ($\chi^2=0.11$, $df=1$, n.s.) (see Fig. 2).

4.2 Access to Explanation

To test Hypothesis 2 which aimed at investigating the relationship between failures and access to explanation, a repeated-measures ANOVA was performed using as dependent

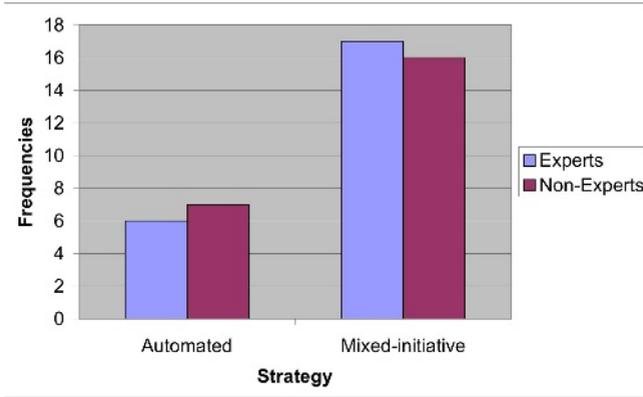


Fig. 2. Strategy selection preferences: difficult problems

variables the indexes *access_failure* and *access_correct*, introduced in Section 3. Results show a significant effect of failure on the access to explanation ($F(1,41)=24.32, p < .001$). In particular users rely on explanation more frequently in case of failure than in case of success (see Table 1).

Moreover, a correlation analysis between number of failures (and successes) and number of accesses to explanation was performed in order to test Hypothesis 2b. Results show a significant correlation between failures and number of accesses to explanation ($r=0.623, p < .001$). Conversely there is no significant correlation between number of correct choices and number of accesses to explanation ($r=.01, n.s.$).

Table 1. Access to explanation: failure and success (statistics)

	N	Mean	Std. Deviation
<i>access_failure</i>	42	.6794	.4352
<i>access_success</i>	42	.3141	.3174

To test the relationship between the selected strategy and the access to explanation (Hypothesis 3), an ANOVA for independent groups was performed separately for the two levels of difficulty. The indexes *access_low_difficulty* and *access_high_difficulty*, defined in Section 3, were used as dependent variables. A significant effect of the strategy selection on the recourse to explanation was found. In particular the access to explanation is higher when the automated strategy is chosen both in case of easy problems ($F(1,43)=67.22, p < .001$), see Table2, and in case of difficult problems ($F(1,44)=10.97, p < .05$), see Table 3.

Finally, to test Hypothesis 4 and 5 a mixed-design ANOVA was performed considering *expertise* as a between-subjects factor and *difficulty* as a within-subjects factor. The indexes *access_low_difficulty*, and *access_high_difficulty* were used as dependent variables. Results show a significant effect of expertise on recourse to explanation

Table 2. Index of access to explanation: easy problems

	N	Mean	Std. Deviation
<i>automated</i>	23	.8043	.2915
<i>mixed-initiative</i>	22	.1667	.2242
<i>totale</i>	45	.4926	.4128

Table 3. Index of access to explanation: difficult problems

	N	Mean	Std. Deviation
<i>automated</i>	13	.5769	.3444
<i>mixed-initiative</i>	33	.2625	.2666
<i>total</i>	45	.3513	.3204

($F(1,43)=4.76, p < .05$). Experts were found to access explanation significantly more than non-experts. An effect of problem difficulty on the recourse to explanation was also found ($F(1,43)=6.48, p < .05$). Access to explanation was found to be significantly higher when an easy problem is to be solved. No significant interaction effect was found ($F(1,43)=0.89, n.s.$) (see Table 4).

Table 4. Index of access to explanation: effect of expertise and problem difficulty

	<i>expertise</i>	N	Mean	Std. Deviation
<i>access_low_difficulty</i>	Non-experts	23	.3696	.3407
	Experts	22	.6212	.4489
	Totale	45	.4926	.4128
<i>access_high_difficulty</i>	Non-experts	23	.2768	.2994
	Experts	22	.4193	.3360
	Total	45	.3465	.3223

5 Discussion

The overall results of the present research are consistent with the expectation that non-expert users prefer the mixed-initiative approach rather than the automated strategy, while experts rely more frequently on the automated strategy. Moreover, the explanation is frequently used and the frequency of access is higher in case of failure than in case of success. More specifically, the study showed that non experts prefer the mixed-initiative procedure independently from the problem level of difficulty. Conversely, experts prefer the automated strategy when solving easy problems, while tend to move to the mixed-initiative approach while solving difficult problems.

As expected non-expert users show a tendency to actively solve problems keeping control over the problem solving process. This result can be considered in accordance

with the idea that non-experts tend to be skeptical toward the use of an automated system, probably because they do not completely trust the solver capabilities. One possible explanation, consistent with [16], refers to the tendency to anthropomorphise machines and to believe that they can make mistakes just like human beings.

Conversely, expert users show a higher trust toward the automated solver. Nonetheless, they find it stimulating to actively participate in the problem solving process when a difficult task is given. The relevance of the problem difficulty emerged to be a key variable in their choice. Expert users are usually system designers and are used to implementing automatic algorithms, thus knowing how effective machines can be in solving problems. When an easy task is to be solved they are likely to consider the mixed-initiative approach as a time-wasting choice. On the other hand the idea of facing a puzzling problem can drive them to conceive alternative methods to generate solutions.

Results confirmed previous studies [6,17] according to which access to explanation is more frequent in case of failure and the main reason for accessing explanation seems to be the will to understand the artificial solver. Interestingly we found that, as expected, the more the failures the more the accesses to explanation; on the other hand no relationship was found between success and access to explanation. As a consequence it is possible to assert that success is not predictive of any specific behavior with respect to access to explanation.

In contrast with what was found in [7], our Hypothesis 3 asserting a greater use of explanation in case of automated solving strategy selection was confirmed. However a key difference between the two studies needs to be highlighted. In this contribution participants were asked to choose their preferred strategies while in [7] the automated vs. interactive strategy was treated as a between-subjects condition. In both sessions of our experiment it was found that participants who choose the automated strategy access explanation more frequently than subjects who chose the mixed-initiative approach. It is possible to speculate that by selecting the mixed-initiative approach, subjects actively participate in the problem solving and keep a higher control on the solving process. As a consequence the need for explanation might decrease. Conversely, participants who chose the automated strategy delegate the artificial solver but at the same time they need to understand solvers's choices and decisions. A somewhat surprising finding of the study was that experts access explanation more frequently than non experts; in addition the access to explanation is more frequent when facing an easy problem than in case of a difficult problem.

6 Conclusions

This paper introduces an experimental approach to evaluate key features of mixed-initiative systems to problem solving.

In particular we have given attention to basic users attitude concerning the choice of automated rather than interactive strategies and the bias toward the use of explanation. As a result, we have empirically proved that the mixed initiative approach responds to the willingness of end users to keep control over automated systems. Additionally, evidence has been found that non expert users prefer to adopt a mixed-initiative approach rather than delegate the automated system to solve a problem. Conversely, expert users

prefer to entrust the system with the task of problem solving. The existing difference between individuals with different levels of expertise highlights the need for different styles of interaction in the development of intelligent problem solving systems.

It was also demonstrated the utility of explanation during problem solving, and the achievement of a *failure* state has been identified as a main prompt to increase the frequency of explanation access.

Several points remain open for future investigation, in particular we would like to use the same experimental apparatus to evaluate different types and depths of explanation. Additionally we are interested in the automated synthesis of explanation from the internal representation of the intelligent scheduling system. (see [18] for some preliminary results).

References

1. P. Baptiste, C. Le Pape, and W. Nuijten. *Constraint-Based Scheduling*, volume 39 of *International Series in Operations Research and Management Science*. Kluwer Academic Publishers, 2001.
2. Mark Burstein and Drew McDermott. Issues in the development of human-computer mixed-initiative planning. In B. Gorayska and J.L. Mey, editors, *Cognitive Technology*, pages 285–303. Elsevier, 1996.
3. B. Chandrasekaran and Sanjay Mittal. Deep versus compiled knowledge approaches to diagnostic problem-solving. *Int. J. Hum.-Comput. Stud.*, 51(2):357–368, 1999.
4. R. Cohen, C. Allaby, C. Cumbaa, M. Fitzgerald, K. Ho, B. Hui, C. Latulipe, F. Lu, N. Moussa, D. Pooley, A. Qian, and S. Siddiqi. What is initiative? In S. Haller, S. McRoy, and A. Kobsa, editors, *Computational Models of Mixed-Initiative Interaction*, pages 171–212. Kluwer Academic Publishers, 1999.
5. R. Davis. Application of Meta-Level Knowledge to the Construction, maintenance and Use of Large Knowledge bases. In R. Davis and D. Lenat, editors, *Knowledge-Based Systems in Artificial Intelligence*. McGraw-Hill, Inc., NY, 1982.
6. N. Gilbert. Explanation and dialogue. *Knowledge Engineering Review*, 4(3):205–231, 1989.
7. S. Gregor. Explanations from knowledge-based systems and cooperative problem solving: an empirical study. *International Journal of Human-Computer Studies*, 54:81–105, 2001.
8. P. Hayes and R. Reddy. Steps toward Graceful Interaction in Spoken and Written Man-Machine Communication. *Internal Journal of Man-Machines Studies*, 19:231–284, 1983.
9. F. Hayes-Roth and N. Jacobstein. The state of knowledge-based systems. *Communications of the ACM*, 37:27–39, 1994.
10. D. R. Jones and D. Brown. The division of labor between human and computer in the presence of decision support system advice. *Decision Support Systems*, 33:375–388, 2002.
11. L. Karsenty and J. Brezillon. Cooperative problem solving and explanation. *Expert Systems with Applications*, 8(4):445–462, 1995.
12. E. J. Langer. Matters of mind: Mindfulness/mindlessness in perspective. *Consciousness and Cognition*, 1:289–305, 1992.
13. M. C. Lovett and J. R. Anderson. Making heads or tails out of selecting problem-solving strategies. In J. D. Moore and J. F. Lehman, editors, *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society, Hillsdale, NJ*, pages 265–270. Erlbaum, 1995.
14. P. G. Polson. M. E. Atwood, M. E. Masson. Further explorations with a process model for water jug problems. *Memory and Cognition*, 8(2):182–192, 1980.

15. J. Mao and I. Benbasat. Exploring the use of explanations in knowledge-based systems: a process tracing analysis. Working Paper 96-MIS-002. Faculty of Commerce, University of British Columbia, Canada , 1996.
16. Clifford Nass and Youngme Moon. Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56:81–103, 2000.
17. R. C. Schank. Explanation: A first pass. In J. L. Kolodner and C. K. Riesbeck, editors, *Experience, Memory and Reasoning*, pages 139–165. Erlbaum Associates, Hillsdale, NJ, 1986.
18. S.F. Smith, G. Cortellessa, D.W. Hildum, and C.M. Ohler. Using a scheduling domain ontology to compute user-oriented explanations. In L. Castillo, D. Borrajo, M.A. Salido, and A. Oddi, editors, *Planning, Scheduling, and Constraint Satisfaction: From Theory to Practice*. IOS Press, 2005.
19. S.F. Smith, D.W. Hildum, and D.A. Crimm. Interactive Resource Management in the Comirem Planner. In *IJCAI-03 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco Mexico, August 2003.
20. E.P.K. Tsang. *Foundation of Constraint Satisfaction*. Academic Press, London and San Diego, CA, 1993.
21. L. R. Ye. The Value of Explanation in Expert Systems for Auditing: An experimental Investigation. *Expert Systems with Applications*, 9(4):543–556, 1995.