

Control and Initiative in Collaborative Problem Solving Dialogues

Pamela W. Jordan
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
jordan@isp.pitt.edu

Barbara Di Eugenio
Learning Research & Development Center
University of Pittsburgh
Pittsburgh, PA 15260
dieugeni@cs.pitt.edu

Abstract

In this paper we first question whether control and initiative are as synchronous as current usages appear to make them. We provide some evidence from a corpus of collaborative problem solving dialogues that control should apply to the dialogue level and initiative to the problem solving level. We then ask what role these two phenomena have in mixed initiative dialogues and present an example of how the recognition of lack of initiative at the problem solving level is cued by control at the dialogue level.

Introduction

In this brief position paper, we would like to pose two questions we hope are of interest to the symposium participants:

1. Are *control* and *initiative* the same phenomena? We provide some evidence that they are not.
2. What is the importance of such a distinction for mixed initiative interactive dialogues? We present some preliminary thoughts on this issue.

Control vs. Initiative

Control and *initiative* are two terms that appear to be used almost as synonyms in the literature. In fact they are used this way in the Call for Papers for this symposium. However, it is not clear that they indeed identify the same phenomena.

One definition of control is given in terms of rules for the allocation and transfer of control (Whittaker & Stenton 1988; Walker & Whittaker 1990). First, utterances are classified into 4 types: assertions, commands, questions and prompts. The speaker has control when the last utterance is 1) an assertion unless it is a response to a question, 2) a command or 3) a question unless it is in response to a question or command. With a prompt the hearer has control because the speaker uses a prompt to surrender control.

Initiative is defined as taking the conversational lead where a conversational partner responds to what has been said but also provides new, unsolicited information. Then as initiative passes between the conversational partners so does control (Walker & Whittaker 1990). Since control is defined as it is, it corresponds

to unsolicited information. So initiative and control are viewed as being at least synchronous.¹

These definitions of control and initiative provided useful insights for advisory dialogues and task-oriented dialogues where a master-slave assumption holds (Walker & Whittaker 1990). However, for collaborative problem solving dialogues where the master-slave assumption does not hold it is not clear that the above definitions are sufficient.

We argue that the notion of *control* should apply to the dialogue level, while the use of *initiative* pertains more appropriately to the participants' problem solving goals. The two are related in ways we will discuss below. A first step in this direction is provided by (Guinn 1993), who defines initiative in terms of when to relinquish control. Guinn observes that control should be passed to the participant who is most capable of solving the current goal. However, we observe that in the dialogues we are studying the participants aren't always in a position to decide who is most capable of solving the current goal.² In this case the decision to relinquish control should be preliminarily modified to: relinquish control so that your partner can take the initiative when you can't offer a solution to any of the shared goals, you have no more information to share that is relevant to the shared goals and you don't know that your partner can't offer solutions to any of the shared goals. Thus, the notion of *control* retains its importance, if we can model *control-shifts* as invitations to the partner to take initiative in addressing the shared goals.

Our corpus

Before providing the evidence that supports our claim that initiative and control are different, we will briefly describe our corpus. We have been collecting and studying computer-mediated conversations between pairs of subjects who are collaboratively solving a problem. The participants are of equal status:

¹Marilyn Walker stated in personal communication that control and initiative are the same thing in her view.

²This corpus concerns task-oriented communication in which the participants are all non-experts at the task. See the section on the corpus for more information regarding the communication task and corpus.

they have been briefed together on the domain knowledge that is needed to solve the problem and neither is an expert in solving this type of problem. The problem solving task requires the subjects to buy furniture for the living room and dining room of a house. (This design is based on tasks used in (Walker 1993; Whittaker, Geelhoed, & Robinson 1993)). Each subject is given a separate budget and inventory of furniture that lists the quantities, colors, and prices for each available item. Although each subject's budget and inventory are private information, the subjects can share this information during their conversation. By sharing this information, the subjects can combine their budgets and can select furniture from each other's inventories.

The subjects' main goal is to negotiate the purchases for the two rooms; the items of highest priority are a sofa for the living room and a table and 4 chairs for the dining room. The subjects also have specific secondary goals which further complicate the problem solving task. Subjects are instructed to try to meet as many of these goals as possible. The secondary goals are: 1) Match colors within a room, 2) Buy as much furniture as you can, 3) Spend all your money.

In addition to collaboratively solving the problem, the conversational participants maintain private representations of their incremental agreements in the form of iconic floor plans. We use this additional information as partial evidence of what a speaker meant by an utterance and what was understood by the hearer.

Differences between control and initiative

Let's now go back to the issue of whether control and initiative are necessarily synchronous. Following the rules for determining control provided by (Walker & Whittaker 1990), consider dialogue (1), between Penny and Jill.

Dialogue (1)

P-1: Jill, I have only 1 table - yellow \$400, no yel chairs, but 2 blue chairs-\$75, 2 green-\$100, and 2 red-\$100....My budget is \$400.

J-1: my bankroll is \$550 this time. ... i have 1 table. it's green - \$400. I also have 2 green chairs for \$100 each and 2 yellow chairs for \$75 each.

P-2: hmmm, we can go for the matching green set for \$800, but our total budget is only \$950. I have the yellow \$400 high table and your chairs ...

J-2: i was thinking about your yellow table and 2 blue chairs and my 2 yellow chairs. that would leave us with \$250.

P-1, J-1 and J-2 all demonstrate that the speaker has control of the dialogue. P-2, while still demonstrating that the speaker has control, is also an example of an explicit control-shift according to (Walker & Whittaker 1990), since P-2 can be considered a repetition of what

has been stated so far. However, just observing who has control of the dialogue doesn't seem to be as helpful as trying to determine whether the speaker has problem solving *initiative*, namely, whether s/he is taking it upon him/herself to address the domain goals. In our domain, this amounts to either proposing a (possibly partial) solution or reformulating the goals themselves. In dialogue (1), only J-2 amounts to proposing a solution, even if all the utterances demonstrate that the speaker has control.

A related problem with equating control and initiative is that the definition of control is not adequate in the case of assertions.

In P-1, Penny is addressing the goal of selecting a table and matching chairs — in fact, as this is the first utterance of this dialogue, Penny is putting forward a composite goal (primary goal: get table and 4 chairs; secondary goal: match colors) for consideration. From P-1 and J-1 we see that Jill should realize there are just two tables and that Penny has listed all of the chairs she has available. Jill is in a position to make a proposal or change the goal if it cannot be satisfied. But she does not take the initiative to do so. According to the definition of control given above, Jill has taken control since she ends with an assertion but, while she is supplying new information, she isn't taking the initiative to propose a solution or change the goal.

In P-2 from dialogue (1), Penny should also realize that they cannot solve the goal in its current form but she does not initiate a modification of the goal. She is defined as having control since she ends with an assertion that is not in response to a question but she is also passing control back since she is not providing any new information. (Walker & Whittaker 1990) claims that a repetition of this sort is one way to achieve a control-shift since it signals that the speaker wishes to relinquish control. We claim that P-2 is not a proposal since Penny shows both options as alternatives in her private representation.

Notice that in this dialogue, the explicit control-shift in P2 plays an important role. It indicates that the partner should take initiative in addressing the domain goals because the speaker has reached an impasse in the problem solving task. Note that Jill does take initiative in J-2 by dropping the color match goal.

The computational importance of the control/initiative distinction

Up to now we have pointed out that control and initiative are different and that they are both useful theoretical tools for the analysis of collaborative problem solving dialogues. We turn now to the question of their computational importance.

Consider a mixed initiative interaction between a computer user and a problem solving agent. We claim that in the case of a collaborative problem solving effort in which the information needed to solve the problem

is distributed between the user and the artificial agent, the agent at least needs to be able to recognize when:

1. problem solving is stalled. A lack of initiative when it is expected may signal that the goals can't be achieved and should be modified.
2. work can begin on a new part of the problem. Reaching an agreement is one such possible signal.
3. an earlier commitment may be blocking the attainment of a satisficing solution.

Here we will consider only the problem of recognizing that problem solving is stalled.³

If the system can recognize that initiative was not taken when expected, it can analyze the possible impasses and take appropriate actions to move the problem solving forward. An impasse could indicate a need to revise the goals or the problem solving strategy. A control-shift might be one cue that there is a stall but the system still needs to recognize when to expect initiative and when initiative is actually taken.

The Problem Solving Model

We claimed earlier that initiative should apply to the problem solving level and control to the dialogue level. Taking our claim about initiative seriously, we must explain the problem solving model we propose for the collaborative task associated with the corpus.

The task is more readily modeled as a constraint satisfaction problem than as a planning problem since the temporal ordering of *buy* actions does not affect the solution.⁴ We view the problem space as a set of variables that must have a single value or a set of values of a certain cardinality assigned to them. Since the set of possible values is not known at the outset of problem solving, the model must recognize when to treat the set of values as open, when to treat it as closed and when to reopen it.

The inputs provided by the dialogue to the problem solving model include:

1. the variables to be solved
2. the accumulated values for these variables
3. the current constraints

We characterize the solution space indicated by an utterance as indeterminate if there is no solution or the shared value set for some variable is open, otherwise it is determinate. For example, the solution is indeterminate if the speaker supplies appropriate values for a variable but does not know what the hearer has available for this variable (i.e. the value set is open).

³We address the problems of recognizing agreement in (Jordan *et al.* 1997)

⁴We use the SCREAMER constraint logic programming language (Siskind & McAllester 1993) to model the problem solving. Although SCREAMER does not handle dynamic variables, we temporarily resolve this by setting up the variables and constraints anew with each change.

Recognizing Lack of Initiative

If lack of initiative is a cue that problem solving is stalled then in general initiative is linked to domain reasoning. We will need to describe the domain conditions under which initiative should be expected and the verbal indicators of initiative. These descriptions will be specialized to our problem solving model.

We claim that initiative should be expected when a value set for a variable has been closed. If the hearer interprets the speaker's utterances as an indication that the speaker has provided all the relevant values for the unsolved variable under discussion then the hearer should expect the following verbal displays of initiative:

- The speaker will propose a solution if he has one.
- If the speaker cannot find a solution then he will suggest ways in which to change the problem solving goals so that variables are reopened.

Unfortunately, in collaborative problem solving dialogues it is not always clear who is taking the initiative. A proposal is actually a *joint* speech act; it is equally directive and commissive (Hancher 1979). Equating a directive with initiative is not sufficient since the directive aspect can be hidden by the commissive aspect (Hancher 1979). This makes it difficult to say who is initiating the joint action. Consider dialogue (2), between Dolph and Matt, where the strategy is to find the cheapest sofa.

Dialogue (2)

D-1: so I say let's start with a blue sofa; I have one for \$300.

M-1: ... I do not have a sofa for a better price ...

In D-1, Dolph informs Matt of an option that is relevant to the domain goals. We claim he is not seriously proposing it as the solution since at this point he doesn't know what options Matt has. If Matt has a better option then Dolph expects him to present it, otherwise he expects Matt to accept his option. Since Matt expects Dolph to have presented his cheapest sofa, Matt immediately knows the solution. In this case, since Matt doesn't have a better option to propose, he accepts Dolph's sofa as if it had been seriously proposed.

The point is that by uttering M-1, Matt has reconstructed Dolph's inform in D-1 as a propose (Fox 1987). But in this case Dolph actually expected the reconstruction; it isn't just the result of an ambiguity. We see that if an option is reasonable in the hearer's context then the utterance that introduced the option has the effect of a propose action (Thomason & Moore 1995).

We offer a solution to the problem of recognizing reconstructed proposals in (Jordan *et al.* 1997). In this work we found that we can identify an utterance that reconstructs a prior one using co-reference relations, topic features, and the state of the solution space. Recognizing the utterance that causes the reconstruction avoids

the necessity of somehow previously labelling the prior utterance as being ambiguously an inform or a propose.

At this point, it seems plausible to claim that the utterance initiating the reconstruction process is the one displaying initiative. Now our modified description of the indicators is that an utterance displays initiative if it:

- proposes or reconstructs as a proposal, or
- modifies the problem solving goals in such a way that a variable is reopened.

Looking again at dialogue (1), we expect a display of initiative during J-1 since the value sets for the table and chair variables have been closed during this turn. But since no proposal is made, the turn does not display initiative. Penny should notice this lack of initiative and either reconstruct a proposal or change the constraints during P-2. Instead she does a control-shift by repeating information. Notice that while she does change some of the problem solving constraints (e.g. find a solution that costs < \$800) there are no direct changes to the problem solving goals of finding a matching set of 4 chairs and a table. In J-2, Jill shows that she recognizes that the domain problem solving is stalled and that the problem solving goals should be modified. The modification does not reopen any variables so the proposal is needed to display initiative.

In this section we've proposed a way of recognizing lack of initiative that will apply to collaborative problem solving tasks in which the information needed to solve the problem is distributed among the agents. We plan to test this definition against our corpus once we have enough information coded.⁵ So far we have coded roughly half of the corpus for topic, coreference relations, solution state and informational relations such as initiate⁶ and response-to. However we have not yet begun to code for changes in problem solving goals. Once we have this data we can test the indicators we have proposed for expected initiative and displays of initiative. Since we have informational relations coded, we can also check control and control-shifts in order to confirm that control and problem solving initiative are not synchronous, to verify that control-shifts are cues for lack of initiative and to find other dialogue level cues for lack of initiative.

Acknowledgments

This material is based on work supported by the National Science Foundation under Grant No. IRI-9314961. We wish to acknowledge the contributions of fellow project members Rich Thomason, Johanna Moore, Megan Moser and Jerry Hobbs.

⁵Our coding scheme is based on the one being defined by the Discourse Resource Initiative. See the draft manual at <http://www.cs.rochester.edu:80/research/trains/annotation>.

⁶This usage of initiative is at the dialogue level and not the problem solving level so this coding indicates control only.

References

- Fox, B. 1987. Interactional reconstruction in real-time language processing. *Cognitive Science* 11:365-388.
- Guinn, C. I. 1993. A computational model of dialogue initiative in collaborative discourse. In *AAAI Technical Report FS-93-05*.
- Hancher, M. 1979. The classification of co-operative illocutionary acts. *Language in Society* 8(1):1-14.
- Jordan, P. W.; Di Eugenio, B.; Thomason, R. H.; and Moore, J. D. 1997. Reconstructed intentions in collaborative problem solving dialogues. See <http://www.isp.pitt.edu/intgen/research-papers.html>.
- Siskind, J. M., and McAllester, D. A. 1993. Non-deterministic lisp as a substrate for constraint logic programming. In *Proceedings of AAAI 93*. CA: AAAI Press.
- Thomason, R. H., and Moore, J. D. 1995. Discourse context. In *Proceedings of the AAAI Fall Symposium on Formalizing Context*. CA: AAAI Press.
- Walker, M. A., and Whittaker, S. 1990. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proc. 28th Annual Meeting of the ACL*, 70-79.
- Walker, M. A. 1993. *Informational Redundancy and Resource Bounds in Dialogue*. Ph.D. Dissertation, University of Pennsylvania.
- Whittaker, S., and Stenton, P. 1988. Cues and control in expert client dialogues. In *Proc. 26th Annual Meeting of the ACL, Association of Computational Linguistics*, 123-130.
- Whittaker, S.; Geelhoed, E.; and Robinson, E. 1993. Shared workspaces: How do they work and when are they useful? *IJMMS* 39:813-842.