Conversational Cooperation: The Leading Role of Intentions

G. Boella, R. Damiano, L. Lesmo and L. Ardissono
Dipartimento di Informatica - Università di Torino

Abstract

In this paper, we consider the role played by the notion of agent intentions in dialog and, particularly, in conversational cooperation.
We present a model, where intentions arise from goals on the basis of two factors: the advantage of achieving a goal, and the cost of the actions to achieve it.
In the model, the adoption of the intention of cooperating at the conversational level is motivated by the fact that the refusal to do low-cost actions, is usually interpreted as an offense to the requester of the action, and that preventing a partner from becoming offended is one of the goals the speaker considers when taking a decision about what to do.
Finally, we show that the decision to cooperate towards the success of the communication accounts for grounding, requests of repair, repairs to misunderstandings, and other related phenomena.

1 Introduction

We assume that dialog is a goal-directed behavior, where the participants realize the existence of some goals of another agent and decide to cooperate (to a certain extent) with him. However, two different types of goals can be identified: domain goals (e.g. knowing the time, or the place where you can buy some gasoline, or getting your friend's car for this night) and dialog goals (e.g. receiving an answer, or a confirmation that a statement of yours has been understood properly; these goals will later be classified as Linguistic or Conversational). In each turn, a speaker expresses one or more goals of both types and the hearer either accepts to cooperate (i.e. adopts that goal) or does not accept.

As noticed in (Airenti et al., 1993), a participant, even when he does not adopt the domain goals of his partner, typically continues to cooperate at a conversational level, by informing him that he has not the required information or by explicitly refusing to perform the requested actions. So, it seems that two different aspects enter into play: one concerns domain behavior, and leaves the hearer free to adopt the partner's goals, one concerns conversational behavior, and often forces the hearer to cooperate independently from his own goals.

(Traum and Allen, 1994) have challenged intention-based approaches to dialog modeling ((Cohen and Levesque, 1990), (Lambert and Carberry, 1991), (Airenti et al., 1993)) arguing that, in non-cooperative settings (i.e. when the speakers do not already share a goal) these approaches leave unexplained why a speaker should bother adopting the goal of being cooperative both at the conversational level and at the behavioral one. In order to overcome these difficulties, (Traum and Allen, 1994) resort to the idea that speech acts pose obligations on the hearer. Obligations provide the motivation for the hearer to act even in non-cooperative situations.

In order to account for these remarks, we provide, in our model, two separate levels of knowledge for conversational and domain behavior (Ardissono et al., 1998b; Ardissono et al., 1999a; Ardissono et al., 1998a). The conversational level takes into account grounding, refusals, and other phenomena, while the domain level includes plans specifying how an agent can achieve his goals in a given domain. At both levels, however, the agent's behavior is driven by his intentions. So, an agent has to form conversational intentions as well as domain intentions out of a set of goals (either his own goals or a partner's goals), and, under some conditions, he decides to adopt the conversational goals of the partner without committing himself to the partner's domain goals.

It is the process of intention formation that determines an agent's behavior. When an agent decides to act, he must carefully (more or less) balance his interest in attaining the goal against the effort required to attain it and its compatibility with other goals. So, even in non-cooperative situations, the adoption of domain or dialog goals is explained in terms of intentions: social goals like preventing the speaker from becoming offended play a role in the decision process, leading an agent to adopt a goal of a speaker even if they do not already share any other goal.

In the next section we present the basic concepts that underlie the proposed model, and in Section 3, we show how they apply to the notion of politeness in agent interaction. Then, we show that, when conversational cooperation has been set up, the model predicts that agents may adopt different conversational strategies. The fifth section describes an implementation of the model and discusses an example. A Comparison with Related Work and a Conclusion Section close the paper.

2 The Underlying Model

In general, we assume that every agent has a set of goals, each associated with a relevance measure, representing the importance of the goal for the agent. Besides goals, an agent has a set of preferences towards states of affairs, with associated a utility function representing the advantage for the agent of (even partially) achieving a certain state of affairs. A goal may also occur among the preferences of the agent and, hence, it may be related to some utility function for evaluating the advantage of achieving it.

Goals and preferences, however, play different roles: goals are the input to the planning process, while the utility functions expresse the preferences of the agent among the plans which satisfy a certain set of goals. In this way, the advantage of a plan is not only evaluated with respect to the the fact that it achieves the goals the plan has been built for, but also with respect to the advantages provided by the side effects of the plan (for instance, consuming less resources).

So, the "goal situation" of an agent A, i.e. G_A, is defined as a set of couples:

$G_A = \{<g_1, r_1>, <g_2, r_2>, \dots <g_n, r_n>\}$

where g_i is a conjunction of literals, the relevances r_i are positive real numbers.

The preferences of an agent are expressed as a set of couples as well:

$U_A = \{<s_1, u_1>, <s_2, u_2>, \dots <s_m, u_m>\}$.

where s_i is a conjunction of literals and u_i is a function from world states (set of literals) to real numbers.

The overall utility function that an agent exploits to evaluate the world state S resulting from the execution of a plan is a weighted sum of the utility functions associated with each preference:

$UF_A(S) = \sum_{i=1}^{m} w_i u_i(S)$

The relevance expresses an a-priori evaluation from agent A about the increase of his well-being in case he decides to adopt a line of behavior aimed at achieving the goal expressed by g_i. Since an agent cannot at any moment build a plan for satisfying all of his goals, he has to focus on a subset of them: the relevance of a goal is the criterium for chosing to plan for satisfying only a subset of them. The set of currently focused goals R_A is composed of those goals in G_A whose relevance is greater than a given threshold r_A (that can be varied according to the resources he can allocate to planning and deciding what to do).

Moreover, each agent knows about a set of action schemas $AS_A = \{AS_1, AS_2, \dots, AS_k\}$ where each AS_i is a 5-tuple: $AS_i = <Name_i, Prec_i, Constr_i, Eff_i, Decomp_i>$:

1. $Name_i$ belongs to a predefined set of Action symbols, each with an associated arity.

2. $Prec_i$, $Constr_i$, Eff_i are conjunctions of literals (preconditions, constraints, and effects).

3. $Decomp_i$ is a set of sequences $Subact_{il}$, where the elements of $Subact_{il}$ belong to AS_A; each $Subact_{il}$ represents a possible decomposition of AS_i; if $Decomp_i$ is the empty set then AS_i is called Executable, otherwise it is called Complex.

On the basis of his focused goals R_A and of the actions he knows, an agent A can build a set of plans, by selecting the actions which have among their effects one (or more) of the goals in R_A. This is done by means of an algorithm which examines the possible decompositions of the actions. This algorithm builds trees of instantiated actions where, for each node AS_m, the (ordered) sequence of daughters corresponds to one of the decompositions in $Decomp_m$. Each such tree is called a Plan.
A Full Plan is a tree of instantiated actions where all the leaves of the tree are Executable Actions. A plan where some leaves are Complex is a Partial Plan.

According to the definition, an agent is able to obtain an evaluation of the utility of a plan both for Full Plans and for Partial Plans, since the utility is evaluated with respect to the resulting final state: the partial plans just abstract away from some of the side effects resulting from the different possible ways of refining the plan.

When the search algorithm concludes the search either because all plans have been refined or the time allocated to the planning process has been consumed, the agent may choose the plan with the greatest associated utility: he becomes committed to that plan, which constitutes his current intention I_A.

The process of intention formation for single agents can be summarized in the following way:

1. The set R_A is composed of those goals in G_A whose relevance is greater than the threshold r_A.
2. A builds the set of plans P for achieving the set of goals R_A.
3. A choose sthe p_i in P such that its outcome maximizes his utility.

The situation becomes more complex when a group of agent interact. In particular, it often happens that some goal g_B of agent B becomes known to agent A. A special case of this situation arises in case B has explicitly asked A for help. In these cases, it is possible that agent A comes to insert this goal among his intentions even if A has no preference towards this goal. The reason for doing so is the fact that doing something for achieving the goal of someone else can anyway have some utility for A: as stated above, a plan is evaluated also with respect to the side effects it produces.

What's more important is that the whole process of intention formation has now to take into account the presence of another agent. In particular, whenever an agent A has included among its goals a goal originating from another agent B, we say that A is pre-cooperating with B. Notice that, in general, it may happen that A will never do anything to satisfy g_B. However, the fact that A takes into account g_B is a (weak) form of cooperation: at least the agent spends his time and allocates planning to take into consideration the alternative of adopting the partner's goal. In case A chooses a plan that leads to the satisfaction of the partner's goal, we say that A cooperates with B. In case that B is aware of the formation of this intention, then the agents are cooperating to a shared goal (Boella et al., 1999).

Besides considering the possibility of doing something for achieving other agents' goals, in an interaction process, an agent should try to foresee their possible reactions. Not only in case their goals have been adopted and achieved but also in case the agent decides that it is not worth for him to do anything for satisfying those goals.
Only by considering the situation resulting after the reaction of other, the agent can really assess the advantage of a given plan.

A kind of "anticipation feedback" of the reaction of the partner is, therefore, required. (Ndiaye and Jameson, 1996) has adopted this form of reasoning in a dialog system in order to evaluate the goodness of the alternative moves of the system. In (Boella et al., 1999), we have exploited a similar form of reasoning in order to evaluate the goodness of an action during the execution of a shared plan by a group of agents. In this paper, we will exploit the ability of an agent to foresee the possible reactions of another agent in order to decide whether it is worth for him to decide to satisfy the other agent's goal.

If an agent evaluated the utility of a plan that achieves a goal requested by another agent only on the basis of its outcome, he would never choose that plan in a non-cooperative setting (i.e., the agent does not already share any goal or preference with the other agent): achieving it does not produce any utility for him. As stated above, the utility of a plan is provided by summing the utility deriving from the preferences of the agent.
Often, performing an action for achieving another agent's goal results only in a negative utility, since the side effects of the action affects other states as resource consumption. In case of cooperation, instead, the partners share some goal which is preferred by both of them, so that achieving it increases the agent's own utility and the utility of the group's as well.

By evaluating the utility of a plan according to the world state resulting after the other agent's reaction, the agent can evaluate how this reaction affects his preferences, for instance not offending the partner and other social goals.

The planning process of the agent in a situation involving interaction, therefore, must also include considering the alternative of doing something for the other agent when his goals are recognized and evaluating the utility of the world state resulting from the partner's reaction.

Considering a goal of the partner during the planning process means assigning it a sort of relevance since planning is one of the resources of the agent. Moreover, this assignment of relevance must not necessarily result in the intention of performing the goal: the agent chooses to achieve the goal only if he can achieve some more utility in doing it with respect to not doing it.

The intention formation process in case of interaction with other agents is the following one:

1. The set R_A is composed of those goals in G_A whose relevance is greater than the threshold r_A.

2. If A knows that there is a goal g_B of agent B such that *relevance-eval*(g_B) $> r_A$, then he creates set R'_A as the union of $\{g_B\}$ and R_A.

3. A builds the set of plans P as the union of the plans resulting from building plans for achieving the set of goals R_A and R'_A.

4. For each plan p_i in P, A considers the possible reaction of B: the world state resulting from the reaction becomes the outcome of p_i.

5. A chooses the p_i in P whose outcome maximizes his utility.

The assignment of relevance by the function *relevance-eval* can be conditioned to the situation: the agent can have no time to waste for considering if it is useful to cooperate or he can know by experience that some goals can be discharged without consideration.
Note that the behavior of considering the goal of other agents does not derive from some social obligation, but it is just relevant in the sense that an agent thinks he can gain a better utility in doing so.

Within this rather wide framework, this paper focuses on the role of intentions on two rather specific problems, i.e. the adoption of other agents' goals during the conversation.

In particular, we claim that:

1. In interactive settings, the involved goals are of three types: Linguistic goals, Conversational Goals, and Domain goals.

2. Agents consider only a limited number of goals at a time depending on their relevance.

3. Agents choose what to do on the basis of their ability to foresee how a given plan changes the situation (i.e on the basis of the utility of that plan).

4. In a group activity, the evaluation of the utility of a plan must involve also the effects resulting from the reaction of the other members of the group.

5. These effects can be accounted for by trying to foresee the actions that a partner will undertake in response of an agent's action (*Anticipation Feedback*).

3 To Be Polite or not to Be

In this section, we assume that:

1. There is an agent A, who is carrying out some plan on the basis of a current intention. According to the model outlined in the previous section, this means that A has previously chosen a plan P_A, which maximizes a utility function.

2. Another agent B has addressed a request to A of satisfying a goal g_B.

The goal of this section is to explain how A can choose to be polite or impolite by adopting or not the goal of B. In particular, since there is, in principle, no utility for A in case B achieves one of his goals, the aim is to show why, in some (possibly most) cases, A chooses a polite line of behavior.

We have mentioned a "request" addressed by B to A. This should not be viewed as a restriction on the kind of thing that B could have said to A. In fact, we consider a question as a request to provide an agent with some information and a statement as a request to the hearer to update his knowledge base.

As stated above, the model involves three types of goals. In the case of a request they are:

1. Linguistic goals: the goal to understand the propositional content of what B has said, together with the subsidiary goals of hearing the voice (in case of vocal communication) of B, of understanding the words composing the sentence uttered by B and of putting the words together in a meaningful expression.

2. Conversational goals: the goal of maintaining the communication channel open, together with the subsidiary goals of keeping B informed that you have understood (grounding) and of maintaining the coherence of the dialogue.

3. Domain goals: any goal of B concerning the external world extracted from the comprehension of the current utterance or of the sequence of utterances composing the dialogue.

We are not going here to describe how these various goals are kept together in a unified structure: in (Ardissono et al., 1999a) we have shown that the problem solving activity expressed as *Agent Modeling Plans* could constitute the required glue; we just assume that as soon as A hears B, these goals are set up inserting them among the current set of goals of A if they are considered relevant.
Of course, domain goals are taken into account only as soon as they are recognized, i.e. depending on the success of the linguistic goals.

Why is it much more common that an agent refuses to cooperate at the domain level than at the conversational level? We claim that the choice of cooperating does not depend on obligations established by the speech acts (as (Traum and Allen, 1994) claim instead); on the contrary, it depends on the type of the action needed to establish cooperation: during the process of forming intentions out of goals, if the cost of the required action is low, a refusal can be motivated only by a negative attitude towards the requester. Therefore, the requester, which can infer the requestee's reasoning, will be offended by a refusal.

The effect on the requester can be foreseen by the planing agent in the anticipation of his reaction: the agent simulates the reaction of the partner depending on his choice to adopt or not to adopt the requested goal.

This explanation is in line with the notion of politeness as defined in (Brown and Levinson, 1987). They show that politeness can be explained with reference to the notion of face. In (Ardissono et al., 1999b), we have provided a formal explanation of how indirect speech acts prevent the hearer from being offended by referring to the interactants' face wants; moreover, we have described how the type of the requested actions affects the offensiveness of a request; compare, e.g., asking the time vs. requesting some money.

In a similar way, we can explain why refusing cooperation at a conversational level may as well threaten the speaker's face wants: paying attention to people, listening and understanding them are “free goods” (Goffman, 1967), so no one can refuse them without threatening the speaker's face. Indeed, ignoring someone and not responding to his speech act is usually perceived as a very aggressive behavior.

In our model, the social goals are expressed as the preference of not offending the other agents, a preference with associated a certain benefit: this utility is part of the overall utility function of the agent, hence it is used to select plans that do not have offending as a side effect.

Note that offending is not a direct effect of not adopting a goal or of action as refusing cooperation; instead, the state in which the partner feels offended is a result of the reaction of the partner when he is faced with the refusal of a low cost action.

As we showed in the previous section, if the goal is relevant, the agent has to compare the plans that achieve just his private goals, with those satisfying the partner's goal as well. Moreover, he will consider also the reaction of the partner in case the goal has been adopted or not. Hence, he has the information for deciding whether it is better to save his own resources and risk to offend the partner, or to spend some of his efforts and preserve the social relationship with him. The choice of the agent depends on the relative weighting of the utilities deriving from the preferences of saving resources and not offending the partner.

There is another situation, however, in which the other agent's goal are not adopted: when, for some reason, a low relevance is assigned to that goal. This behavior expresses the fact that the agent does not consider worthwhile to allocate planning resources for deciding what to do about the goal.

So, the basic tenets of the model are that:

- the relevance is determined without involving means-ends reasoning, but just on the basis of an a-priori evaluation of goals (possibly based on learned rules);
- the utility is determined involving means-end rationality, and it is what always determines the actual behavior of agents,

  but
- the relevance can be assigned values that prevent the evaluation of utility, thus producing (partially) "irrational" behavior (that is, had the agent have enough resources to evaluate all possible plans, he would have found that doing something for achieving the partner's goal was the optimal solution).

A last point needs to be made clear before going back to the computational model. *Relevance-eval* is a heuristic function: it is not, and should not be, independent of the final utility evaluation. This is rather obvious for single-agent intentions: if the goal g is assigned high relevance, then it means that the agent believes that it affects very positively his well-being. So, in the evaluation of any plan leading to g, a highly positive contribution to the global utility should come from the truth, in the resulting state, of g.

Less obvious is the fact that also in two-agents settings the same must happen (if *Relevance-eval* is a good heuristics). If A is used to take into account in its utility evaluation the positive or negative reactions of his partner(s), then the *Relevance-eval* function he uses must assign a reasonable level of relevance to the partner's goals (otherwise his utility function is of no use). Conversely, if *Relevance-eval* assigns high relevance to other agents' goals, the utility function must take into account the partner's reaction (otherwise, considering the partner's goal involves a waste of resources in computing the utility of plans that will seldom be chosen).

4 Conversational Phenomena Accounted for by the Model

We have discussed in the previous section the motivations for starting up conversational cooperation. Once an agent has decided to cooperate with another agent at the conversational level, a group is formed around the conversational shared goals. In our model, shared goals don't compel agents to act, but, instead, within the context of a group, they drive the evaluation of the utility of alternative actions on the part of the agents, given the probability that the actions result in certain outcomes.

In general, we assume that - homogeneously with other behavioral levels - the conversational behavior of an agent can be modeled by a library of conversational plans including actions that encode the knowledge about speech acts, adjacency pairs, and so on, and also optional grounding actions, like notifications, requests for notifications, repairs to misunderstanding, and so on (see (Ardissono et al., 1999a)). Here, we don't address the problem of how an agent chooses the form or the content of his speech-acts; instead, we focus on the intention generation process. In the case of conversational cooperation, our model prescribes that the agent builds and evaluates two alternative plans: the first of them includes the conversational action under consideration, and we will call it *rich plan* (*RP*); the second one does not include it, and we call it *simple plan* (*SP*).

The presence of the conversational goal in the *RP*s is in some cases explained by the fact that the agent has the goal of checking whether his action succeeded (for instance, the success of a speech act), as well as that, on the other side, the partner considers to adopt this goal when he knows that the other agent wants to achieve it.

These plans are then evaluated in the light of the group's utility, and the *SP* or the *RP* is chosen depending on the result of the evaluation: when evaluating the utility of a plan, the agent takes his partner potential reaction into account.

4.1 The goal of understanding a speech act

As long as the conversation proceeds smoothly, the task of checking the partner's understanding in a conversation is normally accomplished by each participants by exploiting the implicit feedback that the partner's next turn provides. However, participants sometimes detect difficulties in understanding the partner, or in making themselves clear, and, in some circumstances, they act in order to prevent such difficulties, when possible. In all the cases, the advantage that is more or less directly traded off (by means of the utility function) against the effort required to undertake some kind of grounding action is the predictable saving of resources that stems from avoiding, in the future, a clarification or negotiation phase to restore the lost common ground.

- The speaker (A), given his beliefs about the common interpretation of the interaction up to that moment, is normally aware of the difficulties the hearer (B) may find in interpreting his turn. Moreover, he has the goal to know whether B succeeded in his interpretation. So, A can decide to take the initiative and explicitly check B's understanding; in order to do so, A may solicit feedback using discourse markers like "right?", "ok?", etc. Based on these factors (predicted difficulties, importance of the turn, etc.), A evaluates the utility of the *RP* in which he checks the B's interpretation and the utility of the *SP* where he doesn't. In case the *RP* gets a higher utility value, then A will commit to the plan that includes an "ask for notification" action.

- Knowing that A has the goal of checking whether B's action of understanding succeeded, B sometimes adopts this goal and notifies to A the success of the interpretation process ("mhm", "right" and repetitions) or his failure ("what?"). However, since *B*'s next turn usually displays his understanding, these notifications are often by-passed. From the perspective of our model, *B* has to compare the utility of the *RP* including the action "Ground" with the utility of the *SP* consisting in going on with the new turn without any notification.

- Finally, *A* often tries to facilitate *B* in his understanding process. In a coherent dialog, each new turn is usually related in some way to the current dialog focus (Carberry, 1990). In general, topic shifts are costly for *B* because in order to identify them, he has to search for a link within the context starting from the current focus. Moreover, the clearer *A*'s turn, the higher the probability that *B* understands it correctly, with a valuable saving of resources for the group cooperating at the conversational level. *A*'s *RP* consists in a plan where he exploits the linguistic resources at his disposal to make the focus shift apparent, while the *SP* is one where he simply changes the focus without explicitly signaling it to *B*. If *RP* yields a higher utility, because *B* can interpret the sentence with less effort, *A* will introduce cue phrases like "however", or will notify that a new topic is introduced ("by the way").

4.2 The goal of maintaining a common interpretation

For the dialog to proceed coherently, the interactants need to share the same interpretation of the previous part of the interaction, or better, the mutual belief must hold that their respective interpretations are reasonably aligned.[1]

The lack of interpretation problems is, by itself, a symptom that there is a common interpretation, i.e., that participants' interpretations are reasonably aligned. Of course, this does not mean that the interpretation is really the same, but only that the possible differences fall within the standard individual differences (Ardissono et al., 1998a). In absence of specific signs of misalignment, this normally makes the utility of any active effort to check the interpretations' alignment very low, because the effort required would not be compensated by the low risk of having to initiate a negotiation phase in the following, to restore the lost alignment.

On the contrary, the loss of dialog coherence normally means that a misunderstanding has occurred, i.e. that at least one of the two speakers has chosen a wrong interpretation of a turn (Ardissono et al., 1998a). If a participant *A* realizes that a misunderstanding has occurred, he will form and compare two alternative plans: the *RP* includes a appropriate action for addressing and solving the misunderstanding, the *SP* has no extra action to address misunderstanding and consists in simply going on with the next turn. If the comparison between the two alternatives ends in favor of the former one, the participant will act in order to find out who is the misinterpreting agent, and correct the wrong interpretation. Intuitively, the utility function here embodies the idea that a misunderstanding is addressed if it is deemed relevant, and is at risk of posing difficulties for the subsequent interaction.

1. if *A* realizes that his *B* is the misinterpreting agent, the *RP* will include a request for repair ("No, I mean that..."); *B*, in turn, is expected to repair his

interpretation as requested, and notify the execution of the repair to A ("Oh, ok ...").

2. if A realizes that he has misinterpreted B, his RP will include the repair and the notification to B that now he holds the right interpretation ("Oh, you meant that ...").

4.3 Other conversational phenomena

In our model, a number of other phenomena occurring in a conversational context stem from behavioral cooperation, i.e. once it has been established the shared goal of performing a (linguistic or domain) action requested by the partner:

- The hearer's notification that it is impossible to perform the requested action:

  A: Could you tell me the time, please?
  B: I'm sorry, I don't have my watch on me.

  In this example the RP includes the notification that the requested action is impossibile to perform. These notifications are captured in our model of cooperation (Boella et al., 1999) by the fact that informing the partner that requested action cannot be performed usually results in a higher group utility, because it prevents the other group members from wasting further efforts in vain (for instance, by repeating the request).

- Conversely, B often notifies that he has done the requested action:

  A: Could you press the key for me?
  B: Done!

- When B has adopted the speaker's perlocutionary goal that he (B) intends to update a certain belief, B sometimes notifies that he believes what has been communicated, possibly after a solicitation by A:

  A: we have a boxcar of oranges in Amsterdam, right?
  B: right

  Here, A's choice to include in his plan the solicitation of a feedback is motivated by his goal to check that B, after recognizing his communicative goal (the mutual belief that the A wants B to intend to adopt the belief conveyed by the assertion), actually tries and updates his "knowledge base" with this belief.[2]
  B knows the partner's goal, so, he can consider to adopt it even if he is not solicited: providing spontaneously the notification results in a less costly dialog in that A has not to check explicitly whether B succeeded.

- Negative notifications of update are motivated by the same goal:

  A: obligations do not exist
  B: I don't think so

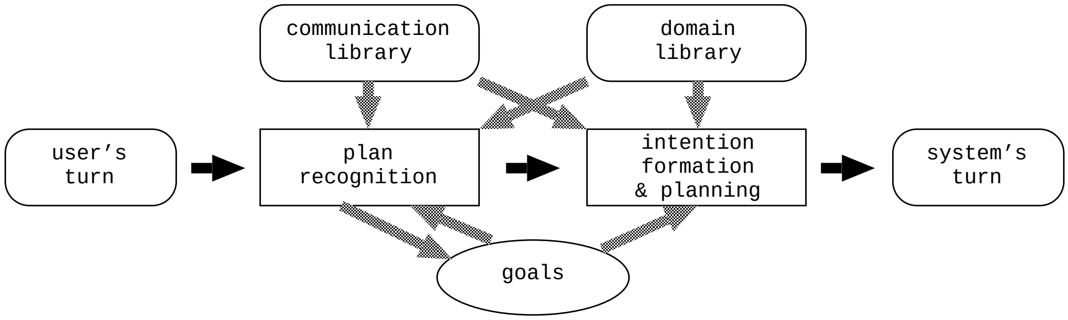

Figure 1: The architecture of the dialog system.

5 The Agent Architecture

The dialog component of our system is described in Figure 1. Although a separation is shown between the plans for achieving Linguistic and Conversational goals (Communication Library) plans and the plans concerning domain actions (Domain Library), the internal organization of the two libraries is exactly the same, the main difference being that the communication library is domain-independent. Both libraries are controlled by a reactive planner, which selects the actions that the agent will do next; the selected action constitutes his intention.

As stated in Section 2, the system (S) applies a planning process driven by the goals and the utilities associated with possible courses of actions. When S realizes that another agent (U) wants S to adopt a new goal by means of the plan-recognition process (Ardissono et al., 1999a), S applies *Relevance-eval* in order to get the relevance value to associate with that goal. If the relevance is greater than the current threshold r_S, then S inserts the goal in the set of the goals to consider, in order to see whether it is worthwhile to adopt it.

The box in Figure 1 labeled "intention formation & planning" is expanded in Figure 2 in order to show more precisely how the utility of a plan is evaluated.

S, given his goals and the one identified during the plan recognition phase, selects some applicable plans and considers the possible outcomes deriving from their execution (S_1, ... S_n). If he limited himself to the outcome of his own actions, when evaluating the alternative plans, he would never adopt any goal from U, since this would always result in a waste of resources, without any other immediate benefit.

Therefore, S has to consider also the potential reactions of U, given U's goal that S performs the action U requested (see the connection in Figure 2 between the recognized goal and the simulation of the planning of S): this is done by means of another instance of the agent architecture which emploies U's point of view. The simulation of U's reaction results in a new set of outcomes: only by considering the utility of these outcomes, S can assess the advantage of his choice; in fact, the reaction of U may affect some of the social goals of the system (for instance, the desire not to offend the partner). Intuitively, not adopting one of U's goals may result in U being in a bad-tempered mood, a fact that could have serious consequences for S. So, the adoption of the goal prevents the offense of the partner, and a social resource is saved at the expense of some effort for executing the action. A similar algorithm, but applied to cooperative settings, has been described in (Boella et al., 1999).

In order to verify the feasibility of exploiting social goals for motivating cooperation, we have implemented a prototype using the decision theoretic planner DRIPS (Haddawy and Hanks, 1998). DRIPS exploits hierarchical plans to find the optimal sequence of actions under uncertainty, based on a utility function. Goals can be partially satisfied, by trading them off against cost (waste of resources) and against each other.[3]

In the experiments carried out to validate the model, we have introduced four different attributes to depict the situation in which S is interrupted by U, while he is executing an action Act formed by two steps $Step_1$ and $Step_2$:

- *time*: it models the time as a bounded resource. The utility gets a fixed value in case its two steps are completed before time t_{end}, otherwise, the utility decreases more or less sharply when the action is completed after the deadline as a function of *time*;
- *ground*: it models the U's goal of knowing that S has successfully interpreted the request; when it is true, U infers that S is cooperating at the conversational level.
- *goal*: the action requested by U has been successfully executed;
- *res*: it models the use of (generic) resources. The utility decreases when the value of *res* decreases;
- *offended*: it models the (foreseen) degree of offense of U. The utility decreases when the value of *offended* increases;

Moreover, we have introduced different actions, the most important of which are:

- Action *Ground*: it makes U know that his request has been successfully interpreted.
- Action *Adopt*: it represents S's cooperation with U at the domain level; the decrease of available time and of generic resources depends on the particular action needed to satisfy U's request.
- Action *Refuse*: it represents S's act of communicating to U that he does not intend to do what U requested of him. Among its effects, there is the fact that U comes to know about S's decision (*refused*) and that *ground* is true.

Other actions concern the possibility of providing U with extra help (by having recognized some of his implicit goals).

On the side of U, we have introduced the single action *Evaluate* (see Figure 4), modeling the change of the *offended* parameter, depending on S's action. Notice that

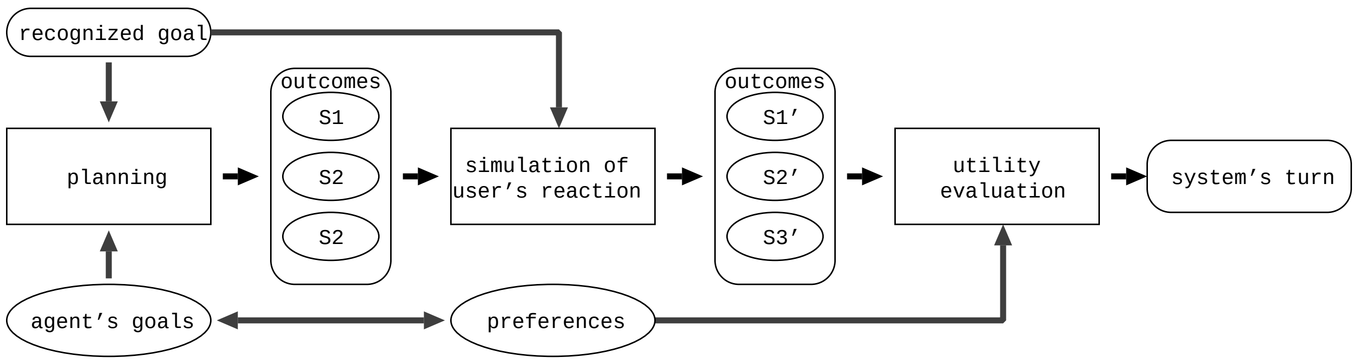

Figure 2: The architecture of the planning module.

```
(action Ground
        (time = time + 1)
        (res = res - 1)
        (ground = 1))

(action Adopt
        (time = time + length(action))
        (res = res - cost(action))
        (ground = 1)
        (goal = 1))

(action Refuse
        (time = time + 5)
        (res = res - 5)
        (ground = 1)
        (refused = 1))
```

Figure 3: Some of the system's actions. The *ground* parameter represents the fact that the goal of grounding the new turn has been accomplished.

what *Evaluate* actually represents is how S believes U will react to his choice. The key parameter affecting the level of offense is the cost of the requested actions: the less the cost of the requested action, the greater the offense; this follows the principle by (Goffman, 1967) that people get offended when they are refused low-cost help.

The *goal* predicate represents the fact that the effect of the requested action has been achieved. The lack of grounding is interpreted by U as S is not cooperating at the conversational level: since cooperating at the conversational level (interpreting the sentence, grounding it) has a low cost, it is offensive not to do it. Note that, besides the *Ground* action, also the adoption of the domain goal or a refusal show S's cooperation at the conversational level. On the other hand, in case of domain action there are different possible costs and also differences between the results of cost evaluation by the two interactants (see the parameter *action*). The *higher-goal* parameter in *Evaluate* is set to 1 as an effect of the action *Extra-Help* mentioned above.

Now let's consider in detail the current situation, i.e, the one where U has just asked to S to do something while S was going to perform the first step of *Act*. Assume that the heuristic function of S (*Relevance-eval*) has returned a value such that U's recognized goal is considered during the subsequent planning phase. In order to explore the different alternatives, the planner builds and evaluates some plans. These plans differ in the actions for pursuing the partner's recognized goal: they can be included or omitted. It then tries to foresee the reaction of U, and commits to the plan which yields the greater utility according to S's preferences.
The experiments have shown how a change in some parameters affects the behavior of the system:

- By associating a greater utility to the completion of *Act* than to *offended*, and by keeping t_{end} tight (i.e. t_{end} coincides with the sum of the durations of $Step_1$ and $Step_2$) we obtained that S decided not to cooperate at the domain level. The

```
(action Evaluate
       (time = time + 1)
       (res = res - 1)
       (offended = offended + (20 * not(ground)) +
                              (20 * not(goal) / cost(action)))) -
                              (20 * higher-goal))
```

Figure 4: The partner's reaction.

Ground action was executed depending on how fast the utility of *Act* goes down to zero after the deadline and on the utility of not offending the partner.

- By relaxing t_{end}, so that S has time enough to answer to U and complete *Act* within the time constraint, we obtain a fully cooperative behavior: the sequence <*Adopt*, $Step_1$, $Step_2$> was the preferred new plan (the requested action is not expensive and the grounding is assured by adoption);
- By increasing the utility of the resource *res* with respect to *offended*, no adoption was produced, or at best the planner explicitly refused behavioral cooperation: the agent was more worried about not wasting his efforts than about preventing U from being offended.

6 Related Work

In order to overcome the difficulties of existing intention-based approaches, (Traum and Allen, 1994) defined an interaction model based on the notion of obligation. Obligations are pro-attitudes that lead humans to act but impose less commitment than intentions; their social character explains why humans are solicited to act. (Poesio and Traum, 1998) provide an axiomatization of speech acts where communicative actions have as an effect the imposition of obligations on the two interactants.

Our model has some advantages with respect to obligation-based approaches. First of all, the same phenomena are accounted for without introducing further propositional attitudes. Moreover, the notion of conversational cooperation together with the definition of shared goal (Boella et al., 1999) provides an explicative treatment of these different phenomena that is uniform with the rest of the model.
(Traum and Allen, 1994) claim that discourse obligations are learned social norms; on the contrary, in our model, the motivation for adopting a conversational goal is given by the goal of preventing a potential offense for the speaker together with its potential consequences. All the speaker has to learn is under what conditions humans happen to be offended, and this same knowledge explains the use of indirect speech acts (see (Ardissono et al., 1999b)).

In (Traum and Hinkelman, 1992) speech acts are modeled as "multiagent actions [...] that are not completed unless/until they are grounded". Instead, in our model, speech acts are single agent actions; similar to domain actions, they succeed only if their effects hold, so agents check in the world looking whether the effects hold. There is no clue that linguistic actions should be modeled in a different way from domain ones.

The only difference is due to the fact that the effect produced by a linguistic action is sometime difficult to detect: the effect of an illocutionary act is to make the speaker's communicative intentions mutually believed. In some cases, the only viable way the speaker has at his disposal to understand whether the linguistic action succeeded is some form of explicit grounding by the hearer.
Finally, the ordering among the alternative sources of deliberation is provided by the means-ends relation involving the different linguistic actions, instead to be explicitly stated as in the model of (Traum and Allen, 1994); for example, explicit grounding actions often are avoided in presence of answering: the explanation relies on the fact that, if the hearer has answered coherently, then the goal of showing his understanding of the question has been achieved and an explicit grounding action is superfluous and results in a further consumption of resources.

The strategy of anticipating the partner's potential reaction, when choosing the next action to be performed in a conversational context, has been exploited also by (Ndiaye and Jameson, 1996) in the dialog system Pracma. In this work, the authors propose a technique called "global anticipation feedback" to select the next system's move, in association with an evaluation of the utility of the available moves; this technique consists in the system temporarily assuming the role of the user, in order to make an assessment of the potential user's reaction to a given system's move. Once the user's reaction has been generated, for each alternative move available to the system, the utility of both system's and user's move is evaluated, and the move that maximizes the sum of the two is chosen. However, in both cases the utility esteem is made from the point of view of the system; moreover, utility is an attribute of the move, i.e., is not evaluated on the basis of explicit goals motivating the system's behavior.

In order to make the system's choices more principled at a higher level, an obligation based approach has been introduced in (Jameson and Weis, 1996), based on (Traum and Allen, 1994). In this model, obligations can be set by certain move types, during the course of interaction, and their status, ranging from acceptance to rejection, can be subsequently affected by other moves. Every move has some methods associated to it, that allows to specify what obligations it creates or addresses. Obligations, in turn, have some properties like a penalty that is imposed to the system for rejecting the obligation itself, or, possibly, an expiration period. For each obligation, the moves that address can be specified, together with the corresponding changes in the status of the obligation. This knowledge is exploited to constrain the system's planning of the next move, in order to imitate the way human participants address obligations in real interactions, especially in non cooperative settings. However, the obligations considered here are domain-dependent, since they are tailored to a particular dialog type, and they utility with respect to the system's goals is accounted for only in an implicit way. The utility of a move is then evaluated with reference to how it affects discourse obligations, both existing and new ones. In order to show the improvement resulting from taking obligations into account, the authors depict different behavior types, from the simplest, where obligations are not considered, to the most complex, where the system accounts for the obligations of both speaker and hearer. However, no anticipation of the moves of the speaker is used in the latter hypothesis.

Although anticipation feedback and dialog obligations can be easily integrated in the same system, none of the two techniques constitutes, by itself, a general and uniform explanation of the conversational behavior of the system.

7 Conclusions

In this paper we proposed an intention-based approach to dialog that aims to overcome the critics posed by (Traum and Allen, 1994). Our solution does not employ the notion of obligation, even if some similarities can be found with (Traum and Allen, 1994), since they use obligation for explaining why agents adopt goals of other agents, even in non-cooperative settings: the reason for not exploiting the notion of obligation is to have a uniform source of motivations for explaining the behavior of agents.

Some considerations on possible critics must be done. First, reducing the number of propositional attitudes, the reasoning process becomes more complex; our model wants to be an explanation and it does not exclude the possibility of compiling the reasoning in more compact form: as (Brown and Levinson, 1987) notice for what it concerns their model of politeness, "there is a rational basis for convention".

Second, here we do not deny the existence of obligations in general: the social character of some dialog behaviors which in (Traum and Allen, 1994) are explained as obligations is accounted for by the reference to social goals (for instance, not offending the partner).

Notes

[1]It can be argued that speakers don't pursue the achievement of mutual belief that common understanding was unachievable, as a consequence of the shared goal to reach mutual understanding. However, when a misunderstanding is not repairable, notifications of the breakdown actually occur (Ardissono et al., 1998a).

[2]Note that the same phenomenon is caused by the shared goal of executing domain actions; e.g. consider "Save the file ... Have you done it?".

[3]In the following example, the DRIPS capability of managing with probabilistic effects and expected utility is not exploited for the sake of brevity.

References

Airenti, G., Bara, B., and Colombetti, M. (1993). Conversational and behavior games in the pragmatics of discourse. *Cognitive Science*, 17:197–256.

Ardissono, L., Boella, G., and Damiano, R. (1998a). A plan-based model of misunderstandings in cooperative dialogue. *International Journal of Human-Computer Studies*, 48:649–679.

Ardissono, L., Boella, G., and Lesmo, L. (1998b). An agent architecture for NL dialog modeling. In *Proc. Second Workshop on Human-Computer Conversation*, Bellagio, Italy.

Ardissono, L., Boella, G., and Lesmo, L. (1999a). Plan based agent architecture for interpreting natural language dialogue. *International Journal of Human-Computer Studies*, to appear.

Ardissono, L., Boella, G., and Lesmo, L. (1999b). The role of social goals in planning polite speech acts. Workshop at 7th Int. Conf. on User Modeling.

Boella, G., Damiano, R., and Lesmo, L. (1999). Cooperating to the group's utility. In *Proc. of ATAL Workshop*, Orlando, FL.

Brown, P. and Levinson, S. C. (1987). *Politeness: some universals on language usage.* Cambridge University Press, Cambridge.

Carberry, S. (1990). *Plan Recognition in Natural Language Dialogue.* MIT Press.

Cohen, P. and Levesque, H. (1990). Rational interaction as the basis for communication. In Cohen, P., Morgan, J., and Pollack, M., editor, *Intentions in communication*, pages 221–255. MIT Press.

Goffman, E. (1967). *Interaction Ritual.* Penguin, Harmondsworth.

Haddawy, P. and Hanks, S. (1998). Utility models for goal-directed, decision-theoretic planners. *Computational Intelligence*, 14:392–429.

Jameson, A. and Weis, T. (1996). How to juggle discourse obligations. In *Proc. of the Symposium on Conceptual and Semantic Knowledge in Language Generation*, pages 259–267, Heidelberg.

Lambert, L. and Carberry, S. (1991). A tripartite plan-based model of dialogue. In *Proc. 29th Annual Meeting of ACL*, pages 47–54, Berkeley, CA.

Ndiaye, A. and Jameson, A. (1996). Predictive role taking in dialog: global anticipation feedback based on transmutability. In *Proc. 5th Int. Conf. on User Modeling*, pages 137–144, Kailua-Kona, Hawaii.

Poesio, M. and Traum, D. (1998). Towards an axiomatization of dialogue acts. In *Proc. of 13th Twente Workshop on Language Technology*, pages 207–222, Enschede.

Traum, D. and Allen, J. (1994). Discourse obligations in dialogue processing. In *Proc. 32nd Annual Meeting of ACL*, pages 1–8, Las Cruces, New Mexico.

Traum, D. and Hinkelman, E. (1992). Conversation acts in task-oriented spoken dialogue. *Computational Intelligence*, 8(3):575–599.

Acknowledgements

We would like to thank Massimo Poesio for the interesting discussion and his suggestions that made this work possible and the anonymous referees that let us improve the final version of this paper.

Authors

Guido Boella, Rossana Damiano, Leonardo Lesmo and Liliana Ardissono
Dipartimento di Informatica - Università di Torino
C.so Svizzera 185, 10149 Torino ITALY
tel. +39-011-6706711
fax. +39-011-751603
email: {guido,rossana,lesmo,liliana}@di.unito.it