Generating Information-Sharing Subdialogues in Expert-User Consultation*

Jennifer Chu-Carroll and Sandra Carberry
Department of Computer and Information Sciences
University of Delaware
Newark, DE 19716, USA
E-mail {jchu.carberry} @cis udel edu

Abstract

In expert-consultation dialogues, it is inevitable that an agent will at times have insufficient information to determine whether to accept or reject a proposal by the other agent This results in the need tor the agent to initiate an information-sharing subdialogue to form a set of shared beliefs within which the agents can effectively re-evaluate the proposal This paper presents a computational strategy for initiating such information-sharing subdialogues to resolve the system s uncertainty regarding the acceptance of a user proposal Our model determines when information sharing should be pursued se lects a focus of information-sharing among multiple uncertain beliefs chooses the most effective information-sharing strategy and utilizes the newly obtained information to re-evaluate the user proposal Furthermore our model is capable of handling embedded informauon sharing subdialogues

1 Introduction

We have been studying a particular kind of collaborative dialogue in which two participants (a consultant and an executing agent) collaborate on developing a plan to achieve the executing agent's domain goal In such an environment, the consultant and the executing agent have different knowledge about the domain and about the executing agent's particular circumstances and preferences that may affect the domain plan being constructed Thus it is inevitable that an agenl will not always immediately accept the actions or beliefs proposed by the other agent However, an agenl should recognize the collaborative nature of the interaction and the fact that each agent has pnvate knowledge that is not shared by the other agent Thus, rather than indiscriminately rejecting proposals that she does not have sufficient reasons to accept a collaborative agent should both share her private knowledge with the other agent and solicit relevant information from the other agenl in order for both agents to effectively re-evaluate the proposal and come to the most beneficial decision

This material is based upon work supported by the National Science FoundaUon under Granl No IRI 9122026

Such collaborative in formation-sharing behaviour is Illus trated in the following dialogue segment based on transcripts of naturally occurng dialogues [SRI Transcripts, 1992] In this dialogue a travel agent (T) and a customer (C) are constructing a plan for two other agents to travel from San Francisco to Los Angeles This segment follows a proposal that the travelers be booked on a particular USAir flight

(1) T Can we put them on American?
(2)C Why?
(3) T We re having a lot of problems on the USAir seat maps so I don I know if I can get them together
(4) But American whatever request pretty much we get
(5) C I don t know if they care if they sit together
(6) Let s go ahead and stick with USAir

In this dialogue T proposes putting the travelers on American Airlines instead of USAir in utterance(I) In (2) Cquestions T s motivation for this proposed action — i e the support that T s pnvate knowledge provides for this proposal After T provides her motivation C informs T in (5) that she rejects the motivauon re-evaluates the proposal and in (6) rejects the actions proposed by T

This paper presents a computational model for collaborative information-sharing during proposal evaluation Our model first uses the system s existing beliefs along with evidence pro vided by the user to evaluate user proposals and to determine whether they should be accepted or rejected If the system has insufficient information to make this decision it initiates an in formation-sharing subdialogue to form a set of shared beliefs within which the agents can effectively re-evaluate the proposal and come to agreement This may lead to evaluation of an agent s reasons for a proposal and further information-sharing about an agent s beliefs supporting these reasons thus leading to an embedded information-sharing subdialogue

Our research contributes to response generation in collaborative interaction by I) providing an algorithm for identifying when an information sharing subdialogue should be initiated during proposal evaluation 2) providing a selection algo nthm for determining the beliefs that should be the focus of information-sharing 3) formulating information sharing

strategies and identifying the eritena a for invoking each stral egy and 4) capturing the process in a *Propose Evaluate-Modify* cycle that enables embedded information-sharing subdialogues

2 Modeling Collaborative Activities

In modeling collaborative activities it is essential that the system captures the agents intentions conveyed by their utterances Our model utilizes an enhanced version of the dialogue model described in [Lambert and Carberry 1991] to represent the current status of the interaction The enhanced dialogue model has four levels the *domain* level which consists of the domain plan being constructed for later execution the *problem solving* level which contains the actions being performed to construct the domain plan the *belief* level which consists of the mutual beliefs pursued to further the problem-solving intentions and the *discourse* level which contains the communicative actions initiated to achieve the mutual beliefs [Chu-Carroll and Carberry 1994]

In our earlier work we developed a plan-based model that captures collaborative planning in a *Propose-Evaluate Modify* cycle of actions [Chu-Carroll and Carberry 1994] This model treats a collaborative planning process as a sequence of the following actions agent A s *proposal* of a set of actions and beliefs to be added to the *shared plan* [Grosz and Sidner 1990 Allen 1991] being developed, agent B s *evaluation* of the proposed actions and beliefs and B s proposed *modifications* to the original proposal in cases where the proposal is rejected Notice that B's proposed modifications will again be evaluated by A and if conflicts arise A may propose modifications to B s proposed modifications resulting in a recursive process

However, our previous research assumed that an agent s evaluation of a proposal always results in the proposal being accepted or rejected and did not take into account cases in which the agent initially has insufficient information to determine whether or not to accept the proposal, as shown in the example in utterances (l)-(6) This paper extends our earlier work by providing a computational strategy for collaborative information sharing during proposal evaluation We focus on situations in which the system s lack of knowledge occurs during the evaluation of proposals at the belief level of the dialogue model [1]

3 Information-Sharing During Collaboration

Since a collaborative agent initiates information-sharing subdialogues to help determine whether to accept or reject a proposed belief the information sharing process is captured as part of the *evaluation* process in the *Propose-Evaluate Modify* cycle for collaborative activities Thus the evaluation of a proposed belief involves the agent 1) determining the acceptance of the proposed belief based on the information currently available to her, and 2) in cases where she cannot decide whether to accept or reject the belief, initiating an information-sharing subdialogue so that the agents can exchange information and re-evaluate the proposed belief The following sections describe these two processes

[1] We are concerned with situations in which the system recognizes [he user s proposal but cannot decide whether to accept or reject it not those where the system initiates a clarification subdialogue to disambiguate the user s proposal [van Beck *el al* 1993 Logan *et a!* 1994 Heeman and Hirst 1992 RaskulU and Zukerman 1993]

Evaluate-Belief(_bel}

1 evidence set - *be*l (appropriately endorsed as conveyed by the user) and the system s beliefs that support or attack _bel
2 If _bel has no children return Evaluate(_bel evidence set)
3 Evaluate each of _bel s children _beh $\_bel_n$
  3 1 belief_resuln-Evaluate-BeUer(_bel,)
  3 2 rel_resulL Evaluate-BeUeT(supporls(_bel, _bel))
  33 If belief-result = rejector rel_rcsuli= reject ignore _bel, and supports(-bel, _bel)
  3 4 Else if belief .result = rel_result = accept add{-bel, supports(_bel, _bel)} to the evidence sei
  3 5 Else if belief_resuli = unsure or rel_resuli = unsure add {_bel, supports(_bel, _bel)} to the potenUal evidence set
4 Evaluate _bel
  4 1 uppcrbound *— Evaluate(_bel evidence set + potential evidence sel)
  4 2 lowerbound *4-* Evaluate(_bel evidence set)
  4 3 If upperbound = lowerbound = accept, accept _bel
  4 4 Else if upperbound = lowerbound = reject rejeel _bel
  4 5 Else unsure about -bel annotate _hel with upperbound lowerbound evidence sel and potential evidence set

Figure 1 Algorithm for Evaluating a Belief

3 1 Evaluating Proposed Beliefs

Our system maintains a set of beliefs about the domain and about the users beliefs Associated with each belief is a *strength* that represents the agent s confidence in holding the belief We model the strength of a belief using *endorsements* [Cohen, 1985] followingIGalliers 1992 Loganer*al* 1994] based on the semantic form of the utterance used to convey a belief the level of expertise of the agent conveying the belief stereotypical beliefs, etc

The belief level of our dialogue model consists of one or more belief trees where the belief represented by a child node is intended to support that represented by its parent When an agent proposes a new belief and gives (optional) supporting evidence for it, this set of proposed beliefs is represented as a belief tree The system must then evaluate the proposed beliefs in order to determine whether to accept the proposal reject it, or pursue information-sharing to allow the agents to re-evaluate it The algorithm *for* evaluating proposed beliefs is shown in Figure 1 and is applied to the root node of each proposed belief tree (the top-level proposed beliefs) Since the acceptance of a child belief may affect the acceptance of its parent before determining the acceptance of a belief or evidential relationship its children in *the* proposed belief tree must be evaluated (step 3) Thus for each child belief of _bel, the system evaluates both the belief (step 3 1) and the evidential relationship between the belief and _bel (step 3 2) A piece of evidence is marked as 1) accepted if both the child belief and the evidential relationship are accepted 2) rejected if either the child belief or the evidential relationship

is rejected, and 3) uncertain otherwise

To determine the status (accepted, rejected, or uncertain) of a belief _bel the algorithm constructs an *evidence set* that contains the user s proposal of _bel, endorsed according to the user s level of expertise in that subarea as well as the user s strength in the belief as conveyed by the semantic form of the utterance (step 1), the system s own beliefs pertaining to _bel (step 1), and evidence proposed by the user that is accepted by the system (step 3 4) It also constructs a *potential evidence set* consisting of evidence proposed by the user whose acceptance is undetermined (step 3 5) The algorithm must then determine whether the potential evidence could have an impact on the system s decision-making It first evaluates _bel by invoking the Evaluate function[2] to compute an upperbound and a lowerbound for the system s acceptance of _bel The upperbound is computed by invoking the Evaluate function with evidence from both the evidence set and the potential evidence set I e treating all uncertain evidence as accepted and the lowerbound is computed by invoking Evaluate with only the evidence set i e , treating all uncertain evidence as rejected (steps 4 1 and 4 2) If _bel is either accepted or rejected in both cases, indicating that the uncertainty of the evidence if any, does not affect the acceptance of _bel the system accepts or rejects _bel (steps 4 3 and 4 4) Otherwise the system has insufficient information to determine the acceptance of _bel and it is marked as uncertain (step 4 *5)* If the top-level proposed belief is marked as uncertain, an information-sharing subdialogue will be initiated as described in the next section

3 2 Initiating Information-Sharing Subdialogues

A collaborative agent when facing a situation in which she is uncertain about whether to accept a proposal should attempt to share information with the other agent so that each agent can knowledgably re-evaluate the proposal and the agents can come to agreement — to do otherwise is to fail in her responsibilities as a collaborative agent Furthermore a collaborative agent should engage in effective and efficient dialogues thus she should pursue the information sharing subdialogue that she believes will most likely result in the agents coming to an intelligent decision about the proposal The process for initiating information-sharing subdialogues involves two steps selecting a focus of information-sharing from the proposed beliefs marked as uncertain during the initial evaluation process, and selecting an effective information-sharing strategy

Selecting the Focus of Information Sharing

The possible combinations of the upperbound and lowerbound values produced by the Evaluate-Belief algorithm (Figure 1) are shown in Figure 23 Cases 1 and 6 correspond to steps

	Upper	Lower	Action
1	accept	accept	accept _bel
2	accept	unsure	attempt to accept uncertain children in order to accept _bel
3	accept	reject	actions in cases 2 and 5
4	unsure	unsure	resolve uncertainty regarding _bel itself
5	unsure	reject	attempt to reject uncertain children in order to reject _bel
6	reject	reject	reject _bel

Figure 2 Combinations of Upperbounds and Lowerbounds

4 3 and 4 4 in Figure I respectively in which the decision to accept or reject is the same whether or not beliefs in the potential evidence set are accepted In these cases the uncertainty in the child beliefs need not be resolved since their acceptance will not impact acceptance of the parent belief that they are intended to support and thus will not affect acceptance of the top-level proposed belief that is important to the plan being constructed[4] In case 4 the system will remain unsure whether to accept or reject _bel regardless of whether the uncertain child beliefs if any are accepted or rejected i e resolving the uncertainty in the child beliefs will not help resolve the uncertainty in bel Thus me system should focus on sharing information to resolve the uncertainty about _bel itself instead of its children In cases 2 and 3 acceptance of the child beliefs has the potential to influence acceptance of _bel and in cases 3 and 5 rejection of the child beliefs can lead to rejection of _bel Thus in all three cases the system should initiate information-sharing that will allow the agents to come to agreement about the currently uncertain child beliefs

However there may be more than one uncertain child belief Thus when the system initiates information-sharing it must first select a belief on which to focus during the information sharing process Our algorithm for selecting the focus of information-sharing is shown in Figure 3 Select Focus-Info-Sharing is initially invoked with _bel instantiated as the top-level proposed belief Step 3 of the algorithm corresponds to case 4 in Figure 2 where the uncertainty in the child beliefs is irrelevant to me acceptance of _bel thus the focus of information-sharing is _bel itself Steps 4 and 5 of the algorithm correspond to cases 2 3 and 5 in Figure 2 where the system attempts to share information to resolve the uncertai nty in the child beliefs and perhaps thereby accept or reject -bel

Step 4 of the algorthm is concerned with cases where the potential acceptance of uncertain child beliefs may lead to the acceptance of _bel (caies 2 and 3 in Figure 2) In selecting the focus of Inlormation-shanng in such cases two factors should come into play 1) how strongly the acceptance of each piece of evidence affects the acceptance of _bel — the stronger the impact that the potential evidence can have on the acceptance of _bel, the more useful it is to expend effort on resolving the

[3]Evaluate utilizes a simplified version of Galliers belief revision mechanism [Galliers 1992 *Loganeral* 1994] which given a set of evidence compares the endorsements of the beliefs that support and attack bel and determines whether or not bel should be accepted

[3]Our model assumes that a child belief is always intended to provide support for its parent belief (a piece of counter evidence is represented as a child belief supporting the negation of the parent belief thus only six out of the rune theoretically possible combinations may occur

[4] Young et al [Young *et al [*1994] argued that if a belief is accepted even though a child belief that is intended to support it is rejected the rejection of the child belief need not be addressed since it is no longer relevant Our strategy extends this concept to uncertain information

Select Focus-Info-Sharing(_bel)

1 If _bel is accepted or rejected focus ← {} return focus
2 If _bel has no children focus ← _bel return focus
3 If upper = lower = uncertain focus ← _bel return focus
4 If upper = accept
  4 1 Assign each piece of uncertain evidence (_bel, and supports(_bel, _bel)) to a set and order the sets according to how close the beliefs in each set were to being accepted Call them $\_set_1$ ... $\_set_m$
  4 2 For each set in ranked order do until new_result, = accept new_result, ← **Evaluate**(_bel evidence set + set_i)
  4 3 If new_result, ≠ accept increase set size by one re rank sets and goto 4 2
  4 4 Else focus ← $\bigcup_{\_el_j \in \_set_1}$ **Select-Focus-Info-Sharing**($\_el_j$) return focus
5 If lower = reject
  5 1 Assign each piece of uncertain evidence (_bel, and supports(_bel, _bel)) to a set and order the sets according to how close the beliefs in each set were to being rejected Call them $\_set_1$ ... $\_set_m$
  5 2 For each set in ranked order do until new_result, = reject new_result, ← **Evaluate**(_bel evidence set + potential evidence set set_i)
  5 3 If new_result, ≠ reject increase set size by one re rank sets and goto 5 2
  5 4 Else focus ← $\bigcup_{\_el_j \in \_set}$ **Select Focus-Info-Sharing**($\_el_j$) return focus

Figure 3 Selecting the Focus of Information-Sharing

uncertainty about the proposed evidence and 2) how close each piece of evidence was to being accepted during the initial evaluation process — the closer a piece of evidence is to being accepted, the easier n is for the system to gather sufficient information to accept the evidence Our algorithm first constructs a singleton set for each piece of uncertain evidence for _bel where a piece of evidence includes a pair of beliefs a child belief _bel, and the evidential relationship between _bel, and _bel *supports(_bel_bel)* The sets are ordered according to how close the beliefs in a set were to being accepted in the initial evaluation process (step 4 1) The first set (_seti) is then added to the evidence set and _bel is re-evaluated with respect to the augmented evidence set (step 4 2) thus considering the potential effect of the acceptance of beliefs in $\_set_1$ on the acceptance of _bel If the result of the evaluation is to accept _bel indicating mat resolving the uncertainty of the beliefs in _seti is sufficient to resolve the uncertainty of _bel then Select-Focus-Info-Sharing is recursively applied to each belief in _seti in order to determine the focus for resolving the uncertainty of beliefs in $\_set_1$ (step 4 4) On the other hand if the evaluation indicates that accepting $\_set_1$ does not result in the acceptance of _bel the next set ($\_set_2$) is tried This continues until either the uncertain evidence in a set is predicted to resolve the uncertainty of _bel or all of the uncertain evidence is tried and none suffices for acceptance of _bel In the latter case, the set size is increased by one, sets of uSe requisite size are constructed by combining individual pieces of evidence the new sets are ordered and the same process is repeated (step 4 3) Thus our algorithm guarantees that the fewest possible beliefs are selected as the focus of information-sharing and that these beliefs require the least effort to achieve among those that are strong enough to affect the acceptance of _bel[5]

Step 5 of the algorithm corresponds to cases 3 and 5 in Figure 2 The procedure for step 5 is similar to that for step 4 except that in predicting the effect of resolving a piece of uncertain evidence _bel is evaluated under the assumption that the set of uncertain evidence under consideration is rejected while the other uncertain beliefs are accepted (step 5 2)

Selecting an Information-Sharing Strategy

We have identified four strategies which a collaborative agenl may adopt in initiating an information-shanng subdialogue to allow the agents to share information and re-evaluate a belief or evidential relationship _bel

1 Agent A may present a piece of evidence against .bel and (implicitly) invite agent B to attack it Such a strategy focuses B s attention on the counterevidence and suggests that it is what keeps A from accepting _bel Thus in collaborative activities this strategy should only be employed it As counterevidence is *critical* i e if proving that the counterevidence is invalid will cause A to accept _bel This strategy also allows the possibility of B accepting the counterevidence and perhaps both agents subsequently adopting _bel instead of _bel
2 Agent A may query B about his reasons for believing in _bel This strategy is appropriate when A does not know B s support for _bel and also does not have evidence against _bel herself It would result either in A gathenng evidence that contributes toward her adopting _bel or in A discovering B s invalid justification for holding _bel and attempting to convince B of ->_bel
3 Agent A may query B for his evidence for _bel and also present her reasons for believing in ->_bel This strategy is adopted when A does not know B s reasons for believing _bel but does have non-critical evidence against accepting _bel In this case B may provide his support for _bel, attack As evidence against _bel or accept As counterevidence and perhaps subsequently adopt ->.bel
4 Agent A may indicate her uncertainty about _bel and present her reasons against _bel This strategy is adopted when A is least certain about how to go about sharing information to resolve the uncertainty — when A already knows B s reasons for believing _bel and only has non-cntical evidence against accepting _bel In a collaborative environment, A s indication of the uncertainty in her decision should lead B to provide information that he believes will help A re-evaluate the proposal

The process for initiating informal]on-sharing subdialogues is performed by invoking the *Share-Info Reevaluate Belief*

[5]In cases where the focus set contains multiple beliefs additional processing is needed to determine the most coherent order in which to address the beliefs

Action	Reevaluate After Invite Attack(_s1 _s2 _bel1 _top-belief _belief tree)
Type	Decomposition
Appl Cond	uncertain(_s1 _bel1) believe(_s1 _bel2) believe(_s1 supports(bel2 ¬_bel1)) results in(believe(_s1 ¬_bel2) believe(_s1 _bel1))
Constraint	member of(_bel1 _belief tree) root-of(_belief tree _top-belief)
Preconds	MB(_s1 _s2,_bel2) ∧ MB(_s1 _s2 supports(_bel2 ¬_bel1)) ∨ MB(_s1 _s2 ¬_bel2) ∨ MB(_s1 _s2 ¬supports(_bel2 ¬_bel1))
Body	Evaluate-Proposed Beliefs(_s1 _s2 _top-belief)
Goals	believe(_s1 _top-belief) ∨ believe(_s1 ¬_top-belief)
Action	Reevaluate After Ask Why(_s1 _s2 _bel1 _top-belief _belief tree)
Type	Decomposition
Appl Cond	¬knowref(_s1 _bel2 believe(_s2 supports(_bel2 _bel1)) ¬knowref(_s1 _bel3 supports(_bel3 ¬_bel1))
Constraint	member of(_bel1 _belief tree) root of(_belief tree _top-belief)
Preconds	knowref(_s1 _bel2 believe(_s2 _supports(_bel2 _bel1)))
Body	Evaluate-Proposed Beliefs(_s1 _s2 _top-belief)
Goals	believe(_s1 _top-belief) ∨ believe(_s1 ¬_top-belief)

Figure 4 Specializations of *Share Info Reevaluate Belief*

problem solving action on the focus identified by Select-Focus-Info-Sharing (Figure 3) It initiates an information sharing subdialogue using the most appropriate of the four information sharing strategies and re-evaluates the top-level belief taking into account the newly obtained information The recipe8 for Share Info Reevaluate Belief specifies that in order for the action to be invoked it must be the case that the system believes in neither a top-level proposed belief (_bel) nor its negation — that is the system cannot determine whether to accept or reject _bel The body of Share Info Reevaluate Belief consists of alternative subactions which correspond to the aforementioned strategies that a collaborative agent can use to pursue information sharing The recipes for two of these subactions are shown in Figure 4

The first specialization Reevaluate After Invite Attack corresponds to the first information sharing strategy in which the system (_sl) has a piece of critical evidence (_bel2) against believing _bel 1 a belief proposed by the user (_s2) and about which the system is uncertain This criterion is captured in the applicability conditions7 of the acuon the conditions that the system is uncertain about the acceptance of _bel 1 that the system believes in _bel2 which provides support for --bell and that the systems disbelief in _bel2 will result in its adoption of _bell The preconditions of Reevaluate After Invite-Attack however show that the action cannot be performed until one of the following conditions is true 1) the system and the user mutually believe (MB) in _bel2 and mutually believe that _bel2 supports -_bell 2) the system and the user mutually believe in _bel2 or 3) the system and the user mutually believe that _bel2 does not support —bel In order to satisfy the preconditions the system will adopt the Express Doubt discourse action [Lambert and Carberry 1992] in which the system expresses doubt at _bell by contending _bel and the evidential relationship between _bel2 and --_bell as an attempt to achieve MB(S U _bel12) and MB(S U supports(J?el2 _bell)) 8 Thus the system will initiate an information-sharing subdialogue by expressing its evidence against the proposed belief and inviting the user to comment on it If the outcome of the in form an on-sharing subdialogue satisfies one of the preconditions of Reevaluate After Invite Attack the system can perform the body of the action and re-evaluate -top-belief (the root node of the proposed belief tree of which _bel I is a pan) taking into account the newly obtained information Notice that the user s response to the Express Doubt discourse action is again considered a proposal of mutual beliefs and will be evaluated by the system The system may again have insufficient information to determine whether to accept or reject the new proposal which was intended to resolve the uncertainty of the previous proposal It will then initiate another information sharing subdialogue to resolve the new uncertainty resulting in embedded information sharing subdialogues

The second specialization of Share Info Reevaluate Belief is Reevaluate After Ask why which corresponds to the second information-sharing strategy in which the system attempts to find out the user s justification for believing -bell The action is applicable (Figure 4) it the system (_sl) does not know the users (_s2 s) justification for holding -bell and also does not have any evidence against _bel 1 itself We argue that a collaborative agent should not accept a proposed belief merely because of the lack of evidence to the contrary Instead, she should only accept a belief if the evidence supporting the belief is strong enough to warrant acceptance For instance suppose a student informs his advisor that the AI course scheduled for next semester has been canceled withouot giving any justification for it (such as attributing the source of the knowledge) although the advisor may not have evidence against believing in the cancellation she does not immediately accept the proposed belief because given the student s presumed low expertise in the domain the endorse ment attached to the proposed belief is not reliable enough to warrant acceptance The precondition of Reevaluate After Ask Why indicates that the action can be performed only if the system knows the users evidence for holding -bell In order to satisfy this precondition the system will adopt discourse actions to query the user for such information thus

6A recipe [Pollack 1986] is a template for performing actions It contains the applicability conditions for performing an action the subactions comprising the body of an action etc

7 Applicability conditions are conditions that must already be satisfied in order for an action to be reasonable to pursue whereas an agent can try to achieve unsatisfied preconditions

8These two mutual beliefs are selected as preconditions to be satisfied because the system itself holds these beliefs The alternative preconditions are present in order to capture situations in which the user in response to the Express Doubt acuon convinces the system that _bel2 is false or that _bel2 does not serve as justification for —_bel I The use of these alternative preconditions will be demonstrated in Section 4 2

initiating an information-sharing subdialogue

4 Example

Suppose that the system an expert in the university course advisement domain has proposed the options of taking Logic or Algorithms to satisfy the user s core course requirement Consider the following continuation which illustrates many of the features of our strategy for information-sharing during proposal evaluation

(7) U Logic is a better choice than Algonthms
(8) Dr Smith is teaching Logic
(9) S Isn t Dr Smith going on sabbatical next year[7]
(10) U I thought he postponed his sabbatical until 1996
(11) S Why do you think Dr Smith postponed his sabbatical until 1996[7]
(12) Isn t he spending next \ear at I B M '

In utterance (9), S initiates an in formation-sharing subdialogue to determine whether to accept or reject the belief that Dr Smith is teaching Logic proposed by U in (8) by expressing a strong but not warranted belief[0] thai Dr Smith is going on sabbatical next year In (10) U initiates an information-shanng subdialogue to determine whether to accept S s claim that Dr Smith is going on sabbatical next year by expressing his weak belief that Dr Smith has postponed his sabbatical Finally in (11) and (12) S initiates an information-shanng subdialogue to determine whether to accept U s claim that Dr Smith has postponed his sabbatical by explicitly querying U s reasons for holding this belief and expressing her belief that Dr Smith is spending next year at I B M The following sections describe how our model will produce these information-sharing subdialogues

4 1 Evaluating Utterances (7) and (8)

Utterances (7) and (8) propose two mutual beliefs Better ThanfLogic Algonthms) and Teaches(Smirh,Logic) as well as an evidential relationship that the latter provides support for the former When presented these proposed beliefs, the system will first determine whedier to accept or reject the proposal by invoking Evaluate-Belief (Figure I)on the top-level proposed belief Better Than(Logic Algonthms) The system will evaluate the proposed evidence as part of evaluating the belief (step 3 in Figure I) thus recursively invoking Evaluate-Belief on TeachesfSmith Logic) (step 3 1) and the proposed evidential relationship (step 3 2) Since Teaches(Smith Logic) has no children in the proposed belief tree u will be evaluated by a simplified version of Galliers' belief revision mechanism [Galliers, 1992] (step 2) Suppose that the system has the following evidence pertaining to Teaches(Smith Logic) 1) a strong belief that Dr Smith usually teaches Logic, 2) a strong belief that Dr Smith is going on sabbatical next year and a warranted belief that going on sabbatical implies that a faculty member is not teaching courses and 3) the user's belief diat Dr Smith is teaching Logic The strengms of evidence for and against Teachesf Smith, Logic) will be combined and compared In this case, the strengths of the two sets of evidence are relatively comparable thus the system will not be able to decide whether to accept or reject Teaches(Smith Logic) based on the available information The system will then evaluate the proposed evidential relationship (step 3 2) Since the system believes that 1) the user believes that Dr Smith is a good teacher and 2) students generally prefer courses taught by good teachers Dr Smith teaching Logic provides support for die user preferring Logic to Algonthms thus the proposed evidential relationship will be accepted Since the proposed evidential relationship is accepted while the child belief is uncertain this piece of evidence will be added to the potential evidence set (step 3 5)

The system will men evaluate die top-level proposed belief taking into account the result of evaluating its only piece of evidence provided by the user The system's evidence set for Better Than(Logic Algonthms) consists of a warranted belief mat Algorithms is a pre-requisite for more CS courses than Logic is which provides some support for Algonthms being a better choice than Logic as well as the user s statement that Logic is a better choice dian Algonthms The potential evidence set consists of a pair of beliefs the uncertain belief that Dr Smith is teaching Logic and the accepted evidential relationship that Dr Smith leaching Logic provides support for Logic being a better choice than Algonthms When both the evidence set and the potential evidence set are included in the evaluation the system will compute die upperbound of the acceptance of Better-Than(Logic Algonthms) to be accept When considering only evidence from the evidence set, however, the system will be uncertain about the acceptance of the proposed belief The result of this evaluation corresponds to case 2 in Figure 2 and results in the need for the system to initiate an information-shanng subdialogue to resolve the uncertainty

Since die system cannot decide whedier to accept either of the proposed mutual beliefs, it will select a focus of information-sharing by invoking Select-Focus-Info-Sharing (Figure 3) on Better Than(LogicAlgonthms) Since acceptance of die only piece of evidence provided by the user results in acceptance of the top-level proposed belief (step 4 2) the algonthm will be applied recursively to the child belief Since it in turn has no children, TeachesfSnuth Logic) itself will be selected as the focus of information-sharing

The system will now invoke Share Info Reevaluate Belief on the identified focus Since the system s belief that Dr Smith is going on sabbatical and its belief in the evidential relationship that being on sabbatical implies that Dr Smith is not teaching Logic constitute the only obstacle against its accepting Teaches(Smith Logic) they are considered a piece of cntical evidence Thus Reevaluate After Invite-Attack will be selected as the specialization of Share Info Reevaluate Belief Figure 5 shows the dialogue model that will be constructed for this process In order to satisfy the preconditions of Reevaluate-After-Invite-Attack (Figure 4) the

[9]The strength of a belief falls into one of three categories war ranted strong or weak hased on the endorsements of the belief

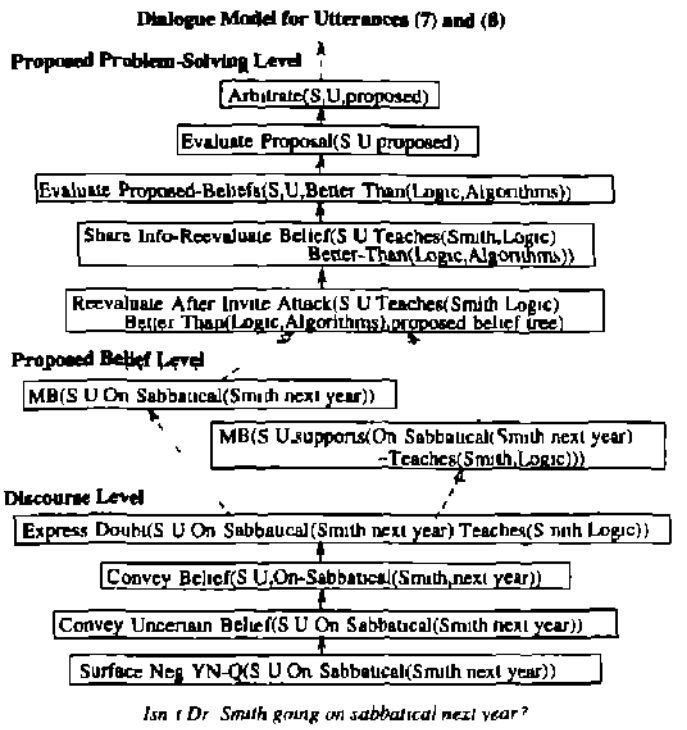

Figure 5 Dialogue Model for Generating Utterance (9)

system will post *MB(S U On Sabbatical(Smith next year))* and *MB(S U supports(On Sabbatical(Smith next year) ¬Teaches(Smith Logic)))* as mutual beliefs to be achieved The system will adopt the *Express Doubt* discourse action to convey its strong belief in *On Sabbatical(Smith next year)* and its warranted belief in *supports(On Sabbatical(Smith next year) ¬Teaches(Smith Logic))* However the latter will not be explicitly stated because the system believes that the user will derive the evidential relationship based on stereotypical knowledge and the structure of the discourse [10] Thus the system would generate the following utterance

(9) S Isn t Dr Smith going on sabbatical next year?

4.2 Possible Follow-Ups to Utterance (9)

Consider the following alternative responses to (9)

(10a) U Oh you re right I forgot about that

(10b) U He was planning on it but he told me that he had decided to postpone it until 1996

(10c) U Yes but he got a good offer from the department to teach Logic while he s on sabbatical

(10d) U I thought he postponed his sabbatical until 1996

Utterance (10a) will be captured as the user accepting the proposal in Figure 5 and acknowledging it Utterances (10b) and (10c) are both the result of the user rejecting the system's proposed mutual beliefs and attempting to modify them

[10] Walker [Walker 1994] has shown the importance of IRU S (Informationally Redundant Utterances) in efficient discourse We leave including appropriate IRU s for future work

[Chu-Carroll and Carberry 1995] In (10b) the user rejects *MB(S U On Sabbatical (Smith next year))* while in (10c) the user rejects the evidential relationship *MB(S U supports(On Sabbatical(Smith, next year) -Teaches(Smith Logic)))* In the case where the user responds with (10b) if the system accepts the proposed belief that Dr Smith postponed his sabbatical until 1996, the system and the user achieve the mutual belief *MB(S U -non-Sabbattcal(Smith next year))* Thus, the precondition of the *Reevaluate After Invite Attack* action in Figure 5 is satisfied and the system would re-evaluate *Better-Than(Logic Algorithms)* taking into account the newly obtained information Notice that in this case the mutual belief achieved to satisfy the preconditions of *Reevaluate After-Invite Attack* is different from the ones the system attempted to achieve — utterance (9) was generated as an attempt to achieve *MB(S, U On SabbaticaH Smith next \ear))* but the result is that both the system and the user accept *MB(S U _On Sabbattcal(Smith next year))* Ithis case although the goal of the *Express Doubt* discourse action is not satisfied the agents mutual belief achieves the higher-level goal that the *Express Doubt* action is intended to achieve namely a pie condition of *Reevaluate After-Invite Attack thus the Express Doubt* action is abandoned This example shows how the precondition of *Reevaluate After Invite Attack* captures situations in which the user presents counterevidence to the system s critical evidence and changes the system s beliefs

4 3 Evaluating Utterance (10d)

Utterance (10d) will be interpreted as a case in which the user is uncertain about whether to accept or reject the system s proposal in (9) and attempts to share information with the system to re-evaluate the proposal It proposes a mutual belief *Postpone Sabbatical (Smith 1996)* which will be evaluated by Evaluate-Belief Suppose the system believes that Dr Smith is spending next year at I B M which is evidence against Dr Smith postponing his sabbatical Then the system cannot determine the acceptance of the proposed belief, resulting in the need to initiate an information-sharing subdialogue The focus of information-sharing is *Postpone Sabbancal Smith 1996)* since it is the only uncertain belief The system will then select an appropriate information-sharing strategy Since the system does not know the user's reasons for believing *Postpone-Sabbatical(Smith 1996)* but does have a piece of non-critical evidence against the proposed belief the third information sharing strategy will be selected Thus the system would query the user for support for the proposed belief and also provide its evidence against the belief, leading to the generation of the following utterances

(11) S Why do you think Dr Smith postponed his sab batical until 1996?

(12) Isn t he spending next year at IBM?

5 Related Work

Grosz, Sidner and Lochbaum [Grosz and Sidner 1990 Lochbaum, 1991] developed a SharedPlan approach to modelling collaborative discourse and Sidner [Sidner, 1994]

formulated an artificial language for modeling such discourse Sidner viewed a collaborative planning process as proposal/acceptance and proposal/rejection sequences Her artificial language treats an utterance such as *Why do X?* as a proposal for the hearer to provide support for his proposal to do X However, Sidner s work is descriptive and does not provide a mechanism for determining when and how such a proposal should be made nor how responses should be formulated in information-sharing subdialogues

Several researchers have studied the role of clarification dialogues in disambiguating user plans [van Beek *et al* 1993 Raskutti and Zukerman, 1993] and in understanding referring expressions [Heeman and Hirst 1992] Loganelal ILogan et *al,* 1994] developed an automated librarian that could revise its beliefs and intentions and could generate responses as an attempt to revise the user s beliefs and intentions Although their system had rules for asking the user whether he holds a particular belief and for telling the system s attitude toward a belief the emphasis of their work was on conflict resolution and plan disambiguation Thus they did not investigate a comprehensive strategy for information-sharing during proposal evaluation For example, they did not identify situations in which informarjon-shanng is necessary, did not address how to select a focus of information-sharing when there are multiple uncertain beliefs did not consider requesting the user s justifications for a belief etc In addition they do not provide an overall dialogue planner that takes into account discourse structure and appropriately captures embedded subdialogues

6 Conclusion

This paper has presented a computational strategy for collaborative information-shanng in situations where the system s current knowledge does not allow it to make a decision about whether to accept or reject a user proposal Our model in eludes algorithms for determining when in formation-sharing subdialogues should be initiated and for selecting a focus of in formation-sharing The latter algorithm takes into account both the effect of the acceptance of a piece of evidence on the acceptance of the top-level belief and the difficulty in resolving the uncertainty about acceptance of a piece of evidence Furthermore we have identified four alternative information-sharing strategies and the criteria under which each should be invoked thus allowing the agents to share the most pertinent information in order to re-evaluate a proposal In addition by capturing in formation-sharing as part of the evaluation process in a Propose *Evaluate Modify* cycle of actions our model can handle embedded information sharing subdialogues

Acknowledgments

The research has benefitted from discussions with Stephanie Elzer, Kathy McCoy and Candy Sidner

References

[Allen 1991] James Allen Discourse structure in the TRAINS project- In *Darpa Speech and Natural Language Workshop* 1991

[Chu-Carroll and Carberry, 1994] Jennifer Chu-Carroll and Sandra Carberry A plan-based model for response generation in collaborative task-onented dialogues In *Proceed ing of the 12th AAAI* pages 799-805 1994

[Chu-Carroll and Carberry, 1995] Jennifer Chu-Carroll and Sandra Carberry Response generation in collaborative ne gotiation In *Proc of the 33rd ACL,* 1995 To appear

[Cohen 1985] Paul R Cohen *Heuristic Reasoning about Uncertainty An Artificial Intelligence Approach* Pitman Publishing Company 1985

[Galliers 1992] Julia R Galliers Autonomous belief revision and communication In Gardenfors editor *Belief Revision* Cambridge University Press, 1992

[Grosz and Sidner 1990] Barbara J Grosz and Candace L Sidner Plans for discourse In Cohen, Morgan and Pollack editors *Intentions in Communication* chapter 20 pages 417-44 MIT Press 1990

[Heeman and Hirst 1992] Peter A Heeman and Graeme Hirst Collaborating on referring expressions Technical Report 435 University of Rochester 1992

[Lambert and Carberry 1991] Lynn Lambert and Sandra Carberry A tripartite plan-based model of dialogue In *Proceedings of the 29th ACL,* pages 47-54 1991

[Lambert and Carberry 1992] Lynn Lambert and Sandra Carberry Modeling negotiation dialogues In *Proceed infs of the 30th ACL* pages 193-200 1992

[Lochbaum, 1991] Karen E Lochbaum An algorithm for plan recognition in collaborative discourse In *Proceedings of the 29th ACL,* pages 33-38,1991

[Logan *et al* 1994] Brian Logan Steven Reece Alison Cawsey Julia Galliers and Karen Sparck Jones Belief re vision and dialogue management in information retrieval TR339, Univ of Cambridge Computer Laboratory 1994

[Pollack 1986] Martha E Pollack A model of plan inference that distinguishes between the beliefs of actors and observers In *Proc ofthe24thACL* pages207-214 1986

[Raskutti and Zukerman 1993] Bhavani Raskutti and Ingnd Zukerman Eliciting additional information during coop erativc consultations In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society* 1993

[Sidner 1994] Candace L Sidner An artificial discourse language for collaborative negotiation In *Proceedings of the I2thAAAI,* pages 814-819 1994

[SRI Transcripts 1992] SRI Transcripts Transcripts derived from audiotape conversations made at SRI International Menlo Park.CA 1992

[van Beek *et al*, 1993] Peter van Beck, Robin Cohen, and Ken Schmidt From plan critiquing to clarification dialogue for cooperative response generation *Computational Intelligence,* 9(2) 132-154,1993

[Walker 1994] Manlyn A Walker Discourse and deliberation Testing a collaborative strategy In *Proceedings of COLING 94,* 1994

[Young *etal* 1994] R Michael Young, Johanna D Moore andMarthaE Pollack Towards a principled representation of discourse plans In *Proceedings of the Annual Meeting of the Cognitive Science Society,* pages 946-951, 1994