

ESTABLISHING THE RELATIONSHIP BETWEEN DISCOURSE MODELS AND USER MODELS

Ethel Schuster

Department of Computer and Information Science
The Moore School
University of Pennsylvania
Philadelphia, Pennsylvania 19104

1 INTRODUCTION

Many current research efforts have focused on building cooperative systems that interact with their users in a natural language such as English. To be effective, these systems must be robust, their dialog must be coherent, and their responses must be helpful to the user. A *user model* (UM), which can be modified during the interaction to represent updated beliefs about the current user, is one mechanism that can contribute to a robust, coherent, and cooperative dialog.

In general, when we as speakers describe certain situations, we try to communicate these situations to our listeners. As proposed by some researchers (Webber 1978, Kamp 1984), speakers do so by attempting to get their listeners to construct an appropriate model: a *discourse model*. A discourse model (DM) is viewed as containing representations of entities, along with their properties and relations they participate in. The key, then, in successful communication is for the speaker to transmit as much information about those entities, their properties and relations to the listener so as to achieve the goals of the current interaction. From the point of view of a system, a computational discourse model is used by the system to generate and/or interpret a discourse.

This paper focuses on the relationship between DMs and UMs. It starts by describing what a DM is, and the role it plays in a coherent dialog. It then describes what a UM is, and the role it plays in a cooperative dialog. I argue that the DM should be viewed as one part of the UM—that is, as one part of the system's model of the user. The examples of the natural language interactions are presented in the context of a natural language interface to an expert system that provides advice on cooking with chilies. (Part of the data was taken from the section "Cooking with Chilies" that appeared in *Bon Appetit* magazine, December 1986. The expert

system can provide information about the different varieties of chili peppers as well as descriptions of how to "turn down the heat" of the chilies (make them less spicy), and how to cook with them without getting any kind of skin or eye irritations.) I justify this by showing how DMs can be viewed as part of UMs and how both models can affect each other. In other words, part of the UMs that systems have correspond to the DM, that is, a representation of what is talked about in a specific interaction. This piece, which changes with each discourse, affects the UM and varies from interaction to interaction.

2 DESCRIPTION OF DISCOURSE MODELS

A piece of discourse is a collection of utterances that are spoken by one or more speakers. Usually the sentences in a discourse are connected in a way that makes them comprehensible and coherent. One way in which sentences in a piece of discourse are connected is via the use of anaphoric expressions. In general, anaphoric expressions refer to things that have been mentioned previously in clauses. (Note that there may be cases in which the anaphoric expression may refer to an entity that will be mentioned afterwards (i.e., cataphora) rather than before it, as in the following:

- i. After he finished the race, John went drinking to celebrate his victory.

In this paper, I am concerned only with anaphoric expressions that refer to entities previously mentioned, such as:

- ii. After John finished the race, he went drinking to celebrate his victory).

In English, anaphoric pronouns contribute to coherence in the discourse by avoiding repetitions of entities already mentioned. Consider the following:

1. John went to the store and bought a pepper. He then went home to cook with it,

Copyright 1988 by the Association for Computational Linguistics. Permission to copy without fee all or part of this material is granted provided that the copies are not made for direct commercial advantage and the *CL* reference and this copyright notice are included on the first page. To copy otherwise, or to republish, requires a fee and/or specific permission.

0362-613X/88/0100●-●\$03.00

where instead of repeating "John" and "pepper" we have used the pronouns "he" and "it", respectively.

We use sentences in the discourse to describe certain situations to our listeners. This we do by attempting to get our listeners to construct an appropriate model: a discourse model. A speaker's DM enables him to generate what he believes will be coherent utterances. Similarly, the listener's DM enables him to comprehend discourse in an organized manner. Several researchers (Webber 1978, Kamp 1984, Heim 1982, Sag and Hankamer 1984) have been concerned with how DMs can be used to identify the referent of an anaphoric expression. (Not all these authors use the term "discourse model". For instance, Kamp (1981) describes the utterances as being represented in a *discourse representation structure* (DRS). The entities mentioned in the sentence are represented in the DRS and they are called *discourse referents* (DRs). Heim's (1982) framework is the File Change Semantics.) They have suggested that speaker and listener each build a model of the discourse from the incoming sentences, including representations of the entities introduced by the discourse, their properties, and the relations they participate in. When an entity is later referred to via an anaphoric expression, the discourse participants can use their DM to make the appropriate link to an entity and hence interpret that anaphoric expression correctly.

Some of the work on anaphora has concentrated in describing what characterizes the entities in the DM. For instance, Webber (1978) looked at the problem of definite noun phrases (where the references are to individuals and sets) and Schuster (1986, 1988) is looking at references to events and actions. The description of how things, sets, events, actions, facts, and so on, are represented in a discourse model and how one can refer to them gives us a clue to what characterizes a discourse model. Because the representations in the discourse model are of specific objects or events which are talked about during the interaction, the discourse model can be viewed as a *temporary* knowledge base. Since a discourse has relatively short duration, the discourse model that supports the interaction contains short term or temporary information.

It is important to note that the representations of entities, as they appear in the discourse have a structure as proposed by Grosz and Sidner (1986). While Grosz and Sidner do not specifically deal with discourse models, their view on discourse is applicable to discourse models. The discourse model reflects the structure of the dialog. In the same way that items are highlighted in the actual discourse, they appear as being more salient in the discourse model. Because some items are more salient than others, the representation is not just a flat representation, but has a hierarchical structure in which the more salient entities are represented in the same way as they appear in the discourse. The structure is needed because the ordering of the representations does not necessarily correspond to the

order in which the entities are mentioned in the discourse. A focusing mechanism plays a very important role in understanding discourse. This mechanism is needed to process sentences at any point in the discourse by indicating which objects, things, events, or facts are more salient at any point in the discourse. When processing a part of discourse, only those entities that are salient come into play.

3 A VIEW OF USER MODELS

In this paper, the UM is viewed as "the system's beliefs about its users". Many views have been proposed to describe what UMs are. The various UMs proposed so far fall under the general category described here. For example, this definition of UMs includes McCoy's (1985) concept of a UM: the system's beliefs about how the user views objects in the domain. It also includes Paris's definition of a UM: the system's beliefs about the user's levels of expertise as well as the definition of UMs as viewed by researchers concerned with plan recognition: the system's beliefs about what the user is trying to do.

Many distinctions have been made when characterizing user models. Kobsa (1985) and Kass and Finin (this issue) distinguish between *user* models and *agent* models. (Kobsa actually uses the term *Akteurmodell* (actor model) since, according to him, the primary meaning of the German *Agent* is "secret (foreign) agent".) For them, the agent model is the model of the person that the system can model and there can be many agent models. The user model is the model of the specific agent that interacts with the system. Often the agent model and the user model coincide. In this paper, I will assume that this is the case. Also, Rich (1979) distinguishes between models of individual users and models for classes of users, as well as between long-term as compared to short-term UMs. This notion of short- and long-term UMs provides a spectrum of parts of the UM, some of which are temporary and some of which remain after the discourse ends. I will show more on this issue in the next section.

In this paper, I assume that the system has representations for three possible stereotypes of users: a beginner, an intermediate, and an expert. The system can modify its own user model as the interaction occurs, as a result of the information that flows out from the discourse model into the user model. Thus the user model is dynamic. In general, information that is relevant to the user and which is represented in the DM becomes part of the UM.

Consider a simulated expert system HOT, which provides information about cooking with chilies. The system provides advice to aficionados (amateurs) about buying, cutting, peeling, storing, and cooking with chilies. The system also has a general UM from which it can identify three possible users: beginner, intermediate, and expert. These are canonical UMs, and they are

representations of three potential classes of users of the system. The beginner stereotype contains information about simple and well-known varieties of chili peppers. It also contains information about storing chilies. The stereotype for intermediate users assumes that the user knows more than a beginner, while an expert is assumed to know about unusual varieties of peppers and to be interested in more sophisticated information concerning chili peppers and detailed information about using different kinds of them.

The users interact with the system by asking questions. From these, HOT can decide how to fit each user into any of the particular UMs that it has available. Also, the sample responses from HOT are used as a way of demonstrating how the UM participates in the discourse. In other words, the responses show evidence of interaction between the UM and the DM. The following example illustrates the interaction between HOT and one of its users.

2. U: Hi! I love to eat spicy food and I love to cook with chilies. I just found some fresh peppers in the health food store called banana-peppers and I was told they are very hot. How can I peel them?

From this introduction the system can deduce that the user is an intermediate user in cooking with chilies, and invokes the stereotype for intermediate users. How does the system decide that this user is an intermediate and not a beginner? Firstly, the user explicitly mentions that he likes to eat and cook spicy food. Also, the system can realize that a more experienced person in spicy food knows about the need to peel hot chilies (sometimes), while a novice may not realize that some kinds of peppers need to be peeled. And an expert would know how to peel hot peppers. These facts trigger the intermediate stereotype in the user modeling system. Notice that the user mentions specific entities (e.g., himself, peppers, health food stores, and so on) as well as events: “user loves to eat spicy food”, “user cooks”, and so on. All these entities and event descriptions are represented in the DM and they are used to infer the correct level of the user in the UM. This fact is evidence that the DM is part of the UM. Once the system has decided that the user is an intermediate, it can respond not only in terms of what the user wants to know, but also what will be most helpful to the user.

4 RELATIONSHIP BETWEEN DM AND UM


In the previous sections I have shown the role of the DM in a user-system interaction. I have also described the role that a UM plays in a user-system interaction. The system uses the information in the UM to decide what kind of user it is interacting with, as well as how to respond to the particular user.

Given the definition of UMs in the previous section, the DM seems to clearly be part of the UM, that is, it is the system’s beliefs about what the user believes about the discourse. The question then is whether the DM is

the system’s beliefs about the discourse or is it the system’s beliefs about the user’s beliefs about the discourse. I would argue that it is the latter. It has been claimed that both dialog participants must be focused on the same subset of knowledge for communication to be successful. If the system has a DM that allows it to comprehend utterances one way and the user has a DM that causes it to interpret an utterance differently, the interaction is going to fail. So if the system is going to use its DM to generate utterances that it believes the user can understand as the system intended, then it must believe that its DM reflects the user’s beliefs about what has been talked about. (One might argue that we have to go all the way to mutual beliefs—namely, that the DM is the system’s beliefs about what is mutually believed about the discourse.)

Furthermore, if the DM were separate from the UM, then an entity introduced by the discourse could always be referred to. But that may not be possible unless the system believes the user knows about this particular entity. On the other hand, if the DM is part of the UM, then only those entities that the system believes the user knows about can be represented implicitly in the DM, since in this case the DM must represent the system’s beliefs about the user’s beliefs about the discourse. Then the system can only coherently refer to entities that it believes the user knows about, since these are the only ones represented in its DM.

In the previous section, I described a view of UMs with three stereotypes. Pictorially, this can be seen as a kernel of information with several possible levels:



The INITIAL-UM is the representation of the UM that the system has initially (before any interaction). During its interaction with the user, the system builds the DM. In turn, information taken from this DM is used to update the INITIAL-UM into an UPDATED-UM. The UPDATED-UM becomes a FINAL-UM when the interaction ends (possibly after several updates). Note that only parts of the FINAL-UM persists for future use after the current interaction ends.

All the information that the user provides is represented in the DM. Consider the following:

3. U: I want to know how to peel banana-peppers. Imagine, my mother was in Mexico and I asked her to buy some for me. She decided to try one of them and she burnt her throat. She had to be rushed to the hospital, blah, blah, blah.

This information is also part of the UM. Given the definition of the UM as the system's beliefs about the user, then this information provided by the user is the system's beliefs about the user's beliefs about what has occurred. For instance, now the system believes that the user believes that you can buy banana-peppers in Mexico.

In replying to its users, the system not only decides what information to include in the reply, but can also use anaphoric expressions (i.e. pronouns) in its responses. The only way the system could have used those pronouns was by having a representation of the discourse in which the mentioned entities were represented and available for reference. Also, since the system responded in terms of its model of the users, only if the DM is part of the UM, can the system take it into account in its responses and its reasoning about the users. Both the UM and the DM were needed in creating the response, not only because of the specific information used in the response, but also in the way in which that information was actually presented to the users. In other words, the DM is part of what the system needs to consult when responding to its users.

One of the ways to identify how the UM contains the DM is by looking for what information might be in the UM but not in the DM. In the earlier examples, the responses generated by HOT made use of information taken from the stereotype invoked for the individual user. This information was not present in (or implied by) the previous discourse. Hence the UM contains information that does not appear in the DM.

Note also that the DM can affect the rest of the UM. Suppose a user comes often in contact with the system, and keeps referring to the same things. After several interactions, these things the user mentions should eventually become part of the long-term UM. The question that is left is whether it is indeed worthwhile to perform this transfer from the DM into the long-term UM. For instance, if a user talks about the same things over a course of several interactions and the information is moved to the UM, what happens if the user stops talking about those specific things? Do we then delete

the information from the UM and allow for new information to come in? Also, with respect to the short- and long-term UMs, we could consider the short-term parts to be the DM, which is removed once it is no longer relevant. The intermediate parts could correspond to the beliefs that the system has about what the user is trying to do. And the long term would be the beliefs about the user's level of expertise, his likes, or dislikes. These are among the many issues that remain to be solved.

ACKNOWLEDGEMENTS

Many thanks go to Kathy McCoy, Sandy Carberry, Bob Kass, Julia Hirschberg, Alfred Kobsa, and Tim Finin for their discussions that helped me in understanding this controversial issue as well as for their comments on earlier drafts of this paper. I would also like to thank the participants of the 1986 User Model workshop at Maria Laach for their comments during the heated discussion that this subject generated.

REFERENCES

- Grosz, B. and Sidner, C. 1986 Attention, Intention, and the Structure of Discourse, *Computational Linguistics* 12: 175-204.
- Heim, I. 1982 *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst, MA.
- Kamp, H. 1984 A Theory of Truth and Semantic Representation. In Groenendijk, J.; Janssen, T.M.V.; and Stokoff, M. (eds.), *Truth, Interpretation and Information*, Foris Publications: 1-41.
- Kass, R. and Finin, T. 1987 Modelling the User in Natural Language Systems. *Computational Linguistics* forthcoming, 1987.
- Kobsa, Alfred 1985 *Benutzermodellierung in Dialogsystemen*. Ph.D. thesis. Springer-Verlag, Berlin—New York.
- McCoy, Kathleen 1985 *Responding to Users Misconceptions*. Ph.D. thesis, University of Pennsylvania, Philadelphia, PA.
- Rich, Elaine 1979 User Modeling Via Stereotypes. *Cognitive Science* 3: 329-345.
- Sag, I. and Hankamer, J. 1984 Toward a Theory of Anaphoric Processing. *Linguistic and Philosophy*: 325-345.
- Schuster, Ethel 1986 *Towards a Computational Model of Anaphora in Discourse: Reference to Events and Actions*. Ph.D. dissertation proposal. Technical Report MS-CIS-86-34, University of Pennsylvania, Philadelphia, PA.
- Schuster, Ethel 1988 Ph.D. thesis, University of Pennsylvania, Philadelphia, PA. Forthcoming.
- Webber, B.L. 1978 *A Formal Approach to Discourse Anaphora*. Ph.D. thesis, Harvard University, Cambridge, MA.