

Constructive Dialogue Management: cooperation based on communicative principles

Kristiina Jokinen
Computational Linguistics Laboratory
Nara Institute of Science and Technology*

1 Introduction

When considering different dialogue systems, three approaches to dialogue management can be distinguished. *Special purpose NL interfaces* aim at providing an appropriate service with simple dialogues constrained by the requirements of a particular task. Usually the system reacts mechanically to the user's request, and the user is required to find an appropriate question strategy to get the information she wants. Many information retrieval systems (library catalogues, flight information, tourist information) are of this type. The second approach aims at building *conversing dialogue systems* which allow active dialogue partners to engage in mixed initiative dialogues. This presupposes the modelling of realistic conversational settings, and thus requires explicit representation of the speakers' beliefs, goals and plans, as well as characterisation of appropriate dialogue behaviour by conversational roles. HAM-ANS [16] and KAMP [2] were early systems in this line, PEA [14] and EDGE [4] are recent examples.

Recent advances in resource-bounded agency [5, 3], conflict resolution [8] and rational agency [7] have led to a third approach, focussing on *conversational agents*. The system models a communicating agent, albeit a simple one, which can reason about communicatively adequate behaviour. Rather than being hard-coded in the system's control structure, rationality and cooperativeness are regarded as general communicative principles which govern the whole reasoning process, and the system can thus adapt itself to different communicative situations.

This paper advocates this new approach to dialogue management, based on a *Constructive Dialogue Model*. Dialogue is understood as a joint activity between rational and cooperative agents. It is initiated to accomplish some independently set real world task which cannot be achieved by the agents acting alone: to obtain required information, the agents need to communicate with each other. Since the agents act in interactive social situations, their behaviour is constrained by obligations which describe rational and cooperative action. On the one hand, there is a need or an obligation to react (quickly) to the changing environment

(open a window when it is too hot, find food when hungry, reply when addressed, etc.). On the other hand, smoothness of interaction requires that agents should not prevent other agents from fulfilling their goals, and thus actions which show the agents' sincerity, motivation and consideration towards the other agent(s) are favoured. Hence, by means of communication, the agent constructs a model of how to achieve her goals and simultaneously take into account contextual requirements and the partners' goals. Communicative contributions are reactions to the partner's previous contribution, evaluated in the changed context with respect to the communicative obligations.

The paper is organised as follows. Section 2 describes the Constructive Dialogue Model (CDM) and its main features. Section 3 presents the system architecture. Section 4 deals with constructive planning and goal formulation. Finally, section 5 draws conclusions.

2 Constructive Dialogue Model

The *Constructive Dialogue Model (CDM)* is based on the following claims:

1. Dialogue is **cooperative negotiation** rather than a simple question-answer sequence, cf. [14, 6].
2. Participants are **rational, cooperative agents** [13].
3. The agent with a **need** initiates the dialogue [10].
4. The agents **push their own goal forward while showing consideration to the partner's goal**.
5. The agents exchange **new information**.
6. Contributions are **reactions** to the immediately previous contribution **in the changed context**.
7. Contributions are **locally planned and realised**:
 - **No predefined dialogue structure**: coherence depends on the domain organisation.
 - **No predefined dialogue acts**: communicative intentions are encoded in expressive and evocative attitudes associated with each contribution.
8. Dialogue **finishes when the agents agree** that their goals have been successfully achieved.

*I would like to thank Yuji Matsumoto for providing an excellent environment to do research and Graham Wilcock for helpful discussions on the topic.

A communication cycle consists of the analysis of the partner's contribution, its evaluation in the changed

context with regard to the agent's own knowledge and intentions, and finally, reporting of the results of the evaluation (cf. [1]). The cycle produces new information which changes the context and obliges the partner to evaluate the change with respect to her knowledge, then report back the result. Dialogue is thus managed locally and collaboratively, and it finishes when both participants report that their goals are fulfilled in the context.

Communicative principles describe rational and cooperative action, and provide a uniform basis for planning both what to say and how to say it. They are encoded as inference rules which operate on the speakers' beliefs and intentions, and on contextual facts such as the central concept of the dialogue, new information to be communicated, initiatives and dialogue expectations, see [12].

3 The system architecture

The CDM system¹ consists of a Dialogue Manager (DM), a Task Manager (TM), and their knowledge bases, as shown in Figure 1. The key resource is the Context Model, a dynamic knowledge base containing information about the mental state of the participants (beliefs, expectations, wants, intentions etc.) as well as the current dialogue and task situation (contributions, topic, new information, dialogue goals, task goals etc.). DM also has access to Communicative Principles (rules and preferences that describe rational and cooperative communication) and World Model (general knowledge about entities and their relations in the world), while TM has access to Application Model (structure of the domain). The Application Model is assumed to be subsumed by the World Model, and thus DM can use wider knowledge than TM in its reasoning about appropriate continuations.




Figure 1: The system architecture.

The Dialogue Manager operates in three phases, corresponding to the three reasoning processes that agents undertake in communication: accept and interpret the partner's goal (DGoal), choose the dialogue strategy and formulate one's own goal (CGoal) by evaluating the

partner's goal with respect to the context, and produce a reaction by specifying one's goal into a surface contribution. The Task Manager, however, tries to prove a given task goal (TGoal) in the current context. If the proof succeeds, the next goal in the goal stack is proved, but if it fails, a recovery strategy is invoked. Recovery strategies deal with task re-planning or asking help from the partner via DM.

The Task Manager and the Dialogue Manager communicate with each other by sending requests and providing answers. TM initiates DM when it thinks that the partner may possess knowledge that is needed to accomplish the given task (claim 3). DM then produces a communicatively appropriate dialogue contribution and awaits the partner's response. The response is analysed and evaluated, and the next CGoal determined. The CGoal can be either to continue the dialogue on the basis of the current dialogue situation, or to tell TM to continue its reasoning. Figure 2 shows the communication between DM and TM.

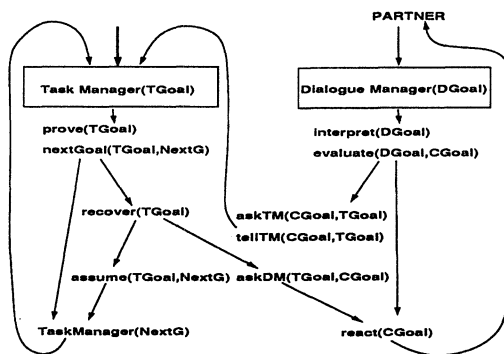


Figure 2: Dialogue Manager and Task Manager.

The difference in the operation of DM and TM is based on the fact that dialogue actions are different from actions in general. As Grosz and Sidner [9] argue, dialogue actions need to be recognised, i.e. the speaker intends the hearer to recognise her communicative plan (intentions behind the dialogue act), which is not the case in general theories of action. Moreover, general actions can be divided into subactions, the performance of which constitutes an act, while this is not the case with dialogue acts. In dialogues, several contributions may constitute an act or a contribution may contain several acts, which requires constructive evaluation of the goal rather than step-wise planning.

4 Constructive planning

A cooperative dialogue partner tries to react in the way her partner intended to evoke. However, if the system always adopts the user's evocative intentions, it can only provide simple answers to clear and unambigu-

¹The prototype is implemented in SICStus Prolog.

response	central concept	goals	speaker initiative	partner initiative
expected	related	unfulfilled	backto	follow-up-old
		fulfilled	finish/start	follow-up-new
	unrelated	unfulfilled	repeat-new, <i>X</i>	new question
		fulfilled	finish/specify	new request
non-expected	related	unfulfilled	subquestion, <i>X</i>	continue
		fulfilled	continue/start	somethingelse
	unrelated	unfulfilled	object, <i>X</i>	notrelated
		fulfilled	object/specify	new indir-request

Figure 3: Possible dialogue strategies for binary valued contextual features. *X* marks conflict situations, where the speaker must decide whether to persist on own goal or take up the partner’s goal.

ous user requests. As claimed above (claim 1), this is a simplified view of dialogues, which are joint activities between rational agents. Moreover, [8] points out that the system is *benevolent* rather than *cooperative*, if it always adopts the partner’s goal simply because the partner has the goal, and thus abstracts away from conflict situations by constraining the context in a particular way.

Given that the partner’s goal is recognised, the first step in the evaluation of the goal is to find an appropriate dialogue strategy. We call the strategy *joint purpose* as it is chosen on the basis of contextual knowledge so that it both advances the agent’s own goal and takes the partner’s goal into consideration (claim 4).² On the basis of our corpus studies³, this reasoning is found to depend on contextual features such as whether the contribution fulfils expectations, whether it is thematically related to the previous dialogue, whether the speaker has the initiative and whether the speaker has unfulfilled goals.

Two of the rules that encode the system’s rationality in determining a joint purpose can be verbally represented as follows (the italicized responses indicate the contributions planned when the constraints hold):

- (1) *If the partner’s contribution is expected and thematically related, and if the speaker has the initiative, then the speaker can go back to her previous goal and push it forward, given that the goal is not fulfilled in which case the dialogue can be closed.*

A1: I want to rent a car.
 B1: Where?
 A2: In Bolton.
 B2: *Ok. Here is a list of car hire companies in Bolton: <>*

²Dialogue participants may have goals which are not shared by their partners. However, they must have at least one common goal for the communication to take place at all, and the joint purpose encodes this.

³The corpus [15] consists of simulated human-computer dialogues where the user needs to find information from car-hire companies and restaurants in a particular area.

- (2) *If the partner’s contribution is not expected nor thematically related, and the speaker has the initiative and unfulfilled goals, then the speaker can repeat the Central Concept of the previous unfulfilled goal.*

A1: I need a car.
 B1: Do you want to buy or rent one?
 A2: Do you know any restaurants?
 B2: *Oh, I thought we were talking about cars, and wanted to know if you want to buy or rent a car.*

The current prototype system has 16 joint purposes, based on binary valued contextual features (see [11]). Figure 3 summarises the joint purposes and their contextual conditions.

After choosing the strategy, the agent formulates a communicative goal (CGoal) which she wants to achieve. Although communication itself arises from the need or wish to make the partner provide necessary information for the accomplishment of a given task, each CGoal is formulated on the basis of communicative knowledge: new information is exchanged so that communicative requirements are fulfilled (claims 5,6).

The repertoire of communicative strategies depends on the agent’s knowledge and experience (flexibility, creativity), but a greater set does not automatically guarantee a greater communicative capability. It provides a greater initial variety for responses but the agent can compensate her limited strategies by reasoning with respect to communicative obligations according to which the CGoal is formulated. Communicative obligations are shared by the agents belonging to the same communicative environment, and they appear to be rather persistent principles: even in the case of strong evidence that the partner does not share the agent’s communication model, the agent may choose to change her application and world models rather than doubt the partner’s communicative cooperation.

CGoals contain several consistent intentions. The set can be augmented with intentions which deal with explanation or compensation, or some of the main attitudes can be dropped if not necessary in the context,

see [12]. Besides enabling the agent to produce different surface realisations from the same CGoal, the filtering of CGoal via communicative principles has a further advantage in that it allows contributions to be multifunctional, i.e. convey several several communicative goals simultaneously.

Communicative contributions carry the speaker's beliefs and intentions, and when uttered in a specific utterance event, not only the conventional meaning of the contribution is conveyed but the communicative context is changed as well. Contextual changes, however, can only be estimated as expectations of the partner's reaction, so the completion of a communicative act is usually achieved via a sequence of utterances where each contribution is a reaction to the previous contribution and specifies the context to the point where the agents agree that their CGoals have been successfully achieved (claims 7,8). Hence, a communicative goal is a complex kind of intention whose different effects on the context are derived from general principles of rational agenthood and cooperative interaction. A similar view is also advocated by [7].

5 Conclusion

This paper has argued that advances in theoretical research are important even when building practical service systems. Although more research is needed to build large-scale dialogue systems that exhibit aspects of rationality in resource bounded agents, conflict resolution etc., it is obvious that the new insights are going to be an integral part of ambitious dialogue systems in future.

Constructive Dialogue Model presents a new approach to dialogue management. The central claim is that dialogues are managed locally by reacting to the changing dialogue context. Communicative principles encode the agents' cooperativeness and rationality, and provide the basis for reasoning about dialogue continuations appropriate to the changing context.

References

- [1] J. Allwood. *Linguistic Communication as Action and Cooperation*. Department of Linguistics, University of Göteborg, 1976. Gothenburg Monographs in Linguistics 2.
- [2] D. Appelt. *Planning Natural Language Utterances*. Cambridge University Press, Cambridge, 1985. Studies in Natural Language Processing.
- [3] M. Bratman, D. Israel, and M. E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
- [4] A. Cawsey. *Explanation and Interaction. The Computer Generation of Explanatory Dialogues*. The MIT Press, Cambridge, Massachusetts, 1993.
- [5] C. Cherniak. *Minimal Rationality*. The MIT Press. Cambridge, Massachusetts, 1986.
- [6] H. H. Clark and D. Wilkes-Gibbs. Referring as a collaborative process. In *Intentions In Communication*. The MIT Press. Cambridge, Massachusetts, 1990.
- [7] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 221–255. The MIT Press. Cambridge, Massachusetts, 1990.
- [8] J. R. Galliers. A theoretical framework for computer models of cooperative dialogue, acknowledging multi-agent conflict. Technical Report 172, University of Cambridge, Computer Laboratory, 1989.
- [9] B. J. Grosz and C. L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.
- [10] C. I. Guinn. *Meta-Dialogue Behaviors: Improving the Efficiency of Human-Machine Dialogue - A Computational Model of Variable Initiative and Negotiation in Collaborative Problem-Solving*. PhD thesis, Duke University, 1994.
- [11] K. Jokinen. *Response Planning in Information-Seeking Dialogues*. PhD thesis, University of Manchester Institute of Science and Technology, 1994.
- [12] K. Jokinen. Communicative principles and utterance planning. In A. I. C. Monaghan, editor, *CSNLP 1995*. 4th International Conference on the Cognitive Science of Natural Language Processing, Dublin City University, 1995.
- [13] K. Jokinen. Rational agency. In M. Fehling, editor, *Rational Agency: Concepts, Theories, Models, and Applications*, pages 89–93. Proceedings of The AAAI Fall Symposium, MIT, Boston, 1995.
- [14] J. D. Moore. *Participating in Explanatory Dialogues. Interpreting and Responding to Questions in Context*. The MIT Press, Cambridge, Massachusetts, 1995.
- [15] J. Nivre. Corpus collection and analysis. Technical Report D2.1, PLUS deliverable., (Ed.)(1992).
- [16] W. Wahlster, H. Marburger, A. Jameson, and S. Buseman. Over-answering yes-no questions: Extended responses in a NL interface to a vision system. In *Proceedings of the 8th IJCAI Conference*, pages 643–646, 1983.