



A model for negotiation in teaching-learning dialogues

Michael J. Baker

► To cite this version:

Michael J. Baker. A model for negotiation in teaching-learning dialogues. Journal of Artificial Intelligence and Education, 1994, 5(2), pp.199-254. hal-00190394

HAL Id: hal-00190394

<https://telearn.hal.science/hal-00190394>

Submitted on 23 Nov 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Model for Negotiation in Teaching-Learning Dialogues

MICHAEL BAKER

*Centre National de la Recherche Scientifique
IRPEACS Laboratory, ENSL, Lyon 69364, France*

Despite the fact that a number of researchers from a wide variety of theoretical perspectives have claimed that negotiation in dialogue is fundamental in learning, little research has been based on a precise notion of what negotiation is and how it relates to learning. This paper describes a model for negotiation, based on analysis of teacher-student and learner-learner dialogues. After discussion of illustrative examples and review of related research, negotiations are defined in terms of four features: what is being negotiated, the initial and final states of negotiation, and negotiation processes themselves. The paper concentrates on negotiation processes, for which a model is developed based on a set of communicative acts and a set of different types of relations between offered propositions. It is argued that the propositional attitude most relevant to negotiation dialogues is *acceptance* rather than *belief*, a claim that has implications for student modelling approaches based on the latter concept. In conclusion, it is claimed that negotiation plays the following principal roles in teaching-learning dialogues: *coordinating* problem-solving and communicative action, establishing *mutual understanding* with respect to possible solutions and *co-constructing* agreed solutions themselves.

INTRODUCTION

We all have an intuitive idea of what negotiation is. The most common everyday case would be that of negotiating a *price* for buying or selling something (a second-hand car, for example). In such a negotiation we would expect each person to make offers which may eventually converge on an

The idea that learning may be negotiated is not, of course, new. It closely relates to what Bennett (1976) termed the "progressive approach" to teaching, which Bruner (1986) described a decade later as follows:

...induction into the culture through education, if it is to prepare the young for life as lived, should also partake of the spirit of a forum, of negotiation, of the recreating of meaning. But this conclusion runs counter to traditions of pedagogy that derive from another time, another interpretation of culture, another conception of authority—one that looked at the process of education as a *transmission* of knowledge and values. (Bruner, 1986, p. 123)

The usual theoretical position in opposition to which the "learning as negotiation" approach is often defined is therefore that of "learning as transmission of knowledge." In fact, if we look at the cognitive science and learning literature, it is quite hard to find the kind of bald statement of the "knowledge transmission" view which its opponents would like to attack. Very often, the "transmission" view is implicitly assumed by much research, and may only be abstracted from it with some difficulty. The view may be found in the classic collection of papers in Stecman and Brown (1982) to the extent that in work which did take interaction into account it was tacitly assumed that the student's goal was to acquire the system's domain model (Collins & Stevens' WHY system) or that the objective was for the system to tailor its explanations for transmission to the student (Clancey). The transmission idea has also resurfaced in Wenger's well-known text (Wenger, 1987) in the form of "knowledge communication." There is a very general sense in which the "knowledge transmission/communication" view is correct, however: Much education is concerned with enabling the student to acquire knowledge as a result of interacting with the teacher. The point is that, except in very special cases, transmission alone is unlikely to enable students' to integrate what is transmitted with their pre-existing knowledge (constructivism). More recently, a number of researchers in cognitive science have claimed that negotiation has something to do with learning. For example, Resnick et al. (1991) state that:

Within cognitive science, challenges to the dominant role of formal logic as an account of human reasoning are producing growing concern for the role of argument, contradiction and negotiation processes. (Resnick et al., 1991, p. 388)

agreed price. A partner who does not want to move from a given price may attempt to argue as to why their offer is "reasonable" or "fair." From this everyday notion of negotiation we can see that: the common goal of negotiation is to reach an *agreement* (compromise on a final price); negotiating agents may have individual and competing goals (agreeing highest or lowest price); negotiations consist basically of sequences of *offers* that may be *accepted* or *rejected*; and two possible negotiation strategies are to refine original offers towards agreement, or to keep an offer fixed and to attempt to persuade the other to accept it by *argumentation*. Negotiations may clearly have other "objects" than prices, such as goals, beliefs, and so on (we refer to the "thing(s) negotiated" as *negotia/negotium*, for lack of existing terms), and in negotiations, successive offers are usually closely *related* to each other in specific ways. As a first approximation, we can therefore characterise a negotiation in terms of four features: (a) the *negotia* (what is negotiated?); (b) the *negotiation initial state* (goals, beliefs, conflicts, etc.); (c) the *negotiation final state* (what is agreed?); and (d) *negotiation processes* (what processes lead from 2 to 3?). This basic analysis schema forms the basis of the model for negotiation to be described later in the paper. Before going more deeply into the nature of negotiation, we must first address the question: What does negotiation have to do with *learning* and *education*?

The Concept of Negotiation in Research on Teaching and Learning

A first guess concerning the role of negotiation in learning would be to recall to mind an enlightened and considerate teacher who took students' views and aims sufficiently seriously to be willing to freely discuss the way in which the teaching interaction should proceed and what might count as an acceptable problem solution. More cynically, it might conjure up the image of a teacher who made a show of such consideration whilst at the same time guiding the interaction along a predefined path (directive teaching with a sympathetic face). In terms of the four features of negotiations mentioned above, negotiations in the specific case of teaching-learning dialogues might be assumed to have the following general characteristics: (a) the *negotia* will be learning and interaction goals, problems to be solved, answers to them, topics to be addressed, explanations, etc., and negotiation may or may not occur with respect to specific *negotia* depending on the social relations concerned (teacher-student, student-student, and so on); (b) the initial state will contain the common goal of interacting in a certain way with a view to some learning taking place; (c) the final state will be agreement with respect to some or all of 1; and (d) negotiation processes may include a critical examination of reasons for and against different points of view.

This concern with negotiation and argumentation is echoed in research on cooperative learning that has been carried out from Vygotskian and from neo-Piagetian perspectives. A number of researchers have recently attempted to go beyond identifying variables of the task and students' prior knowledge to consider variables of the interaction itself. For example, Barbieri and Light (1992) found that verbal measures such as negotiation were significant predictors of the efficacy of pairs, in terms of success on tasks and the "partnering" of success. They concluded that "cooperation," associated with negotiation, rather than "conflict" was favoured as productive for individual learning.

The work which has been conducted within the recent "situated learning" paradigm is also replete with references to the importance of conversation and negotiation in learning within real-life settings (see for example Clancey, 1992; Lave, 1988; Lave & Wenger, 1989; Roschelle & Clancey, 1992; Sandberg & Wielinga, 1990; Seely-Brown, 1990), although the nature of negotiation itself is not made explicit. In this paradigm negotiation is both a means for coordinating perception, action and the environment (thus Seely-Brown, 1990, speaks of experts "conversing with the situation itself") and a mechanism for social construction of knowledge by conversation, within "communities of practice" (Lave & Wenger, 1989).

Finally, although negotiation is an "old" idea in educational research, it is relatively new in Intelligent Tutoring Systems (ITS) research. This is due in part to constraints of formalisability imposed by AI programs, which precluded until recently the implementation of genuinely mixed-initiative interaction mechanisms. Another important restriction of earlier ITS research was that of only considering domains for which a high degree of formalisation and consensus already exists, such as restricted branches of mathematics, physics and computer programming, and "domains" consisting of encyclopaedic collections of facts. Essentially, this suggested that in these carefully chosen sub-domains, there was little "room" or necessity for negotiation with the student—where the "right answers" are known, why negotiate them? The introduction of the notion of negotiation in ITS research is therefore strongly associated with an *extension of the teaching domains considered*, towards social and human sciences, for example. This idea is pursued in a recent collection of papers on the theme of "Knowledge Negotiation" (Moyse & Elsom-Cook, 1992). The main themes presented include provision and negotiation of *multiple viewpoints* on a domain, negotiation of the future course of the tutorial dialogue, consideration of domains where knowledge (or rather "belief") cannot be regarded as certain, uni-dimensional, complete and "closed," and argumentation with respect to justified beliefs as

a means for co-constructing knowledge in tutorial interactions. These themes remain, however, "themes" within what may be termed an "approach" to ITS research. That is to say, the logical or conceptual relationships between them remain implicit, just as the "learning as knowledge transmission" view remained implicit in earlier ITS research.

In summary, although negotiation has emerged as a transversal theme in cognitive science disciplines concerned with teaching and learning, little or no research has provided a systematic analysis of what negotiation is, nor a detailed model for negotiation processes in teaching-learning dialogues. This is necessary if we want to add some substance to the notion of negotiation, and to understand how it relates to learning. In this paper, we therefore aim to present a *theoretical approach for modelling negotiation in learning dialogues*, and to explore the way in which negotiation, thus defined, acts as a mediating mechanism in learning.

Relevance of Dialogue Analysis and Modelling for AI&Ed Research

Analysis and modelling negotiative dialogue is important for Artificial Intelligence and Education (AI&Ed) research in a number of ways. Current interest in studying negotiation and argumentation processes is itself part of an attempt to extend the field of study from one-to-one teaching interactions (between humans, or between individual learners and ITSs) to teaching-learning situations where *groups* of agents interact (human and artificial). The research described here is therefore relevant to the study of many different kinds of teaching-learning situations where verbal or written interaction plays an important part. These include: (a) interaction with artificial co-learners (Chan & Baskin, 1988; Chan et al., 1992; Dillenbourg & Self, 1992), (b) communication in computer-mediated distance learning networks, and (c) interactions between students in cooperative learning situations (see for example Barbieri & Light, 1992; Behrend, Singer, & Roschelle, 1988; Blaye, 1988), as well as (d) dialogues between individual learners and teachers or ITSs.

Given this wide applicability, our own research goal is more general than the above: to produce a cognitive simulation, in the form of an AI program, that will enable us to understand the relationships between negotiative dialogue, problem-solving and learning. The research described here represents one step towards achieving that goal, and consists mainly of an attempt to bring existing approaches to modelling communication closer to real dialogue data. The second step will be to implement the new model in an AI simulation in order to test its completeness and consistency.

Existing AI work on modelling the mental states underlying agents' recognition and generation of speech acts (see, for example, Cohen, Morgan, & Pollack, 1990) has generally relied on invented idealised utterances, and the utterance types that have been most studied are different kinds of assertions, questions and requests. In this paper we shall attempt to show that when we approach the study of real teaching-learning dialogues with these theoretical tools, new types of communicative acts need to be considered and defined, and the existing modelling approaches need to be modified in order to plausibly account for the data. This is one way in which AI&Ed research on dialogue analysis can contribute to AI in general. In addition to analysing dialogue corpora, however, fabricated utterance examples are also used, in order to simplify the presentation of specific theoretical points.

Finally, we would argue that analysis and modelling of teaching-learning dialogues may be of wider interest to AI&Ed research, and particularly to the design of ITSs (alternative "d" above). Despite the fact that there is a reasonable and growing body of AI&Ed research on dialogue modelling (e.g. Baker, 1989; Blandford, 1991; Cawsey, 1989; Collins & Stevens, 1983; Elsom-Cook, 1984; Fox, 1993; Joab, 1990; Petrie-Brown, 1989; Winkels, 1992), and that nearly every ITS or learning environment must have some means for interacting with the student, modelling the interaction itself has rarely been considered as a (or the) primary problem. ITS designers have usually started from modelling teaching domain knowledge and "worked outwards" towards interaction mechanisms, believing that a natural language "front end" could always be added on later. Otherwise, dialogue models in ITS have been viewed as replaceable by graphical interfaces, or as unnecessary within "free discovery environment" and "situated learning" approaches. We argue that, even admitting the alternatives above, it is interesting for AI&Ed research to "work inwards" from analysis and modelling of dialogues in learning situations towards models for knowledge as it is taught and learned. One one hand, a dialogue model can not simply be added on to existing domain and student models: Effective dialogue models establish their own constraints on other components of ITS architectures. In addition, since many of the interactive functions performed in verbal interactions (such as negotiation of turn-taking, focus and time management) have counterparts in graphical ones, it is worth studying the former as a key to the latter.

The paper is structured as follows. In the next section two examples of negotiation occurring in natural dialogues (which are taken to be representative cases) are discussed by way of illustration. Then existing research which has made significant reference to negotiation is critically discussed, in an attempt to identify the current consensus on what negotiation is. In the

third section of the paper a "communicative act" model for negotiation in dialogue is gradually developed, based on detailed analysis of dialogue sequences. In conclusion it is argued that the primary attitude of agents in such dialogues is *acceptance*, rather than *belief*—a claim that has implications for student modelling approaches based on the latter concept. Finally, the question of how learning could take place through negotiation is discussed.

TWO EXAMPLES OF NEGOTIATION IN LEARNING DIALOGUES

We begin by briefly discussing two examples of negotiation sequences in teaching-learning dialogues, in order to concretise what we are talking about.

Example 1 (Figure 1) is taken from the beginning of a set of dialogues conducted via electronic mail, in the domain of learning PROLOG. The teacher was a university lecturer in artificial intelligence and the learner a masters student in the same subject.

```

1/Teacher : I would greatly enjoy doing some prolog. Did you have something in mind?
2/Learner : No, not really. Like I said I haven't really looked at your program with the idea of
    extending it yet. Could we do something simplish, just to make me feel good?
3/Teacher : Certainly. You have a choice of something simple and interactive now, or some-
    thing fairly simple with an interactive session later on today.
4/Learner : Um. I'd like to do something now, cos I really should do the vp stuff later on ... is
    that ok?
5/Teacher : Fine. Ummm... How about writing a predicate called second_word, which takes
    as its argument a string, and instantiates its second argument to be the ATOM corresponding to
    the second word of that string eg
        second_word('fish and chips',M).
        M=and
    yes. How does that seem?
6/Learner : Ok I think. um can you do something else in another window while I think about it?
7/Teacher : Certainly ...
    < ... >
  
```

Figure 1. Negotiation example 1 (PROLOG email dialogues)

Example 2 (Figure 2) is taken from the middle of a one-hour long face-to-face interaction between two students (17 years old), who were given the task of interpreting and explaining results of simple mechanics experiments in physics (the rebound behaviour of balls of different substances) in terms of *energy*. The students were asked to produce a single solution upon which they were agreed. Prior to the extract, the students had allowed two balls of the same size, one made of rubber, the other of plastic, to fall at the same time from the same height. They had agreed that the rubber ball had rebounded 5cm lower than the plastic one. S2 then put forward the explanation that the difference in rebound behaviour was due to the difference in mass of the balls, S1 did not agree, and so they decided to consider what would happen in the simpler case of an inelastic impact.

<...>
 1/ S1 : so, what can we say if there's an inelastic impact ?
 2/ S2 : well, that the energy ... all the energy ...
 3/ S1 : well, that the kinetic energy is theoretically nil !
 4/ S2 : it's nil on arrival, in fact ...
 5/ S1 : since ... since the object stops, in fact, ah yes, especially since there it doesn't move, uh ...
 6/ S2 : it's nil at the start, and it's nil on arrival ...
 ... about energy ... yes, but at the moment of an inelastic impact, what is it that ...
 7/ S1 : we've been doing that for a while now ! <sighs>
 8/ S2 : but we've also ...
 9/ S1 : wait ... inelastic impact, right, you've conservation of momentum, but ... the kinetic energy isn't conserved ! I think that's what we've seen ... the elastic impact, by contrast, both are conserved ...
 10/ S2 : Yes, elastic impact, there's the total energy which is conserved ...
 11/ S1 : Yes
 < ... >

Figure 2. Negotiation example 2 (collaborative physics dialogues)

On the face of it, these two examples appear quite different: In one, there is a teacher and a student engaged in a dialogue via electronic mail, discussing what they should do, and in the other, there are two students (engaged in a face-to-face dialogue) trying to solve a problem together. Why, therefore, should they both be considered cases of "negotiation"? There seem to be two main reasons.

The first reason concerns negotiation initial and final states: In each dialogue the *agents' exchanges are oriented towards the search for agreement* (a specific negotiation *end state*). Each example illustrates different types of negotia ("things that are negotiated"). In example 1, the agreement sought concerns problem-solving *goals* and their specific properties; in the second

example agreement concerns intermediary problem solutions (an explanation for the experimental observations). Negotiations are often associated with resolution of conflicts of some kind between agents (a specific *initial state*), especially in Distributed Artificial Intelligence (DAI) research. This is not, however, the case in these two examples. We claim that it is reasonable to speak of *negotiation in the absence of conflict* (whether explicitly expressed or not), that conflict is a specific case that may arise at any moment in an interaction, and that one of the principal necessary conditions for negotiation is (minimally) the presence in the initial state of the cooperative goal of reaching an agreement of some kind.

The second main reason concerns negotiation *processes*: In each example *both agents contribute to the interaction* in order to *successively refine* the joint proposal towards one that is mutually acceptable, that is, the interactions are relatively *symmetrical* in that each agent uses a variety of utterance types, and each may propose refinements to the joint proposal. In example 1, after an initially proposed goal of solving some problem, both student and teacher successively refine this goal, rendering it more specific and concrete, by proposing a specific level of difficulty, planning timing of problem solving goals, and eventually a specific problem to work on (the student then negotiating "time to think"). In example 2, after an initial (implicitly accepted) proposal to work on the sub-problem of what happens in the case of an inelastic impact, a partial solution is gradually refined by both students. It is therefore the use of the refinement strategy that gives one reason for viewing the two examples as negotiations.

It may appear that these characteristics (search for agreement, interaction symmetry and refinement strategy) are sufficiently general so as to be true for *any* dialogue, whether concerned with learning or not. This is not, however, the case. Consider, for example, the "Socratic" dialogues of Collins and Stevens (1983), where interaction roles are quite different for the teacher (asks questions) and the student (answers questions). What makes the two examples above cases of negotiations is therefore fundamentally the fact that the interlocutors *collaborate* as relative equals in the pursuit of the *co-operative goal*¹ of reaching agreement, and that they use specific strategies to do so, such as joint refinement of a ongoing proposals.

Finally, if negotiation is concerned with reaching an agreement, then the question is raised as to what exactly "agreement" is in this context. The first response which would spring to the mind of anyone familiar with computational speech-act theory would be: "having mutual beliefs." Later in the paper we shall argue that the notion of mutual belief is (alone) inadequate for modelling the phenomena occurring in such dialogues. We should rather speak of something like "willingness to *accept* the current proposal for the

present purposes, as far as it goes." In that case, speakers are said to be *committed* to the propositions, unless they have explicitly denied them (or retracted them, if they themselves uttered them). To put it another way: Agreement is about playing the dialogue game of explicitly accepting; belief is something else.

Now that the basic phenomena to be modelled have been described, we discuss existing work on negotiation.

MULTIDISCIPLINARY PERSPECTIVES ON NEGOTIATION

After discussing an attempt to develop a paradigm for AI&ED research based on negotiation, we review work in Distributed Artificial Intelligence and in Language Sciences on negotiation in other contexts.

Knowledge Negotiation in Artificial Intelligence and Education

A new approach to AI&Ed research has been proposed recently—that of "Knowledge Negotiation" (abbreviated to "KN") (Moyses & Elsom-Cook, 1992). The approach has been defined in opposition to what may be termed the "knowledge-based tutoring" (Clancey, 1987) approach, as follows :

There has recently been a move in research on Intelligent Tutoring Systems (ITS) away from the idea that the tutor takes a directive role controlling the interaction. This view implied that the tutor had the *correct* knowledge about the subject area and was in a position of authority over the student. It has been realised that for many domains there is not a single correct representation, and that the interpretation of the domain, or "viewpoint," must be jointly constructed between teacher and learner. This implies the need for representations and mechanisms to support negotiation. (Moyse & Elsom-Cook, 1992, p. 1)

As Self (1992, p.39) points out (in the same volume), there are two main motivations for appealing to negotiation in ITS research: (a) to enable discussions about how to proceed, what strategy to follow, what example to look at, etc., in an attempt to loosen the tutorial control of standard ITSs, and (b) to permit discussions about the viewpoints (i.e., the beliefs or the knowledge) themselves, on the philosophical grounds that one agent (the tutor) should not be deemed to have infallible insight into the "correct" viewpoint. The KN approach can therefore be summarised as the following

set of claims:

1. **(Incomplete knowledge)** Some (or all) tutors do not possess complete knowledge for the domains which they teach.
2. **(Viewpoints on knowledge)** For some domains there is no single, privileged or correct viewpoint on knowledge.
3. **(Argumentation and belief)** There exist domains, or parts of them, for which there is no "knowledge," but only a number of different competing and arguable sets of justifiable beliefs.
4. **(Multiple solutions)** For some problems in some domains there are multiple correct or acceptable solutions.
-
-
5. **(Non-transmissibility of knowledge)** For domains that possess some or all of characteristics 1 to 4, teachers should not aim to simply transmit their own knowledge (beliefs, viewpoints, solutions, etc.) to the student.
6. **(Control of interaction)** The tutor should not have complete control over the tutor-learner interaction in domains that possess some or all of characteristics 1 to 4 above.
7. **(Knowledge co-construction)** The representation (viewpoint) of the domain to be taught/learned, and problem solutions themselves, should be jointly constructed by teacher and learner.
8. **(Negotiation mechanisms)** Negotiation mechanisms and representations (viewpoints) to support them are required in teaching/learning interactions with respect to domains that possess some or all of the properties 1 to 4.

Claims 1 to 4 correspond to Self's (op.cit.) second motivation and claims 5 to 8 to his first. It is important to note that 1 to 4 are **epistemological** claims—they are about the nature of knowledge in different domains and for different individuals such as teachers and learners—and that 5 to 8 are claims about teaching and learning **interactions**—how they may be controlled and by what processes. In addition, the KN approach, as described in Moyse and Elsom-Cook (1992), claims that the epistemological claims somehow *imply* or *require* the claims about the interaction. In order to clarify the nature of this implication we now need to briefly examine each individual claim.

Claim 1 (incomplete knowledge) is either a statement about limits of human individuals (no teacher has limitless knowledge) or a claim that there exist domains where knowledge (on the part of anyone) is incomplete (which is surely true). With respect to education, the claim becomes more interesting

when we consider the more likely incompleteness of individual knowledge with respect to *how to teach* or *knowledge of individual students' knowledge*. **Claim 2** says that there may be different conceptual viewpoints on a knowledge domain. The classic example is that of physical, geographical, biological or other systems, where one may adopt a functional or a causal viewpoint, for example. There are two main approaches to representing what viewpoints are: the "belief systems" approach, and the "procedural" approach. According to the first, a viewpoint is simply a set of beliefs, sometimes called a "context" where there may be overlapping subsets of a wider set (Self, 1992). The procedural approach goes further in including the reasoning mechanisms which may be applied to a viewpoint, and a set of contexts or problems for which the viewpoint is appropriate, within the viewpoint itself (Moyse, 1992). **Claim 3** (argumentation and belief) transforms viewpoints, and the parts of them within which the system or student is competent, into beliefs rather than knowledge. This implies a special kind of teaching dialogue, since an interactive exchange of justified beliefs is a minimal definition of an *argumentation*. For example, Goodyear and Stone (1992) claim that "Given the epistemological disagreement that characterizes the social sciences... we would expect a domain expert both to be able to present particular arguments in social science domains fairly... and to be able to step outside such a framework in order to offer critiques comparable to those offered by other social scientists" (p. 75). Many of these issues are similarly problematic in the physical sciences, although "...for many purposes, it is possible... to act as if those problematic issues did not exist... without the raising of epistemological issues" (p. 72). **Claim 4** (multiple solutions) is closely related to claim 3, since there are domains where multiple and competing theoretical frameworks exist, but is also true for some parts of more "conventional" domains such as branches of mathematics and physics, where there may be multiple solution *methods* for a given problem solution, as well as multiple solutions themselves.

To sum up, claims 1-4 amount to an appeal to consider a wider range of teaching domains (other than math, physics and computer science), together with a general *weakening* of the knowledge that teachers may possess, and an *epistemological enriching* of domain knowledge itself. It is this latter point which establishes a link with the further claims 5 to 8.

Let us now consider claims 5 to 8, and the extent to which they follow from claims 1 to 4. **Claim 5** (non-transmissibility of knowledge) is effectively a rejection of the "knowledge transmission" view, described in the introduction to this paper. It follows from 1-4 to the extent that given incomplete knowledge, justified belief, multiple viewpoints and solutions, it is a *misrep-*

resentation of what the domain knowledge is in this case, to attempt to transmit or impose a *single* viewpoint, set of beliefs or solution, or to accord a greater degree of certainty to knowledge than it really possesses. **Claim 6** (control of interaction) follows from 1-4 given a certain connection between knowledge or competence with respect to a task, and control of its execution (as demonstrated by distributed AI research, for example), that is, the agent who controls task execution is the agent who has the relevant task-related knowledge. If, as 1-4 suggest, a teacher's knowledge may be weakened under certain circumstances, then the teacher's control over the teaching interaction must also be weakened. This leads to the view that control may be shared between student and teacher. **Claim 7** (knowledge co-construction) is *one* possible manifestation of the sharing of control advocated by claim 6, although it is not the only such manifestation (for example, the student may agree to pursue one of the possible viewpoints presented by the teacher). It does not therefore strictly follow from 1-4, although it bears a "family resemblance" to 5 and 6. Finally, **claim 8** (negotiation mechanisms) follows from claim 6 (sharing of control) to the extent that if control of the interaction is shared, some mechanism, such as negotiation, is required for coordination, reaching agreement, and possibly resolving conflicts. Another possible mechanism is *critical argument* (Baker, 1989a, 1990, 1992b; Goodyear & Stone, 1992), where positive and negative supports for alternative views are jointly explored within the perspective of reaching agreement (which makes it a special form of negotiation).

In summary, the logical relation between the knowledge-negotiation claims, as represented here, is as follows: $\{1-4\} > 5; 5 > 6; 6 > 7; 7 > 8$. The main links between the epistemological claims and those concerning the interaction are relate to **epistemological fidelity** and an **adequation between knowledge and control**. Our aim here has been to attempt to render the KN view more explicit and to highlight relations between elements of the approach. It must be emphasised that the KN approach does not imply that *all* teaching interactions must be negotiative ones, since there may be domains or parts of them for which teachers have perfect competence and knowledge of students' knowledge (although the latter is unlikely). In this case a directive teaching approach may be appropriate. For many domains, however, recognition of the limits of a teacher's knowledge along different dimensions should lead to some form of negotiative teaching interaction.

Given that under certain circumstances negotiation mechanisms are necessary in tutorial interactions, it is surprising that there has been relatively little work within the KN approach on specifying precisely what such negotiation mechanisms are. The KANT ("Critical Argument Negotiated Tutor-

ing") system (Baker, 1989b) incorporated mechanisms for negotiating dialogue focus and interaction style in a computational model for tutorial dialogue. A number of elements of the model were validated with respect to analyses of computer-mediated teaching dialogues (Baker, 1992). For example, it was found that: a large range of different things may be negotiated (topic, problem to be pursued, level of difficulty, interaction style, ending the dialogue, etc.), explicit negotiation is more prevalent at the beginning of dialogues; negotiation generally proceeds from general to specific, with increasing concreteness (one of the principal hypotheses on which KANT was based); and, implicit negotiation generally occurs in situations where conflict or failure is thought to be likely, in an attempt to avoid it. Rimmershaw (1992) has also analysed teaching dialogues (verbal face-to-face) in an attempt to identify how *explanations* are negotiated. Her major findings were that in order for the form and content of explanations to be negotiated, a relatively "equal" and free interaction is required in terms of "power relations" between explanation participants. On this point she cites Edwards and Mercer's (1987) work on educational interactions in schools:

... negotiation is a rather one-sided affair in which the teacher's role as authoritative bearer of the ready-made knowledge simply finds alternative, more subtle means of realising itself than the crudities of brute "transmission." (p. 163)

We now turn to research on negotiation in other domains in an attempt to shed light on the nature of negotiation processes themselves.

Negotiation in DAI, Agent Theory, and Language Sciences

For many AI&Ed researchers, the most obvious place to look for work on negotiation is in DAI research (see, e.g., Bond & Gasser, 1988; Gasser & Huhns, 1989). Multi-agent systems are concerned with coordinating intelligent behaviour among modules that are considered as relatively autonomous "intelligent agents." In order to do this, agents must be able to reason about *processes of coordination* with others. This has led to the specification of "negotiation protocols," for conflict resolution and information exchange, the most well known of which is the "contract net" of Davis and Smith (1983). But what precisely is a "protocol" in DAI research? Rosenschein (1993, p. 794) describes a protocol as "the rules by which agents will come to agreements," where the rules specify "the kinds of deals they can

make, as well as the sequence of offers and counter-offers that are allowed." For our purposes here, the main results of DAI research on negotiation mechanisms may be summarised as follows:

1. Negotiation is viewed as a mechanism for resolving conflicts concerning problem-solving resource allocation (with the exception of Galliers, 1989, who views negotiation as a process for achieving cooperation from initial "postures" of conflict or of "indifference").
2. *There are different conflict resolution styles* (Adler et al., 1988) that vary according to the degree of goal expansion prior to negotiation, the complexity of the information exchanged between agents, the scope of the model each builds of the other, and the degree of control agent has over the negotiation process itself. Negotiation will be a suitable choice when agents have *wide latitude* in problem solving and where they are able and willing to refine their plans to accommodate effects of others' plans.
3. *Argumentation is a possible strategy* (Sycara, 1988, 1989) when an offer is not accepted, in which case one agent may try to convince the other to accept by attempting to change the perceived benefits of the offer.
4. The *basic communication elements* of negotiations are announcements of offers and bids, leading to contract formation.
5. Communication between agents via a protocol deals with the content of messages exchanged, rather than with the mechanisms of communication itself (Rosenschein op. cit.).

Since DAI research concentrates on characterising the type and contents of messages to be passed between artificial agents, and on describing the mental states of agents who exchange such messages, it does not, therefore, give a complete picture of the structure and management of negotiative interactions between humans in natural language. For this we need to look to research in certain branches of language sciences, particularly discourse and conversational analysis. For the purposes of modelling educational interactions involving human agents, perhaps the most important point emerging from DAI research is point 2 above, which relates to the Knowledge Negotiation approach described earlier: For negotiation to be an appropriate strategy, there needs to be wide latitude in the domain (such as justifiable and arguable beliefs), and agents need to be able to relax their constraints.

Work on negotiation from the point of view of language sciences concentrates on precisely those elements that are absent in DAI communication protocols: *negotiation of meaning*, and *negotiation of interaction control*. In other terms, whilst DAI research concentrates on the task-level, linguistics

ties research concentrates on negotiation at the communication-level. Communication-level negotia (things that may be negotiated) may be quite diverse. Bange (1992) thus describes conversational negotiation as follows:

"To negotiate" is to attain agreement on the way in which an activity will be accomplished and on its signification. This agreement is realised across a sequence of stages where the partners intervene: (1) one to propose, (2) a second to accept or refuse the proposition, (3) in the case of acceptance, the first to ratify it. This negotiation is in general automatic and tacit, but it may be thematised. Any activity which must be coordinated must be in this sense negotiated. (Bange, 1992, p. 29)

According to Moeschler (1985), these "activities" to be coordinated include negotiation of facts necessary for the continuation of the interaction, the images of themselves that protagonists want to impose (such as a teacher being "knowledgeable" or "authoritative"), rights and obligations, and the construction of common interpretations for utterances (metadiscursive), as well as interaction management.

It is important to note that from this perspective negotiation is not merely a possible or optional phenomena in verbal interaction, it is in fact *constitutive* of it. As Moeschler (1985, p. 176) puts it, "Without negotiation the dialogue is transformed into monologue, the function of the interlocutor being reduced to that of a simple receptor of the message." In educational interactions, absence of negotiation at the task and communication levels clearly corresponds to the "transmission" view mentioned earlier. One way in which negotiation is constitutive of dialogue relates to what Edmondson (1981) describes as the "strategic indeterminacy" of illocutions. The general idea is that the illocutionary force of an utterance as, for example, an assertion or a request, is not solely determined by the speaker of the utterance and the relevant context at the time of utterance, but may also be determined (i.e. negotiated) retroactively by the hearer's response to it. Thus if X says "It's cold in here" (u1) when entering Y's apartment, the illocutionary force of u1 as an indirect request to put the heating on, or a simple statement, etc. may be negotiated by X as a function of Y's response (e.g. Y u2: "Shall I put the heating on?" / X u3: "No no, I was only remarking"). This means that communicative acts do not have immediate and automatic effects. These effects are *assumptions*, that may or may not be confirmed by further utterances (Perault, 1990).

From a cognitive perspective, Clark and Schaefer (1989) have explained negotiation of meaning of utterances in terms of the following *grounding criterion*: "... The contributor and the partners mutually believe that the part-

ners have understood what the contributor meant to a criterion sufficient for current purposes" (p. 262). The criterion is presented in critical opposition to most "computational speech act" models which, it is claimed, assume that the "common ground" in dialogue simply "accumulates" as the result of uttering the right sentence at the right time. Thus specific interaction structures occur—such as "episodes," "repairs," "collaborative completions," etc.—whose function is to assure grounding in these terms, by giving varying degrees of "feedback" (Allwood, Nivre, & Ahlsén, 1991) on acceptance of proposed propositions.

From this brief review of language sciences research on negotiation it emerges that negotiation may be viewed not only as a mechanism by which cooperative tasks external to the interaction itself may be achieved, but also as the fundamental mechanism by which the interaction itself is maintained on track.

Summary: Dimensions of Negotiation

From this review of research on negotiation in different domains, a large degree of consensus emerges on what negotiation is, what may be negotiated, what negotiation processes are, and the important role that negotiation plays in learning and problem-solving.

Negotia (what is negotiated) are defined on two levels: the *domain-task* level, and the *communicative* level. At the domain-task level, negotia may be (justified) beliefs, goals, problem solutions, sharing of problem-solving resources and viewpoints on knowledge or belief. The task external to the dialogue, or the knowledge domain concerned, must possess a certain level of "latitude" in order for negotiation to be an appropriate strategy for achieving joint goals. At the communicative level, negotiation may concern agreement on meaning of utterances (in a wide sense), images that agents want to impose (for example, the teacher wanting to impose the image of knowing everything that is relevant), rights (who can say what kinds of things, when and how), and aspects of interaction management.

The *initial state* of negotiation requires at least that the parties have the common goal of interacting in order to attain an agreement with respect to some negotia. There may or may not be conflicting beliefs or goals (resolution of which may be one of the reasons for interacting), or individual goals concerning what precisely the final agreed negotia should be. Negotiators need to be *relative equals* (or at least to believe themselves to be) in terms of rights to make proposals. The *final state* is cooperation or agreement with




Figure 3. Model for negotiation in dialogue: general schema

Negotiation: Initial and Final States

The initial state is defined in terms of specific elements (individual and mutual goals and beliefs) of the mental states of dialogue agents. The one necessary element is the mutual goal of agreeing on some set of negotia: (mgoal A1 A2 (agree A1 A2 {negotia})). The final state will obtain when the proposition "(agree A1 A2 {negotia})" is present in the contexts of each agent (or the dialogue terminates due to time constraints, because it is believed that this goal is not achievable, etc.). There may be special relations between individual attitudes ("Ra"), such as mutually believed contradiction between beliefs that each attributes to the other (of various "one-sided" or other forms),³ or the mutual belief that goals of each can not both be fully achieved, but these are not essential to the definition of negotiation. Each agent does not, however, want to agree on *any* set of propositions whatsoever: There will be *constraints* on the final agreement from the point of view of each. These may relate to individual goals, such as "agreeing on all and

respect to some negotia.

Negotiation processes involve making offers, bids or proposals by one agent, and the acceptance or rejection of them by another. In the case of acceptance, the acceptance must be ratified by the other agent. On the *strategic* level, there are two main possibilities: attempt to refine goals by offers and counter-offers, towards agreement, or remain with a given offer, and attempt to persuade the other to accept it on initial rejection. This persuasion may involve, according to the negotiation situation, threats, argumentation, or attempts to change the other's perception of the desirability of the offer.

These features of negotiations will be taken into account by the modelling approach described in the next section.

A MODEL FOR NEGOTIATION IN DIALOGUE

After having motivated the need for a model of negotiation in teaching-learning situations, and illustrated the phenomena to be modelled, we now present the main components of such a model, and describe one negotiation process in detail. The model is illustrated mainly for the case of collaborative learning dialogues (see Figure 2) since these involve the most interesting phenomena at the domain level.²

Components of the Model

Figure 3 below shows the main elements of the model: initial state, final state, negotia, and negotiation processes. Negotia are shown as shaded and empty circles to indicate the fact that task and communicational levels of negotiation proceed simultaneously. Throughout the paper we will write attitudes of agents as "(*<attitude>* *<agents>* *<propositions>*)", and communicative acts, or "CAs," to be defined later, as "CA_n[*<agent>* *<propositions>*]." Different agents or speakers will be written "A1,A2, etc." and "S1,S2, etc." For convenience on some occasions, we shall also refer to different agents and speakers as the current "speaker" or "hearer." Different propositions will be named p1, p2, etc. Other notation will be introduced as necessary.

only the set of propositions that I propose"; they may be logical, such as a constraint that the final set of agreed propositions should not be (internally) contradictory; or they may be social in nature—for example, the solution must "satisfy the teacher." There may also be domain-dependent constraints, such as in negotiation example 1 above, where the students' task was not simply to produce *any* kind of explanation for physical phenomena, but rather an explanation stated in terms of a model for *energy* in physics. We represent the goal state as a set of unstructured propositions. This is of course a simplification, since the propositions will be related in specific ways (elements of a problem solution, an explanation, goals and subgoals, etc.).

We can now state the following general defining criteria for negotiative dialogue, together with definitions for specific types of negotiations.

- (Def 1) **Negotiative dialogue initial state (general).** The interacting agents have the mutual goal that agreement will be reached with respect to some set of propositions in the dialogue, subject to constraints on the set of propositions to be agreed from the point of view of each agent.
- (Def 2) **Competitive negotiative dialogue (specific).** One or each interacting agent wants all and only the propositions which they propose to be in the end state.
- (Def 3) **Disinterested negotiative dialogue (specific).** Neither of the interacting agents has the specific goal that all and only the propositions which they have proposed to be in the end state.

These definitions apply both to complete (negotiative) dialogues as well as to *sequences* within dialogues. Given a goal decomposition structure at the domain as well as at the dialogue level, agents may negotiate some subgoals but not others (they may be assumed to be agreed, or there may be conflict and ensuing argumentation) within a negotiative dialogue, and in dialogues which are not negotiative (there is no high level goal of reaching agreement) negotiation may occur with respect to certain tasks.

If the initial goal of a negotiative dialogue is to reach agreement, then we may ask how *this* goal is established? There seem to be three main possibilities: the goal is *assumed*, *imposed*, or *negotiated*. In the first case, dialogue goals may be assumed in certain concrete situations: For example, in a travel agency it may be assumed that the goal is to agree on a specific price for a specific voyage. Negotiative goals may be *imposed* in very special circumstances, such as in experiments (negotiation example 2) where students may be "given" (i.e., imposed) the task of producing a single agreed solution to the problem. Clearly, some tasks lend themselves more to negotiation and collaboration than others. Finally, negotiation may be "pushed back" (theoretically indefinitely) to the goals of the dialogue itself, that is, agents may

negotiate that they should negotiate. This is rarely done *explicitly*; more often general dialogue goals are negotiated *implicitly*. We define implicit negotiation as follows:

An agent A1 implicitly negotiates acceptance of *f* by agent A2 in performing some communicative act C[A1p] if *f* is not already agreed by A1 and A2, and *p* is mutually understood by A1 and A2 to be relevant to *f*.

This simply means that one of the interacting agents performs some action(s) which may be interpreted by the other as relating to achievement of the (implicitly negotiated) goals in question. Thus, for example, *p* may be mutually understood to be a possible way of achieving a not-yet-agreed goal, and by offering *p* the offering agent is implicitly negotiating acceptance of that goal. Or, in making some utterance concerning "velocity," given that the current focus is "energy," an agent may be understood to be implicitly negotiating a change of focus, since this is one of the direct contextual effects of that utterance. Or again, saying "the kettle has boiled" may, in certain circumstances, be an implicit offer to make tea, a request to make tea, etc. "Relevance" must thus be understood in several senses here (inferential relations in communication, subgoal, semantic, thematic, etc., relations). In some dialogues, there may be repeated implicit negotiation of the goals of negotiative dialogue (for example a dialogue in which a mature student attempts to "put the dialogue on a more equal footing" with a teacher, by attempting to negotiate, but where the teacher steadfastly attempts to retain authority).

Negotia

We have so far described negotia as propositions (or sets of them). However, there are a number of different kinds of things that may be negotiated (as described in the literature review above). The main types are classified in Figure 4 below.

Within the task-level class, agents may negotiate attitudes (what is to be believed, what counts as knowledge, what goals to pursue), or negotiation may take place at the level of performing a task (how it should be conceptualised in a viewpoint, resolution methods, solutions themselves). With respect to attitudes, it seems to be the case that *negotiation of goals negotiation is more fundamental than its epistemic counterpart*, that is, before you can negotiate a belief or a problem-solution you must first be agreed on the goal of doing so (see negotiation example 1).




Figure 4. Classes of negotia (things that may be negotiated)

At the communication level, features "internal" to utterances may be negotiated (perception, or "what was said," meaning of utterances), "external" communicative features may be negotiated (rights in the interaction, images or models of the other), and dialogue control (Buit, 1989) aspects may be negotiated (structuring sequences, topic shifts, stopping or starting the interaction; management of turn-taking, and so on). With respect to "models of the other," this may be important in teaching interactions since students may want to negotiate between their own metacognitive awareness of their competence, goals, and so forth, and the "student model" that the teacher appears to be working with.

In order to simplify the ensuing discussion, negotia are represented simply as a set of propositions. Extensions to the representation of the dialogue context required to take into account communication level negotiation are discussed later.

Negotiation Processes

Negotiation processes have the purpose of achieving the final state from the initial state. The inverted triangle in Figure 3 is meant to indicate *convergence* towards this final state (see Vernant, 1992). There are three main components of negotiation processes: (a) strategies, (b) actions, and (c) rights.

Strategies are general methods adopted by agents to achieve the negotiation goals, within the constraints of each agent. There are three main strategies, as shown in Figure 3 (see Druckman, 1973): (a) refine, (b) argue, and (c) stand pat.

Within the **refine** strategy, an agent responding to an offer of another decides to make another offer which "refines," "builds on," or modifies the original offer in some way (for example, generalising it, adding a new element to the "package," etc.). This strategy is used when the original offer does not violate constraints of the receiving agent's goals and beliefs, and perceived "distance" between the original offer and a proposal that the agent could eventually accept is not too great.

The **argue** strategy is used when one agent's offer is not accepted by another, and the first agent has a strong commitment to what is offered (goal, belief, solution, etc.). In this case, the agent will argue for the offer in an attempt to convince the other to accept it. The **stand pat** strategy is used when an agent has no reason to reject the offer of another, and no other relevant proposal to make. It does not simply mean "doing nothing": often the agent will encourage the other to elaborate their proposal further. Thus in the corpus from which example 2 is drawn, students sometimes spontaneously adopt "elicitor" roles, encouraging the other to develop a proposal.

In this paper we **only consider the refine strategy**. This is because it is prototypical of negotiations (and is already complex), because the "stand pat" strategy is relatively simple to define, and because the "argue" strategy involves complex issues which go beyond the scope of the paper.⁴ Figure 5 summarises the refine strategy, explicitly representing the associated dialogue actions and relations established between successive offers (R1, R2, etc.). A specific set of such relations is presented later in the paper.

Actions realise particular strategies for achieving negotiation goals. Here they are restricted to *communicative acts*. For the "refine" strategy, communicative acts are shown in Figure 5, including initiating moves such as *offers* (of candidate propositions for agreement) and reactive moves or speech acts which communicate *acceptance* or *rejection*. An acceptance which follows an acceptance by the other is called a "ratification."




Figure 5. Model for negotiation in dialogue: schema for refine strategy

Rights are here restricted to rights to perform *actions*, which are themselves related to strategies. So far our definition of negotiation is quite compatible with the case where one agent proposes all the propositions, and the other simply accepts them, or may accept or reject them. Intuitively, this is not a prototypical case of negotiation (1). A stronger condition for negotiation is therefore required, with respect to *symmetrical dialogue rights* to propose, accept, or reject propositions, as follows:

(Def 4) **Symmetrical rights** in negotiative dialogue. Each interacting agent has the right to propose, accept or reject propositions as candidates for the negotiation end state (agreement).

Note that this does *not* say that these rights will be always exercised fully, that each agent will contribute to the same extent or that negotiation requires participants with equal competence and knowledge. Considering our two examples, it is clear that in example 1 there is unequal knowledge between student and teacher; however, the teacher proposes “doing prolog,”

“a simple and interactive session now, etc.,” “writing a predicate...,” and the learner proposes “something simplish,” “do something now,” “you do something else while I think about it.” It is clear that the *kinds* of things which each can and does propose are therefore different. In example 2 the knowledge (or lack of it) of each speaker is approximately at the same “level,” and each contributes in more or less the same way.

We now describe one negotiation process in detail—the **refine strategy**. This is the most used strategy, and most of its component actions are also part of the other strategies. First, communicative acts relating to this negotiation strategy are defined, and the minimal dialogue context within which these acts operate is described. Then we discuss contextual effects of acts when they occur in refinement sequences, and extensions to acts and contexts to incorporate other types of negotia.

COMMUNICATIVE ACTS IN NEGOTIATIONS

A minimal set of communicative acts is described for the “refine” negotiation strategy, preceded by a brief description of the theoretical framework adopted.

Theoretical Framework: Communicative Acts as Functions

We adopt a specific speech act theoretic approach called “communicative act” (CA)⁵ theory. This approach has its problems,⁶ but we believe that it can be extended to meet objections, and that it can provide comprehensive coverage of most communicational phenomena. According to CA theory (Gazdar, 1981; Bunt, 1989), communicative acts are *functions from contexts into contexts*, where a context is the union of epistemic states of dialogue agents—for example, “...an assertion that ‘Q’ is a function that changes a context in which the speaker is not committed to justifiable true belief in $\emptyset$ into a context in which he is so committed” (Gazdar, 1981, pp. 68-69). This circumvents a number of well-known problems in speech-act theory, such as the lack of a simple one-to-one relation between sentence form and illocutionary force (assertions do not always have declarative, nor questions always interrogative sentence forms). Communicative acts are therefore to be identified with their appropriateness conditions (Searle, 1969) and their context-changing effects. Thus what counts as a “question,” “inform,” “request,” etc., does not depend directly on the linguistic form of utterances,

but is any communicative means⁷ which evaluates the corresponding function. Communicative acts are therefore conventional, "packages" of epistemic states which may be communicated in dialogue; communicative functions ("CFs") are names for the immediate or direct⁸ effects that occur when these packages are communicated. Single utterances may realise different communicative functions in different contexts, or multiple simultaneous functions in the same context.

Following Bunt (1989), the effects of a communicative act CA can be represented as follows (A and B are different interacting agents, K is a context of epistemic/doxastic states):

$$K = \{K_A, K_B\} \Rightarrow CA \Rightarrow K' = \{K'_A, K'_B\}$$

where K' is the "updated" context. Note that the CA will have effects on both the hearer *and* the speaker. In the simplest case (which is the case considered here), for every effect E on the hearer, there will be a corresponding effect on the speaker "A" of "A assumes that E." Effects of CAs are thus *de-faults* (Perrault, 1990), to be confirmed or not by subsequent dialogue utterances. In more complete terms, the direct effects of a CA can be represented as follows (Bunt, 1989):

$$\begin{aligned} \{K'_A, K'_B\} &= CA \Rightarrow \\ &\{\text{update}(K'_A, (\text{BMB}(A, B, (A \text{ assumes that } B \text{ believes that } C_j(c))))), \\ &\{\text{update}(K'_B, \{[B \text{ believes that } C_j(c)] \cup [\text{BMB}(B, A, \\ &\quad (A \text{ assumes that } B \text{ believes that } C_j(c))]\})\} \end{aligned}$$

In the formula above, effects E are written " $C_j(c)$ ", which corresponds to the *appropriateness conditions* of the communicative act CA, that is, the elements that must be present in the context of the utterer in order for CA to be an appropriate means of satisfying communicative goals. "c" is the semantic content (not necessarily propositional) of the act. In general, therefore, a CA will have the effects that: a mutual belief ("BMB") is created to the effect that the speaker assumes that the hearer believes the appropriateness conditions, and that the hearer believes those appropriateness conditions.

There are two main classes of CFs: "task-oriented" CFs, that have the principal purpose of achieving the extra-communicational goals of the dialogue (such as information transfer, problem-solving, and so on) and "dialogue control" CFs, that have the purpose of monitoring the dialogue itself, making sure it stays "on track" (such as establishing mutual comprehension

of utterances, opening and closing the dialogue, managing topic shifts, turn-taking and timing, etc). Readers are referred to Bunt (1989) for more details. These different classes of communicative functions relate closely to the types of things that may be *negotiated* (see Figure 4).

Definition of Communicative Acts for Negotiative Dialogues

The basic communicative acts in negotiative dialogues are "offers," "acceptances," and "rejections." Offers mostly realise *task* level CFs, when their purpose is to propose possible solutions to the extra-communicational task, but communication-level offers are also possible. Acceptance/rejection are *dialogue control* CFs, since their purpose is to provide feedback on the existence (or not) of agreement between speakers. These CAs are quite unlike those usually considered by speech act theory—typically, assertions, questions, and requests. Such "elementary illocutionary acts" have the form " $F(p)$," consisting of an illocutionary force F with a propositional content p (Vanderveken, 1990, p. 13). Vanderveken (op. cit.) describes acts such as offers as "conditional," having the form " $p \rightarrow F(q)$." Their "aim is to perform an illocutionary act F(q) not categorically but on the condition that a proposition p is true" (p. 13). We argue that their specificity also resides in the fact that they are directed towards the attitude of "acceptance" rather than "belief." For example, consider the following two utterances from dialogue extract 1:

2/Learner: <...> Could we do something simpler, just to make me feel good?

3/Teacher: Certainly. <...>

In 2/ the learner is clearly not asserting anything, unless we wish to treat this as an "inform of a want," so belief does not seem to be the primary underlying attitude here. The utterance of the teacher in 3/ does not, similarly, express (primarily) belief in anything, but rather *acceptance* of the learner's request or offer. We claim that the same kind of analysis applies to dialogue extract 2, as in the following two exchanges:

9/ S1: wait...inelastic impact, right, you've conservation of momentum, but...the kinetic energy isn't conserved ! I think that's what we've seen...the elastic impact, by contrast, both are conserved...

10/ S2: Yes, elastic impact, there's the total energy which is conserved...

In 9/, S1 may be said to have expressed the "content" "momentum is conserved but not kinetic energy with an inelastic impact, whereas both are conserved in an elastic impact." However, S1 is not primarily expressing *belief* in this "content"; the content is rather *offered* as a possible partial solution to joint problem-solving goals, which does not necessarily imply belief (one could offer something that appears to solve the problem and which one thinks the other may accept without necessarily believing it). Similarly, the "yes" of 10/ communicates *acceptance* of the offer rather than belief in its content. We adopt Cohen's (1992) definition of "acceptance" where acceptance of a proposition *p* implies "willingness to take *p* into account in one's reasoning".⁹ For the present we are concerned with the basic structure of CAs themselves; acceptance will therefore be discussed later in the paper. Acceptance, then, is the fundamental attitude underlying both offers and, obviously, acceptances/non-acceptances.

Edmondson (1981) defines a "propose" act,¹⁰ which we term an "offer," as follows:

Proposes can be interpreted as conditional Resolves i.e. "I will if you will" or as combining in one act the force of a Request and that of a Willing i.e. "You act in my interests and I'll act in yours." (Edmondson, 1981, p. 142)

Similarly, Airenú et al. (1989) define a propose act as one agent committing himself to doing a given act if another will also commit himself. By combining a slight modification of existing work on formalisation of requests (Cohen & Perrault, 1979) with a direct interpretation of "I will if you will," we can arrive at the following characterisation of appropriateness conditions for offers ("bel" is the modal belief operator):

Appropriateness conditions for offers

- C1. (want S1 (bel S1 (accept S2 p)))
- C2. (want S1 (accept S1 p))
- C3. (bel S1 [(accept S2 p) => (accept S1 p)])

This says that the appropriateness conditions *C* for an offer are that: the speaker S1 wants to acquire a belief that the hearer S2 accepts *p* (which can be achieved by the hearer's accepting), the speaker wants to accept *p*, and the speaker believes that if the hearer accepts *p* then the speaker will do so. A stronger form of C1 would be where the speaker has a *persistent* goal of attaining a belief that the speaker accepts a proposition, which would normally be associated with *justified belief* in it, possibly leading to argumentation in order to convince the other.

For the present we shall adopt a very simple characterisation of acceptance and non-acceptance, which is that their appropriateness conditions correspond to acceptance or not:

- Appropriateness conditions for acceptances
C1. (accept S1 p))
- Appropriateness conditions non-acceptances (rejections)
C1. ¬ (accept S1 p))

A stronger and special condition for non-acceptance would be that one speaker S1 believes that if (s)he accepts then another, S2, will not do so (adversarial negotiation where one speaker refuses all propositions offered by another):

- C2. (bel S1 [(accept S1 p) => ¬ (accept S2 p)])

We now consider how these basic CAs operate to effect changes in the contexts of agents in order to achieve the final state of agreement. The definition of "agreement" must be consistent with these CA definitions and with the dialogue data.

Contextual Effects of Communicative Acts in Basic Sequences

Initially, we consider the simple case where an offer is followed by an acceptance. This rarely occurs in real dialogues, but consideration of this hypothetical case is useful for showing how the model needs to be extended in order to take into account more elaborate sequences. We consider the following example (fabricated):

No.	Speaker	Utterances	Communicative acts
[1]	S1 :	Why don't we discuss fractions?	OFFER[S1'discuss fractions']
[2]	S2 :	Ok, let's.	ACCEPT[S2'discussfractions']

We now analyse evolution of the contexts of S1 and S2 prior to and following the communicative acts, using the basic update formula of Bunt, described above, and the CA definitions in terms of appropriateness conditions. In the following, "p" represents the proposition "S1 and S2 will discuss fractions," "bel" is the belief operator, "Asn" is the assumption operator and "BMB" is a mutual belief operator (its precise definition does not matter for our purposes here). Each contextual element is numbered for refer-

ence. We assume for the sake of simplicity that all communicative acts are “successful” and restrict the analysis to the appropriateness condition C3 (“(bel S1 [(accept S2 p) => (accept S1 p)])” of offers. This is reasonable, since C1 and C2, above, are the conditions that make the speaker decide to make the offer in the first place. It is C3 that functions with respect to obtaining agreement.

CA	Context K _{S1}	Context K _{S2}
	Appropriateness condition C3 :	
	1. (bel S1 [(accept S2 p) => (accept S1 p)])	
[1] OFFER(S1 p)		
	Update : 2. (BMB(S1 S2 (Asn S1 (bel S2 (bel S1 [(accept S2 p) => (accept S1 p)])))))	Update : 3. (bel S2 (bel S1 [(accept S2 p) => (accept S1 p)])) 4. (BMB(S2 S1 (Asn S1 (bel S2 (bel S1 [(accept S2 p) => (accept S1 p)])))))
		Appropriateness condition : 5. (accept S2 p)
[2] ACCEPT(S2 p)		
	Update : 7. (BMB S1 S2 (Asn S2 (bel S1 (accept S2 p)]))	Update : 6. (BMB(S2 S1 (Asn S2 (bel S1 (accept S2 p)]))
	Inference : 8. (bel S1 (accept S2 p)) 9. [(bel S1 x) & (bel S1 (x => y))] => (bel S1 y) (axiom of Modal Theory S5) & 1. & 8. => 10. (bel S1 (accept S1 p))	

At this stage we can see that a somewhat one-sided set of contexts has been produced. If we look at elements 5, 8, and 10, we can see that S1 believes that S2 accepts p, and he believes, by inference, that he therefore accepts it as well. However, although S2 accepts p (and believes that he does so, assuming “[(accept S x) => (bel S (accept S x))]) he does not have a belief to the effect that S1 does so. In fact, S2 could only obtain such a belief from an explicit acceptance of S1, which is precisely what the derived element 10 above ((bel S1 (accept S1 p))) allows, since it can be transformed into the appropriateness condition for acceptance [(accept S1 p)]. Our fabricated utterance sequence above therefore needs to be completed to include a second acceptance, or “ratification.” This may be represented as follows.

(continued from above)

No.	Speaker	Utterances	Communicative acts
[3]	S1 :	Fine.	ACCEPT/ratify(S1, “discuss fractions”)

CA	Context K _{S1}	Context K _{S2}
	Inferred element : 10. (bel S1 (accept S1 p)) (assuming axiom of S5, [(bel S x) => x], and that variable x may take acceptances as values) => Appropriateness condition : 11. (accept S1 p)	
[3] ACCEPT(S1 p)	Update : 12. (BMB(S1 S2 (Asn S1 (bel S2 (accept S1 p)]))	Update : 13. (bel S2 (accept S1 p)) 14. (BMB S2 S1 (Asn S1 (bel S2 (accept S1 p)]))

On the basis of this analysis, we can therefore define **minimal agreement** as follows:

Minimal agreement

[(bel S1 (accept S1 p)) &
(bel S1 (accept S2 p)) &
(bel S2 (accept S2 p)) &
(bel S2 (accept S1 p))]
=> (AGREE S1 S2 p)

{element 10 above}
{element 8 above}
{element 5 above, [(accept S p) => (bel S (accept S p))]}
{element 13 above}

This definition possesses a desirable simplicity (no mutual beliefs are necessary) and is consistent with the CA definitions above. A more complete definition would take into account the following mutually believed propositions:

Mutually believed propositions

Elements 6,7 : (Asn S2 (bel S1 (accept S2 p)))
Elements 12,14 : (Asn S1 (bel S2 (accept S1 p)))

In other words, each agent assumes that the other believes that the first accepts the proposition. Perhaps more importantly, this analysis and definition of “minimal agreement” is consistent with the dialogue protocols to the extent that this type of OFFER, ACCEPTANCE, ACCEPTANCE/RATIFICATION sequence is one of the basic building blocks in negotiative dialogues. It may be summarised as follows:

Negotiation basic interaction sequence

- [1] OFFER[S1 p]
 [2] ACCEPT[S2 p]
 [3] ACCEPT/ratify[S1 p] [=] (AGREE S1 S2 p)

This type of sequence occurs mainly at the end of longer sequences, as a kind of “termination” or “closure,” for example as in the last three turns of the examples in Figures 1 and 2. Here, however, there is the added refinement that within turn [2], a “reformulation” occurs in order to check mutual understanding of what is agreed (communication-level negotiation). One of the reasons why this sequence usually occurs after a number of exchanges have already been made is that in order to make an offer that is immediately accepted without refinement, the offering agent needs to have quite a good *model of what the other is likely to accept*, and construction of such a model requires *interaction*. There are, of course, deviations from this basic sequence (for example, there are more extended offer sequences, and other types of CAs are used, such as questions and requests). However, the structure gives a basis on which to approach these differences, some of which are discussed in the next section. Finally, we have not considered the case of non-acceptance here, as mentioned above. However, we may briefly note that following non-acceptance of an offer, three main things may happen: The agent making the original offer may *refine* it, the agent may attempt to persuade the other to accept the original offer by *argumentation*, or the original offering agent may *accept* the non-acceptance (and so the dialogue may terminate¹¹).

Communicative Acts in Extended Negotiation Sequences

In sequences preceding the “basic interaction sequence” mentioned above, initial offers are often jointly and gradually refined towards this termination stage, by *successive related offers*, as well as by other types of communicative acts. For example, CAs in negotiation example 2 may be represented as follows (the terminating basic sequence is marked in **bold**):

Negotiation Example 2: Analysis of CAs

- | | |
|---|---------------------|
| 1/ S1 : so, what can we say if there's an inelastic impact? | WH-QUESTION |
| 2/ S2 : well, that the energy...all the energy... | OFFER |
| 3/ S1 : well, that the kinetic energy is theoretically nil! | OFFER |
| 4/ S2 : it's nil on arrival, in fact... | OFFER |
| 5/ S1 : since... since the object stops, in fact, | OFFER |
| ah yes, especially since there it doesn't move, uh... | OFFER |
| 6/ S2 : it's nil at the start, and it's nil on arrival... | OFFER |
| ... about energy...yes, but at the moment of an | |
| inelastic impact, what is it that... | |
| 7/ S1 : we've been doing that for a while now! <sighs> | WH-QUESTION |
| | <dialogue control |
| | act, time / problem |
| | solving resource |
| | management> |
| 8/ S2 : but we've also... | <incomplete |
| | utterance> |
| 9/ S1 : wait... | <dialogue control |
| | act, time |
| | management> |
| | OFFER |
| | <metacomment about |
| | what is believed |
| | mutually accepted> |
| | OFFER |
| | ACCEPT |
| 10/ S2 : Yes, | <reformulation of |
| the elastic impact, by contrast, both are conserved | what is believed |
| elastic impact, there's the total energy which | mutually accepted> |
| is conserved... | ACCEPT |
| 11/ S1 : Yes | |

In addition to the sequences of consecutive offers in this example, the initial “wh-” question is in fact also an implicit negotiation of a specific problem solving subgoal. We can also observe that in 10/ S2 accepts a specific restricted *reformulation* of 9/, that is, momentum and kinetic energy are accepted as “total energy,” and that this reformulation is itself accepted in 11/. Otherwise, 9/ to 11/ conforms to the basic negotiation sequence described above.

If negotiation is about reaching agreement, then throughout such a dialogue we need to be able to identify at any point *what is agreed and what is not*.¹² We have already considered the case of explicit agreement in the previous section, but we also claim that acceptance and rejection may be communicated *implicitly* within sequences of successive offers. To take a simple example, if you offer to take me to the theatre and I reply by offering to take you to the cinema, then this second offer (a “counter-offer”) amounts to an

implicit rejection of the first. If, by contrast, I offer to buy you a drink during the interval at the theatre, then this amounts to an *implicit acceptance* of your offer to take me to the theatre. These additional contextual effects of successive offers are analysed in the next section. Subsequently, we try to show that such effects depend, in general, on the *relations* (semantic, logical, pragmatic) between offered propositions, and a finite set of such relations for the case of collaborative problem-solving dialogues in physics is described. Finally, we discuss other types of communicative acts and negotiation in negotiations.

Sequences of Successive Offers and Their Contextual Effects

The basic negotiation interaction sequence can now be modified to include repeated offers, as follows:

Extended negotiation interaction sequence

[1] {OFFER[(S1/S2) f(p)]}*
[2] ACCEPT[S2/S1 p]
[3] ACCEPT/ratify[S1/S2 p]

[*x]* = x may be performed 1 to n>1 times ; "XY" = X and Y alternate on successive repetitions; "f" is a functional variable]

[=> (AGREE S1 S2 p)]

In element [1] of this sequence, "*f*(*p*)" means that a specific *function* is applied to the offered proposition, which may itself be the result of the application of several functions. The symbol "*f*" is a functional variable, since different functions may be applied on each iteration. The general idea is that an agent who "receives" an offer of a proposition from another agent, applies a function to it to return a new value, which is then offered in return. For example, in the analysis above, when S2 receives the offer "kinetic energy is theoretically nil" in 3/, S2 applies a function of the kind "render more specific," to give "kinetic energy is theoretically nil *on arrival*" as a new propositional value to be offered in turn. The range of functions of this kind that may be applied is described in the next section. (We shall also say that the application of such a function to a previously offered proposition establishes a *relation* between the consecutively offered propositions.)

We now need to consider the possible contextual effects of offers which succeed others. Specifically, the question arises as to whether in the following sequence, [2] has the effect that the agents concerned mutually believe that *p1* has been *accepted* or not:

(EG1) [1] OFFER[S1 p1]
[2] OFFER[S2 p2] [??? => (accept S2 p1)]

We claim that the answer to this question depends on the type of the relation "*R*" in "*R*(*p1*,*p2*).". An alternative answer is given by the "mathematical models for dialogue" of Hamblin (1971) and Mackenzie (1981, 1985). The approach consists in positing a dialogue rule for updating a "commitment slate," whereby *acceptance is assumed* with respect to a statement in the absence of explicit denial or retraction. As Mackenzie (1985) puts it "...in this [dialogic] game silence means assent" (p. 333). In other words, since in [2] of (EG1) above S2 has not "denied" *p1* (S2 has been "silent" with respect to it), then S2 *accepts p1*. We would argue that this approach is too simple to be able to be of use in modelling real dialogues. Our view may be summarised as follows:

Whether an OFFER[S2 p2] following an OFFER[S1 p1] has the effect that S2 accepts *p1* or not depends, in part, on the (semantic, logical, informational) *relations* between the offered propositions *p1* and *p2* and their sequential placement within the interaction.

This claim may be illustrated by the following hypothetical examples.¹³

(EG2) [1] OFFER[S1 red(object_1)]
[2] OFFER[S2 cube(object_1)] {accepts [1]}

(EG3) [1] OFFER[S1 cube(object_1)]
[2] OFFER[S2 red(object_1)] {accepts [1]}

In this pair of examples, since an object can have the colour red and be a cube, we would say intuitively that in [2] of both EG2 and EG3, S2 *accepts* [1], mainly because S2 *has not rejected it*. In addition, this effect is *not* sensitive to the *sequential ordering* of the interventions, but, as we shall see, *only in the case where the uttered propositions are logically and semantically independent*.

The following examples illustrate the effects of successive offers of propositions that are contradictory.

(EG4) [1] OFFER[S1 cube(object_1)]
[2] OFFER[S2 sphere(object_1)] {rejects [1]}

(EG5) [1] OFFER[S1 sphere(object_1)]
[2] OFFER[S2 cube(object_1)] {rejects [1]}

In these two examples, since an object cannot be both a cube and a sphere—the two are “incompatible”—then in both cases, [2] is a new offer which also *rejects* the previous one. The result is independent of sequential ordering.

(EG6) [1] OFFER[S1 coloured(object_1)]
[2] OFFER[S2 red(object_1)] {does not reject [1]}

(EG7) [1] OFFER[S1 red(object_1)]
[2] OFFER[S2 coloured(object_1)] {rejects [1]}

In these two examples, being “red” is a specific case of “being coloured.” In (EG6) we would say that [2] *does not reject* [1], or that it “accepts it as far as it goes,” and that a new offer is made which is more specific. Since [red(x) => coloured(x)] is common knowledge, it might therefore appear that in [2] of (EG6), S2 accepts [1]. This is not, however, necessarily the case since in [2], S2 does not necessarily *accept* red(object_1) (it is an offer, and hence conditional). We therefore do not posit any additional contextual effects for the case where OFFER[S1 p1] is followed by OFFER[S2 p2], and (BMB S1 S2 (p2 => p1)). This relation between successively offered propositions is, however, sensitive to sequential ordering. In (EG7), we would say that [2] is equivalent to saying “it is coloured, but it is not red,” or “it is coloured but is not necessarily red.” In certain contexts, [2] is therefore an offer that *rejects* [1]. It is important to realise that this effect is not achieved by purely *logical* means, since the assertion COLOURED(a) is not contradictory with the assertion RED(a), even given the fact that “RED(x) => COLOURED(x).” The effect is achieved via a general Gricean dialogue rule concerning the obligation to be “informative”¹⁴. Since (it is common knowledge that) being coloured is a simple consequence of being red, making an offer that something is coloured in this context is interpreted, assuming cooperativity, in a mutually exclusive sense, that is, as saying something “new.”

The same kind of analysis applies to examples (EG8) and (EG9).

(EG8) [1] OFFER[S1 (red(object_1) & cube(object_1))]
[2] OFFER[S2 red(object_1)] {rejects part of [1]}

(EG9) [1] OFFER[S1 red(object_1)]
[2] OFFER[S2 (red(object_1) & cube(object_1))]
{does not reject [1]}

In [2] of (EG8), S2 communicates that S2 rejects part of [1] (that the object is a cube). In (EG9), however, [2] does not necessarily accept [1] in making an

offer that conjoins the two propositions. That this is the case may be seen more clearly from the following example:

(EG10) [1] S1 : You can have the car for £1000
[2] S2 : £1000 for the car and the radio.

In [2] above, S2 clearly does not accept to give £1000 for the car alone, but only on the condition that the radio is included. Since this conditionality amounts simply to that of the new offer itself, we do not propose extra contextual effects.

Consider finally the following limiting case (identity) :

(EG11) [1] OFFER[S1 (red(object_1) & cube(object_1))]
[2] OFFER[S2 (red(object_1) & cube(object_1))]

In this case S1 would probably believe that S2 had not heard correctly, or that S2 wanted to confirm that he had understood. In either case, [2] is performing an indirect *dialogue control act*, that is, it is an attempt to establish mutual perception concerning what is offered. As we shall see below, a slightly different case occurs that is close to “identity,” where the second offer *reformulates* the first, in order to check mutual understanding of what is to be agreed.

We can now summarise these findings as a set of “dialogue rules.” As mentioned above, CAs such as “offers” may achieve certain contextual effects in virtue of their *position in an interaction sequence*. A well known criticism of speech act approaches is precisely that they define utterance types in a manner which is isolated from their interactive context (Franck, 1984). However, on methodological grounds, we prefer precisely to represent only those effects which *are* independent of interactive context in CA definitions, and to place interactive effects elsewhere in “dialogue rules.” In order to define dialogue rules referring to interventions in an interaction sequence we need to introduce the notion of a “stage.” The notion is derived in a modified form from Mackenzie (1981, 1985):

a dialogue consists of a succession of N stages

STAGE = <stage number, {CAs}>

where {CAs} are each represented as CA[<speaker> <proposition>].

We begin with dialogue rules concerning contradictory offers.

Dialogue Rule 1: successive incompatible offers

If a stage $\langle N+1, \text{OFFER}[S2\ p2] \rangle$ follows a stage $\langle N, \text{OFFER}[S1\ p1] \rangle$, and $p1$ and $p2$ are mutually believed by $S1$ and $S2$ to be logically or otherwise incompatible, then stage $N+1$ has the effect (BMB $S1\ S2 \sim (\text{accept } S2\ p1)$), in addition to the standard effects of $\text{OFFER}[S2\ p2]$.

Example: (EG4) and (EG5) above.

In other words: Offering a proposition that is contradictory with a previously offered proposition is understood as a *rejection* of that previously offered proposition.

Dialogue Rule 2: successive offers and set exclusion

If a stage $\langle N+1, \text{OFFER}[S2\ p2] \rangle$ follows a stage $\langle N, \text{OFFER}[S1\ p1] \rangle$, and $\{p2\}$ is an element of, or a subset of $\{p1\}$, then stage $N+1$ has the effect (BMB $S1\ S2 \sim (\text{accept } S2\ p1) - \{p2\}$), in addition to the standard effects of $\text{OFFER}[S2\ p2]$.

Example: (EG8) above.

In other words: Offering a proposition that is restricted to only a part of what has previously been offered amounts to accepting only the part in common between the two offers and not what is not in common.

Dialogue Rule 3: successive offers and generality

If a stage $\langle N+1, \text{OFFER}[S2\ p2] \rangle$ follows a stage $\langle N, \text{OFFER}[S1\ p1] \rangle$, and $p2$ is mutually believed to be a more general case of $p1$, then stage $N+1$ has the effect (BMB $S1\ S2 \sim (\text{accept } S2\ p1)$), in addition to the standard effects of $\text{OFFER}[S2\ p2]$.

Example: (EG7) above.

In other words: Offering a proposition that is more general than a previous offer amounts to rejecting the previous more specific offer (if you claim that we are in Lyon and I claim that we are in France, given that we both know we are, then I'm claiming that we're in France but not in Lyon).

The following rule describes how acceptance may be indirectly communicated:

Dialogue Rule 4: successive offers and logical independence

If a stage $\langle N+1, \text{OFFER}[S2\ p2] \rangle$ follows a stage $\langle N, \text{OFFER}[S1\ p1] \rangle$, and $p1$ and $p2$ are mutually believed to be (logically, semantically, causally, etc.) independent, then stage $N+1$ has the effect (BMB $S1\ S2 (\text{accept } S2\ p1)$), in addition to the standard effects of $\text{OFFER}[S2\ p2]$.

Example: (EG2) and (EG3) above.

In other words, if I offer something which in no sense depends on your offer, and don't explicitly reject your offer, then I implicitly accept it. Although this resembles the dialogue rule "silence = assent" of Hamblin and Mackenzie, mentioned earlier, our rule is restricted to one very special case. Note also that $\langle N, \text{OFFER}[S1\ p1] \rangle \langle N+1, \text{OFFER}[S2\ p2] \rangle$ is not therefore the same as $\langle N, \text{OFFER}[S1\ p1] \rangle \langle N+1, \text{OFFER}[S2\ p1 \ \& \ p2] \rangle$.

In this section we have described how in extended sequences of successive offers, additional contextual effects may be produced in the interaction, depending on the type of semantic, logical and pragmatic *relations* established between successively offered propositions, and depending on the placement of the offers within the sequence. We certainly do not claim to have described *all* the dialogue rules necessary for modelling negotiative dialogues; but it seems likely that contradictory (counter) offers, offers which expand or "contract" previous offers, and offers that are mutually independent are the principal cases. In other terms, we stated that a new offer may result from the application of a "transformation function" to the previous one. In the next section we describe a specific set of such relations and transformation functions which create them, for a specific problem-solving domain (that of negotiation example 2).

Relations Between Successive Offers and "Transformation Functions"

We have made an important simplification in saying that a relation exists between two successively offered propositions, since, the proposition of a given utterance in fact relates more generally to the *context* (partly) created by a previously uttered one. In establishing such relations and transformation functions for given pairs of utterances we may therefore need to identify the contexts in question, since the proposition to which an offer relates may not have been explicitly expressed. Since dealing with this issue would involve a lengthy excursion into general pragmatics, we only consider explicitly expressed propositions in this paper.

We now describe in more detail the different classes of TFs that occur in negotiation example 2, and the corpus from which it is drawn. Since these functions are quite general in nature (they are mostly logical, set and well-known semantic relations) we may make the conjecture that these functions will also apply to a greater or lesser extent to other negotiation domains.

In general, a new offer may **expand** the propositional content of a previous one, it may **contract** it, it may be **neutral** with respect to it, or it may concern the **foundations** of the propositional content of the previous offer. Such "transformation functions" bear some resemblance to "coherence relations"

(Hobbs, 1982) and "rhetorical relations" (Mann & Thompson, 1988) that have been described for texts. Here we use the notion of application of a function in order to capture the active co-construction of relation between utterances by agents in interaction. Although the "elaboration" coherence (Hobbs, op. cit.) and rhetorical (Mann & Thompson, op. cit.) relations, for example, are general cases of the "expansion" relations described below, it is not contentious to claim that relations between utterances are quite different to relations between text segments. Whilst relations between text segments serve a purpose with respect to comprehension and acceptance on the part of an envisaged class of readers, relations established between utterances of different speakers are concerned with *mutual* comprehension and acceptance, and may be dynamically revised by both (writers do not usually revise the intentions and beliefs behind their books whilst other readers are reading them!).

Members of the four main classes of TFs are described below. If we exclude the case of Dialogue Rule 1 (successive incompatible offers), then the only additional contextual effects are assumed to be those specified in the above rules.

Classes of Transformation Functions

Expansion TFs. These "elaborate" on the content of the initially offered proposition, thus "expanding" its content in some way. The following are abstract representations of specific expansion TFs that have been analysed in the corpus of negotiation example 2. Some concrete examples are given at the end.

Sequence	(Expansion) transformation function
[1] OFFER[S1 a] [2] OFFER[S2 (a & b)]	conjunction
[1] OFFER[S1 a] [2] OFFER[S2 (a v b)]	disjunction
[1] OFFER[S1 a] [2] OFFER[S2 (a => b)]	inference
[1] OFFER[S1 (∃a)G(a)] [2] OFFER[S2 (∀a)G(a)]	generalisation
[1] OFFER[S1 a] [2] OFFER[S2 G(a)]	additional-property

[1] OFFER[S1 B(a,b,x?)] [2] OFFER[S2 (x? = c)]	specific-value
[1] OFFER[S1 a] [2] OFFER[S2 (a ∈ G)]	categorisation

The following are examples from negotiation example 2 (a complete analysis of the example will be given at the end of this section when other TF categories have been described).

Negotiation Example 2

3/ S1 : well, that the kinetic energy
is theoretically nil !

4/ S2 : it's nil on arrival, in fact ...

5/ <...>

6/ S2 : it's nil at the start, and it's nil on arrival ...

OFFER[S1 KE(obj_1, nil, loc?)]
OFFER[S2 (loc? = "point of arrival")
(Expansion)TF = specific-value

(Expansion)TF = conjunction

Contraction TFs. These functions "contract" or restrict the content of the initially offered proposition in some way. To a certain extent, they are therefore the *inverses of expansion functions*. The following are some abstract examples.

Sequence	(Contraction) transformation function
[1] OFFER[S1 (a & b)] [2] OFFER[S2 a]	conjunction-contraction
[1] OFFER[S1 (a => b)] [2] OFFER[S2 a]	contra-inference
[1] OFFER[S1 (a => b)] [2] OFFER[S2 b]	contra-reason
[1] OFFER[S1 (a v b)] [2] OFFER[S2 a]	exclusive-disjunction-choice
[1] OFFER[S1 (∀a)G(a)] [2] OFFER[S2 (∃ a)G(a)]	particularisation
[1] OFFER[S1 G(a)] [2] OFFER[S2 a]	contract-property

5/ S1: since...since the object stops, in fact, ah yes, especially since there it doesn't move, uh...
(gives *reason* as to why energy nil on arrival)

The following is an example from negotiation example 2.

Negotiation Example 2

2/S2: well, that the energy... all the energy... OFFER[S2 energy(x)]
3/S1: well, that the kinetic energy <...> OFFER[S1 kinetic_energy(x)]
(Contraction)TF = subtype

Note that in this example, the fact that "kinetic energy" is a subtype of "energy" is not explicitly stated, but is a general common assumption; that is, as stated above, utterance 3/ relates to the context that is partially created by utterance 2/ rather than to the utterance itself.

Neutral TFs. These leave the content of the initially offered proposition either completely unchanged (the new offer is identical to the previous one) or the new offer may be "equivalent in meaning" to the previous one. Since the "content" is the same, these functions cannot be operating at the level of negotiation of the problem solution itself. In fact, they operate at the level of negotiating joint understanding (in which case they often are used *after* explicit acceptances), such as "summaries" at the end of long sequences of the agreed solution, or statements that may be glossed as "if by x you mean y, then I agree...." In the following extract from negotiation example 2, S2 accepts explicitly, then repeats and *reformulates* the previous offer to provide a check mutual understanding of what has been agreed:

9/ S1: <...> elastic impact, by contrast, both are conserved ...
10/ S2: Yes, elastic impact, there's the total energy which is conserved...

In such cases, the offer does not change the standard effects of acceptance, but rather *delays* them until the following utterance, in which the next speaker may *confirm* that these effects are mutually understood or not.

Foundational TFs. These are somewhat different from other categories in that the informational relation does not operate to *construct* the "package" offer to be agreed, but rather to provide *foundations* for it, such as giving *reasons* for accepting it or not, or giving *explanations* as to why it may be the case. The following is an example.

Negotiation Example 2

4/ S2: it's nil on arrival, in fact...




Figure 6. Transformation functions for negotiation example 2

From the analysis in Figure 6 it can be seen that the model presented previously for the application of transformation functions in sequences is idealised in several respects, notably:

- an offer does not necessarily relate to (the context created by) the precedent utterance;
- an offer may relate to (the context created by) several previous utterances;
- an offer by a speaker may relate to (the context created by) their own previous utterances as well as to those of another; and
- since several propositions may be stated within an utterance, transformation functions may be applied within the same utterance.

Our model assumes that in order to fully understand the effects of utterances and to build up a representation of what is mutually accepted, speakers must be able to identify the *elements of contexts to which an uttered proposition relates*, and the *transformation function that has been applied*.

This completes description of the major components of our model for negotiation in teaching-learning dialogues. In the next section some extensions to the model are briefly discussed, concentrating on additional negotiation and their attendant communicative acts.

Other Communicative Acts and Negotia

In the previous discussion we have been concerned with negotiation at the domain or task level, and have restricted the analysis to propositions. Extending the model to other task-level negotia—viewpoints, goals, methods, etc.—does not, we would claim, involve fundamental changes to the set of communicative acts, their combination into sequences, and the direct or indirect contextual effects involved. All that is required is that the agents concerned have an explicit representation of these phenomena as negotia, and that the agents have some means of evaluating their acceptance or not of such negotia. Since means of representing such negotia are relatively well known, we now describe *communication-level negotia* in terms of our model.

Negotiation at the communication level is achieved by what Bunt (1989) terms “dialogue control” (“DC”) acts (see the earlier literature review). The main classes of attendant DC functions are shown in Figure 7.




Figure 7. Classes of dialogue control functions

It is important to realise that a given utterance may have a function at the task level (such as solving a common problem) and at the dialogue control level (for example, it may be an *interruption*, thus relating to turn and/or time management). DC functions are therefore usually realised by other “standard” types of communicative acts, such as different types of “questions.” They differ in that the “content” of the act is not necessarily *propositional* in nature. For example, the utterances “Do you mean the red ball?” and “Did you say ‘the red ball?’” are yes/no questions whose contents are a semantic representation and a linguistic representation, respectively, and which give *feedback* on the *understanding* and *perception* levels. This implies that the dialogue context needs to be enriched to incorporate an explicit representation of stages (dialogue history). Feedback on the “evaluation” level is given by communicative acts such as “acceptances,” discussed earlier, which usually have no propositional or other “content” (for example, the utterances “Yes” or “Ok”).

Turning to “discourse structuring” DC acts, the following example from the same corpus as Negotiation example 1 shows how “terminating the dialogue” can itself be a negotium :

Teacher : What do you want to do then?

Student : Stop?

Teacher : Oh. Really? I'd quite like to talk through this one...

[dialogue delimitation]

Here, as with many other DC acts, the “content” of the act is the termination of a (dialogue) *process* which implies that an agent must have the capability to reason about its own process execution (for example, that a speech recognition module has failed, and why).

The following are examples of “interaction management” DC acts in the two examples previously considered:

Negotiation Example 1

1/Teacher: I would greatly enjoy doing some prolog [topic management]
 <...>
 4/Learner: Um. I'd like to do something now [time management]
 5/Teacher: . Ummm... [turn management]
 <...>
 6/Learner : ...um can you do something else in another window while I think about it? [time management]

Negotiation Example 2

7/ S1 : we've been doing that for a while now! [time management]

Simplifying somewhat, we can represent DC functions and associated types of negotiations for dialogue delimitation and turn/topic management in terms of the concept of a stage in dialogue, described earlier. Representing negotiation of time management would require real-time temporal units to be attached to sequences of stages.

Dialogue delimitation: For “closing,” in stage N the speaker S1 offers that the next two stages should contain no communicative acts,

[1] OFFER[S1 [<N+1, {}> & <N+2, {}>]

Turn management: For self-selection, the speaker S1 offers in stage N that S1 should utter some set of communicative acts in stage N+2 (given acceptance in stage N+1),

[1] OFFER[S1 [<N+2, {CA[S1 N]}>]

Topic management: The speaker S1 offers in stage n that in stage N+1 the hearer S2 makes some communicative act with respect to a proposition p2 that is relevant to the focus¹⁵ p1 of stage N:

[1] OFFER[S1 [<N+1, {CA[S2 p2] & relevant(p2 p1)}>]

We now summarise the main theoretical assumptions of the model presented in this paper, and discuss some remaining research problems.

Summary: Theoretical Assumption of the Model

We shall now attempt to render explicit a more general theory of negotiation underlying the model, and to discuss some fundamental problems that arise. This abstraction from details of the model may be useful in deriving testable predictions with respect to learning.

Agents negotiate when they attempt to reach agreement with respect to something by interacting together, and when they each recognise the right of the other to use similar means to reach such an agreement. They will usually want to agree on some aspect of a *cooperative task*, the achievement of which motivated them to interact, and they will always need to agree on meanings of utterances, coordination and control of dialogue and task-related actions. When one agent makes an offer with respect to reaching agreed-upon actions, the other agent has a number of choices: (a) The receiving agent may *reject* the offer in the case where it violates his fundamental constraints, or is contradictory with his beliefs or aims to which he is firmly committed. In this case, the first agent may try to make the other agent accept the original offer (by *persuasion* or by *argumentation*) if he is similarly committed to it and believes that this may be successful, or he may simply accept the rejection (and thus the negotiation terminates with respect to the things for which agreement was originally sought). (b) Alternatively, if the receiving agent can find no reason to reject the offer and it appears to converge towards something which might be preferable to him, and he has no concrete offer to make with respect to the original offer, he may “stand pat” and encourage the other to further develop his offer. (c) Lastly, the receiving agent may find no reason to reject the offer, he may view it as wholly or partially contributing to a possible agreement, and he may be able to think of some way of transforming the offer to a new one that satisfies his own constraints and the constraints that he attributes to the other. Thus the agent could extend the offer in some way, restrict it to those aspects that he finds acceptable, or he may reformulate it in some way to check that both understand the offer in the same way. Understanding the way in which an offer has been so transformed is necessary if agents are to co-construct a common representation of what is agreed and what is not. Depending on the way the second agent refines the original offer, the first agent may understand the new offer to be a partial or complete rejection or acceptance of his offer, or to be neutral with respect to it. This refinement process may continue towards an agreement, the agents may “leave it at that” if they think that this is the best that can be achieved at present, or one of the other alternatives may be chosen (reject, stand pat). Negotiation is thus a complex interactive process that enables agents to achieve mutual understanding, to pursue common goals, and to collaborate in achieving them.

With respect to the refine strategy, an implementation of the model in future work would require the definition of a number of specific procedures, other than those required for a general model of dialogue. In particular, an agent capable of negotiating in this limited sense needs to be able to make two main decisions:

1. whether to accept, reject or refine an offer made by the other agent, and
2. given a decision to refine the previous offer with a new one, which transformation function to choose.

We discuss how the first decision may be made in conclusion, but it is clearly related to the second since the availability or not of an acceptable transformation is one factor influencing acceptance itself. In general, we may hypothesize that the TF is chosen that is viewed as best satisfying the constraints of beliefs of both speakers. Finally, a complete definition of the logical properties of "acceptance," to which we have appealed in this paper, has not been provided. Rather, we have concentrated on motivating the necessity for such a definition from detailed dialogue analysis. We restrict ourselves to some preliminary remarks on acceptance and belief in conclusion.

CONCLUSIONS

In conclusion we discuss two main issues that arise from the model presented here: implications of the use of the notion of "acceptance" for student and dialogue modelling, and the role of negotiation in learning and problem-solving.

Belief Is Not Enough (for Agent and Student Modelling)

In Cohen's (1992) terms, "acceptance" of a proposition differs from belief in that "Belief is a disposition to feel, acceptance a policy for reasoning." To accept a proposition, is to be willing to use it in one's reasoning, but not necessarily to adopt it in any reasonably stable sense (like belief). In addition, it is claimed, although one can *decide* to accept something (as a policy), one can not genuinely decide to believe something. To this extent, acceptances are similar to the special type of justifications introduced in Doyle's (1979) truth maintenance system called "conditional proof justifications," which are basically hypotheses that may be simply dropped when

the validity of a conclusion derived from them in no way depends on the hypothesis.

A partial illustration is as follows. Assume that a particular agent can derive a proposition p2 from proposition p1 (whether the reasoning processes involved are valid or not): $p1 \Rightarrow p2$. Then, in a "foundations" belief system, for example, p2 is believed when p1 is believed (assuming that this is the only valid justification), and p2 is "dropped" as a belief when p1 is no longer believed. With acceptance, however, according to Cohen, an agent will accept p1 if the agent wants to accept p2, or believes that p2 is the best compromise attainable (and so on, as discussed earlier), since p1 appears to be a fruitful premise in reasoning. Supposing that p2 is not subsequently accepted, for example, then the agent may drop p1 (it serves no useful purpose in joint reasoning). Similarly, a proposition p may be accepted simply because it appears to be the best joint solution at present. An example (from the corpus of Figure 2) is where two students accept p1: "potential energy explains rebound behaviour" since they believe that this implies p2: "mass explains rebound behaviour" (mass is a term in that energy equation, and the balls have different masses), and they both already want to accept p2 (or, it is one of their "hypotheses").

Given our communicative act analysis, therefore, *what a speaker says* with respect to the cooperative task to be achieved (buying a carpet, learning, problem solving, etc.) is not necessarily a direct reflection of what the speaker *believes* but is, we would argue, rather a reflection of what the speaker is willing to *accept* in the current situation, and rules for accepting or revising acceptances and rules for adopting a belief or revising one may, as even Doyle's early work suggests, be quite different. The reason why someone may or may not accept a proposition is very often, of course, related to what the person believes, but there are other factors involved. For example, someone (such as a learner) may accept something, thus manifesting what may be taken to be "belief," because the person believes that this is a proposition that is close enough to his beliefs and is the best "compromise" that could be agreed with the other, given his beliefs about the other's beliefs. Or again, Gutwin (1992) has argued that in some circumstances, a teacher may have good pedagogical reasons for accepting a proposal by the student, or making a proposal himself, even though the teacher does not believe it.

In addition to negotiations, we are also required to posit different attitudes from belief with respect to argumentation (see Dennett, 1981). For example, when one agent *concedes* an argument (because he can produce no acceptable attack or defense), we would not want to say that he necessarily

believes the opponent's thesis, but rather that he is *compelled to accept* it. It may be that such considerations apply to *any* type of dialogue produced with respect to achieving some cooperative task. We would therefore argue that student modelling techniques that depend on analysis of students' utterances in interaction with a teacher (human or artificial) should incorporate mechanisms for acceptance as well as belief revision. In fact, much research on student modelling (e.g., Huang, 1992; Huang et al., 1991; Self, 1992) concentrates on belief revision mechanisms, and simply *assumes* that a belief will be "input" to the revision mechanism in some way from the learning environment. Although this decoupling of students' attitudes from the interactive context from which they are supposed to be derived appears to be a necessary simplification as a research method, it may have some misleading consequences concerning the students' beliefs, since acceptances are adopted and revised differently from beliefs. New (unsolved) research problems therefore arise: Given mechanisms for belief and acceptance revision, *how do beliefs and acceptances interrelate? Under what circumstances does acceptance become belief?*

At present the logical properties of acceptance have been little explored in AI (Cohen's "Essay on Belief and Acceptance," 1992, lays the groundwork for such a project) in comparison with belief, but this appears to be a promising avenue of research for AI and AI and Education research, particularly given the specificities of teaching-learning dialogues. We do not claim to have solved these theoretical problems concerning acceptance in this paper, but we have showed that such a notion is necessary for modelling negotiative dialogues in teaching-learning situations.

Negotiating Learning and Learning by Negotiating

Finally, given a model for how negotiative dialogues are generated in teaching-learning situations, we can begin to understand what functions they serve, and thus advance some hypotheses concerning their role in teaching, learning and problem-solving. Essentially, negotiation is a mechanism for achieving *coordination* of problem-solving and communicative actions, for achieving mutual *understanding* of such actions and for attaining *agreement* at an epistemic level.

In given learning situations one type of negotiation may be emphasised rather than another, depending on the desired learning outcome, the relative competence of the agents involved, and their respective willingness and desires to negotiate at given levels. In other words, it all depends on the *latitude* that each agent has on a number of different dimensions. For example,

in some teaching situations, learning on the part of the student may require coordinated goals with strong constraints on the part of the teacher, in others the teacher may impose the goals, and in interactions between students who are relative equals there may be relatively free negotiation of goals. In some learning situations one agent may have the goal of enabling another to understand a problem in terms of a particular viewpoint (such as understanding a mechanics experiment in terms of energy), and in others a new viewpoint may be co-constructed by the agents. Whatever the desired outcome, it is difficult to see how any interaction can be successful in its general goals without coordination, understanding and agreement. By constraining learning situations, and negotiation possibilities in various ways we can thus favour negotiation on any or all of these levels. Our claim can therefore be seen as a more specific version of the view of Elsom-Cook (1990) that teaching interactions may be characterised along a dimension of freedom-constraint on the learner: For us this dimension becomes freedom-constraint *to negotiate* and *with respect to what*.

A negotiative Intelligent Tutoring System or Collaborative Learning System must therefore know on what aspects it is willing to negotiate and on what aspects it is not, for a given learner. Given that this decision can be made intelligently, we have described the form that such negotiations should take.

References

- Adler, M.R., Alvah, B.D., Weilmayer, R., & Worrest, W. (1988). Conflict-resolution strategies for nonhierarchical distributed agents. In *Distributed artificial intelligence, Volume II* (pp. 139-161). London: Pitman Publishing.
- Airenti, G., Bara, B.G., & Colombetti, M. (1989). Knowledge for communication. In Taylor, M.M., Neel, F. & Bouwhuis, D.G. (eds.) *The Structure of Multimodal Dialogue* (pp. 145-158). Amsterdam: North-Holland.
- Allwood, J., Nivre, J., & Ahlén, E. (1991). *On the semantics and pragmatics of linguistic feedback*. Gothenburg Papers in Theoretical Linguistics No. 64. University of Gothenburg, Department of Linguistics, Sweden.
- Baker, M.J. (1989a). A model for tutorial dialogues based on critical argument. In D. Bieman, J. Breuker, & J. Sandberg (Eds.), *Artificial Intelligence and Education*, (pp. 2-8). Amsterdam: IOS Publishers.
- Baker, M.J. (1989b). *Negotiated tutoring: An approach to interaction in intelligent tutoring systems*. Unpublished PhD thesis, Institute of Educational Technology, The Open University, Milton Keynes, Great Britain.

- Baker, M.J. (1990). Arguing with the tutor: A model for tutorial dialogue in uncertain knowledge domains. In *Guided-discovery tutoring*, (ed.) Elsom-Cook, M., p. 141-170. London: Paul Chapman Publishing.
- Baker, M.J. (1992). Modelling negotiation in intelligent teaching dialogues. In Elsom-Cook, M. & Moyses, R. (Eds.) *Knowledge negotiation* (pp. 199-240). London: Paul Chapman Publishing.
- Bange, P. (1992). *Analyse conversationnelle et théorie de l'action*. Paris: Hatier-Didier.
- Barbieri, M.S., & Light, P. (1992). Interaction, gender and performance on a computer-based problem-solving task. *International Journal of Educational Research*, 2 (3), 199-214.
- Barth, E.M., & Krabbe, E.C.W. (1982). *From axiom to dialogue: A philosophical study of logics and argumentation*. Berlin: Walter de Gruyter.
- Behrend, S., Singer, J., & Roschelle, J. (1988). A methodology for the analysis of collaborative learning in a physics microworld (pp.48-53). *Proceedings of ITS-88*, Montréal 1988.
- Bennett, N. (1976). *Teaching styles and pupil progress*. London: Open Books.
- Bond, A.H., & Gasser, L. (1988). *Readings in distributed artificial intelligence*. San Mateo, CA: Morgan Kaufmann.
- Bruner, J.S. (1986). *Actual minds, possible worlds*. London: Harvard University Press.
- Bunt, H.C. (1989). Information dialogues as communicative action in relation to partner modelling and information processing. In *The structure of multimodal dialogue* (Eds.) Taylor, M.M., Neel, F. & Bouwhuis, D.G., pp. 47-74. North-Holland: Elsevier Sciences Publishers.
- Clancey, W. (1987). *Knowledge-based tutoring: The GUIDON program*. Cambridge, MA: MIT Press.
- Clancey, W.J. (1992). "Situating" means coordinating without deliberation. Paper presented at the McDonnell Foundation Conference "The Science of Cognition," Santa Fe, NM, June 1992.
- Clark, H.H., & Schaeffer, E.F. (1989). Contributing to discourse. *Cognitive Science*, 13, 259-294.
- Cohen, L.J. (1992). *An essay on belief and acceptance*. Oxford: Clarendon Press.
- Cohen, P.R., & Perrault, R. (1979). Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3, 177-212.
- Collins, A., & Stevens, A.L. (1983). A cognitive theory of inquiry teaching. In Reigeluth, C.M. (Ed.), *Instructional-design theories and models: An overview of their current status*. Hillsdale NJ: Lawrence Erlbaum Associates.
- Davis, R., & Smith, R.G. (1983). Negotiation as a metaphor for distributed problem-solving. *Artificial Intelligence*, 20(1), 63-109.
- Dennett, D.C. (1981). How to change your mind. In *Brainstorms: Philosophical essays on mind and psychology* (pp.300-309). Brighton, UK: Harvester Press.
- Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence*, 12(3), 231-272.
- Druckman, D. (1973). *Human factors in international negotiations: Social-psychological aspects of international conflict*. Sage Professional Paper in International Studies 02-020. Beverly Hills & London: Sage Publications.
- Edmondson, W. (1981). *Spoken discourse: A model for analysis*. London: Longman.
- Edwards, D., & Mercer, N. (1987). *Common knowledge: The development of understanding in the classroom*. London: Methuen.
- Emmeren, F. van, & Grootendorst, R. (1984). *Speech acts in argumentative discussions. A theoretical model for the analysis of discussions directed towards solving conflicts of opinion*. Dordrecht/Cinnaminson: Foris Publications.
- Emmeren, F. van, & Grootendorst, R. (1991). The study of argumentation from a speech act perspective. *Pragmatics at issue* [Selected Papers of the International Pragmatics Conference 1987, vol. 1], Verschueren, J. (Ed.), pp. 151-170. Amsterdam: John Benjamins.
- Elsom-Cook, M. (1990). Analysis of a tutorial dialogue. In M. Elsom-Cook (Ed.) *Guided-discovery tutoring: A framework for ICAI research* (pp. 113-131). London: Paul Chapman Publishing.
- Fox, B.A. (1993). *The human tutorial dialogue project: Issues in the design of instructional systems*. Hillsdale NJ: Lawrence Erlbaum Associates.
- Franck, D. (1984). Speaking about speech acts. *Journal of Pragmatics*, 8, 87-92.
- Galliers, J.R. (1989). *A theoretical framework for computer models of cooperative dialogue, acknowledging multi-agent conflict*. Unpublished PhD thesis, Human Cognition Research Laboratory, The Open University, UK.
- Gasser, L., & Huhns, M.N. (1989). *Distributed artificial intelligence, vol. II*. London: Pitman.
- Gazdar, G. (1981). Speech act assignment. In *Elements of discourse understanding*, (Eds.) Joshi, A., Webber, B. & Sag, I., pp. 64-83. Cambridge: Cambridge University Press.
- Goodyear, P., & Stone, C. (1992). Domain knowledge, epistemology and intelligent tutoring in social science. In Moyses & Elsom-Cook, pp. 69-96.
- Grice, H.P. (1975). Logic and conversation. In Cole, P. & Morgan, J. (Eds.), *Syntax and semantics 3: Speech acts*. New York: Academic Press.
- Grosz, B.J., & Sidner, C. (1986). Attention, intentions and the structure of discourse. *Computational linguistics*, 12(3), 175-204.
- Gutwin, C. (1992). The use of pedagogic misrepresentation in tutorial dialogue. *Intelligent Tutoring Systems: Proceedings of ITS'92*, lecture notes in computer science 608, pp. 507-514. Berlin: Springer-Verlag.
- Hamblin, C.L. (1971). Mathematical models of dialogue. *Theoria*, 2, 130-155.
- Hobbs, J. R. (1982). Towards an understanding of coherence in discourse. In *Strategies for natural language processing*, (Eds.) Lehnert, W.G. & Ringle, M.H., pp. 223-244. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Huang, X., McCalla, G., & Neufeld, E. (1991). *Using attention in belief revision*. ARIES Lab. Research Report 91-1, Dept. of Computational Science, University of Saskatchewan, Saskatoon, Canada.

- Huang, X. (1992). Inconsistent beliefs, attention, and student modelling. *Journal of Artificial Intelligence in Education*, 3(4), 417-428.
- Joab, M. (1990). *Modélisation d'un dialogue pédagogique en langage naturel*. Thèse de doctorat de l'Université Paris VI, France.
- Lave, J., & Wenger, E. (1989). *Situated learning: Legitimate peripheral participation*. Institute for Research on Learning, Pittsburgh, PA. Report No. IRL 89-0013.
- Lave, J. (1988). *Cognition in practice*. Cambridge: Cambridge University Press.
- Levinson, S.C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Mackenzie, J.D. (1981). The dialectics of logic. *Logique et Analyse*, 24, 159-177.
- Mackenzie, J.D. (1985). No logic before Friday. *Synthese*, 63, 329-341.
- Mann, W.C., & Thompson, S.A. (1988). Rhetorical structure theory: Toward a functional theory of text organisation. *Text*, 8(3), 243-281.
- Moeschler, J. (1985). *Argumentation et conversation: Éléments pour une analyse pragmatique du discours*. Paris: Hatier-Crédif.
- Moyse, R., & Elsom-Cook, M. (1992). *Knowledge Negotiation*. London: Academic Press.
- Moyse, R. (1992). VIPER: The design and implementation of multiple viewpoints for tutoring systems. In Elsom-Cook & Moyse (Eds.), 97-134.
- Myers, T., Brown, K., & McGonigle, B. (Eds.). (1986). *Reasoning and discourse processes*. London: Academic Press.
- Perrault, R. (1990). An application of default logic to speech act theory. *Intentions in communication*. Eds. Cohen, P.R., Morgan, J., & Pollack, M.E. Cambridge MA: MIT Press.
- Pilkington, R., Hartley, J.R., Hintze, D., & Moore, D.J. (1992). Learning to argue and arguing to learn: An interface for computer-based dialogue games. *Journal of Artificial Intelligence in Education*, 3 (3), 275-295.
- Resnick, L.B., Salmon, M.H., & Zeitz, C.M. (1991). The Structure of reasoning in conversation. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pp. 388-393.
- Rimmershaw, R. (1992). A discourse repertoire for negotiating explanations. In Moyse & Elsom-Cook, pp. 241-276.
- Roschelle, J., & Behrend, S.D. (in press). The construction of shared knowledge in collaborative Problem Solving. To appear in O'Malley, C. (Ed.), *Computer supported collaborative learning*. Berlin: Springer-Verlag.
- Roschelle, J., & Clancey, W.J. (1992, in press). Learning as social and neural. Paper presented at the AERA Symposium, Implications of Cognitive Theories of How the Nervous System Functions for Research and Practice in Education, April 1991, Chicago. To appear in special issue of *The Educational Psychologist*.
- Rosenschein, J.S. (1992). Consenting agents: Negotiation mechanisms for multi-agent systems. *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence [IJCAI-93]*, 792-799. San Mateo, CA: Morgan Kaufmann.

- Roulet, E. (1992). On the structure of conversation as negotiation. In *(On)searle on conversation* (pp. 91-100). Eds. Parret, H. & Verschueren, J. Amsterdam: John Benjamins.
- Sandberg, J., & Wielinga, R. (1990). How situated is cognition? Research report, Social Science and Informatics, University of Amsterdam.
- Searle, J. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge University Press.
- Seely-Brown, J.S. (1990). Toward a new epistemology of learning. Keynote address to ITS '88 (Montréal). In Frasson, C. & Gauthier, J. (Eds.), *Intelligent tutoring systems at the crossroads of AI*, pp. 266-282. Norwood NJ: Ablex Publishing.
- Self, J. (1992). Computational viewpoints. In Moyse & Elsom-Cook, pp. 21-40.
- Sleeman, D., & Brown, J.S. (1982). *Intelligent tutoring systems*. Academic Press: London.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford: Blackwell.
- Sycara, K. (1988). Resolving goal conflicts via negotiation. *Proceedings of the AAAI-88 Conference*, St. Paul, MN (USA), 245-250.
- Sycara, K. (1989). Multiagent compromise via negotiation. In *Distributed artificial intelligence: Volume II*, (Eds.) Gasser, L. & Huhns, M.N., pp. 119-137. London: Pitman Publishing.
- Vanderveken, D. (1990). *Meaning and speech acts. Volume I: Principles of language use*. Cambridge: Cambridge University Press.
- van Rees, M.A. (1992). The adequacy of speech act theory for explaining conversational phenomena: A response to some conversational analytical critics. *Journal of Pragmatics*, 17, 31-47.
- Vernant, D. (1992). Modèle projectif et structure actionnelle du dialogue informatif. In *Du Dialogue : Recherches sur la philosophie et le langage*, No. 14, pp. 295-314, (Ed.) Vernant, D. Paris: Vrin.
- Wenger, E. (1987). *Artificial intelligence and tutoring systems*. Los Altos, CA: Morgan Kaufmann.

Notes

1. According to Roschelle & Behrend (in press, p.1) "collaboration" is distinguished from "cooperation" in that cooperative work "...is accomplished by the division of labor among participants, as an activity where each person is responsible for a portion of the problem solving...", whereas collaboration involves the "...mutual engagement of participants in a coordinated effort to solve the problem together." (ibid, p.1).
2. See Baker (1992) for a discussion of negotiating goals in teaching dialogues.
3. See Elsom-Cook (1990) and Galliers (1989) for definitions of specific forms of "conflict" of this kind.

4. See, for example Barth & Krabbe (1982), van Eemeren & Grootendorst (1984, 1991) on argumentation dialogue, and Pilkington et al. (1992) for an application to AI and Education.
5. We use the term "communicative acts" rather than "speech acts" in order to avoid associations with exclusively verbal, rather than written, interactions.
6. See Levinson (1983) and Franck (1984) for influential critiques; van Eemeren & Grootendorst (1991) and van Rees (1992) for defenses.
7. Clearly, realising the epistemic effects of a weak y/n question by, for example, injecting a drug into the hearer, would not count as having "asked a question" in the sense of communicative action.
8. For any communicative act there may be far-reaching and unpredictable effects on the mental states of addressees (perlocutionary effects). These are important, but are not taken into account in the *characterisation* of communicative acts.
9. Note that "acceptance" as we understand it there is an *attitude* (positive and evaluative) of an agent with respect to a semantic "content." It is therefore not at all the same as "acceptance" in the sense of Clark & Schafer (1989), which is rather a type of feedback that indicates that a contribution is "received," or understood.
10. In Edmondson's terms this is a "discourse external act," meaning that it is addressed to the domain of discourse, here understood as the "task."
11. Roulet (1993) describes this latter possibility in terms of a general constraint on "interactional completeness."
12. The space of what is agreed is a more restricted form of the "Joint Problem Space" of Behrend, Singer, & Roschelle (1988).
13. We use hypothetical examples here simply for clarity of exposition. The conclusions drawn do apply to the real dialogues, some of which are illustrated in the later section on relations between successive offers. Since any model makes a selection from observables, the model does not take all of them into account.
14. For a discussion of how such communicative inferences work, see Grice (1975), chapters by Wilson & Sperber and by Wilks in Myers, Brown, & McGonigle (1986), and Sperber & Wilson (1986).
15. See Grosz & Sidner (1986), and references in this paper to previous work of Grosz, for a discussion of focus modelling.

Acknowledgements

Thanks to Mark Elsom-Cook for providing the PROLOG teaching dialogue discussed in this paper, and to Françoise Langlois and Andrée Tiberghien for providing the physics teaching dialogues. The ideas presented in this paper have been developed over a number of years, and have therefore benefited from discussions with many people, in particular: Mark Elsom-Cook, Andrée Tiberghien, Michelle Joab, Jens Allwood, Cécile Paris, Laurent Karsenty and Pierre Dillenbourg. I would like to thank the anonymous reviewers for their helpful comments. The final responsibility for this work is of course entirely