User Models and Discourse Models: United They Stand . . .

Alfred Kobsa

SFB 314: AI-Knowledge-Based Systems
Department of Computer Science
University of Saarbrücken, W. Germany

Opinions on the relationship between discourse models (DMs) and user models (UMs) are obviously influenced by preassumptions about their respective contents. As far as DMs are concerned, two divergent views have been expressed in the discussion published here:

1. The DM contains only representations of the objects mentioned so far in the discourse (i.e., a *mentioned-object memory*—see Schuster, this issue). The term "object" will be used here in the broad sense of Schuster, thus also denoting events, properties, etc.
2. The DM contains *in addition*
   a. a representation of the purpose underlying the segments of the dialog (i.e. a *dialog purpose*—see Grosz Sidner 1986, Chin, this issue).
   b. an *attentional structure*, which is a subset of the representations mentioned in (1) containing the currently focused objects which are ordered in a focus stack (Cohen, this issue; Chin, this issue, who requires only that the user must be familiar with these objects).

Less disagreement seems to exist about the components of a UM. Generally, it is regarded as containing explicit representations of the system's assumptions about all relevant aspects of the user, i.e., assumptions about his/her "objective situation" (e.g., marital status, number of children), as well as about his/her prior knowledge, goals, plans and false beliefs with respect to the domain of discourse. In order to meet Wahlster's personnel-database counterexample, it must be further required that the user model be separable by the system from the rest of the system's knowledge.

To discuss the relationship between DMs and UMs, a general belief, goal, and plan maintenance system (BGP-MS) will be presented here, the purpose of which is to store and update the beliefs, goals, and plans of both the system and an arbitrary number of other agents, including the system's current user. Specific subcomponents and subfunctions of this system hopefully capture the general consensus on what constitutes a discourse model and a user model, respectively. However, we will see that these subcomponents are strongly interwoven and that—apart from a few rarely occurring exceptions—the DM is part of the UM at least at the level of content. The question arises then, of course, whether it makes sense to separate these notions conceptually.

The belief, goal, and plan maintenance system outlined here is being implemented (in a somewhat simplified form) in XTRA, a natural language access system to expert systems (Allgayer et al. 1988). A previous implementation was VIE-DPM (Kobsa 1985a,b). In the knowledge base of BGP-MS, the representation of the various types of (nested) beliefs and goals (Kobsa 1988) is separated into a number of hierarchically ordered partitions (see Figure 1). If it is shared knowledge between S and U that U possesses certain beliefs (knowledge), then this knowledge or these beliefs are represented in MB(UB).[1] MB(UW) contains those goals and plans of the user, MB(SB) those beliefs of the system, and MB(SW) those goals of the system for which the same holds true. "Private" beliefs of the system about the domain of discourse / about the user's beliefs / about the user's beliefs about the system's goals are represented in SB, SBUB, and SBUBSW, respectively. MB contains the mutual beliefs (knowledge) with respect to the domain, and MW the mutual goals and plans of S and U. The arrows between the partitions denote inheritance relationships.

In the partitions of BGP-MS, the content of the individual beliefs, goals, and plans can be expressed through arbitrary representational structures (e.g., a KL-ONE-like representation as used in XTRA). Various markers for non-belief and uncertainty can be added: For instance, in SBUB it can be expressed, among other

0362-613X/ 88 /0100●-●$03.00

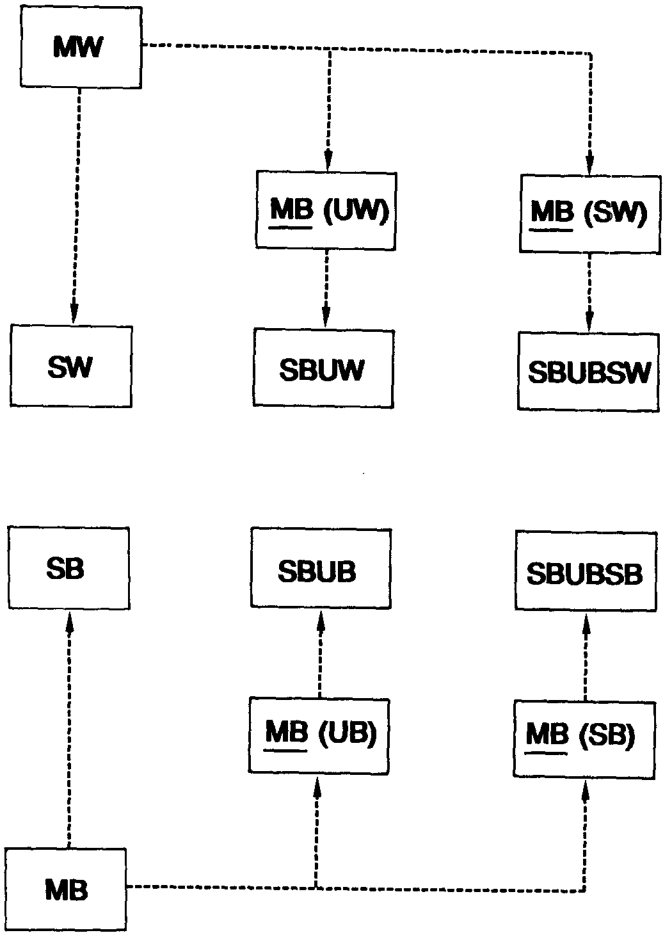

Figure 1. Hierarchical belief, goal and plan representation in BGP-MS.

things, that S is uncertain whether (or does not believe that) U knows some fact; and in MB(UB), that S is uncertain (or does not believe) that a belief of the user is mutually known.

Where are the user model and the discourse model located in this architecture? The UM part of BGP-MS consists of all partitions except SB and SW, plus all representations in SB (and probably SW) in which an individual constant occurs denoting the user (the rest of SB corresponds to Sparck Jones's world model). The DM cannot be so easily identified. In the following sections I will discuss how the different functions of a discourse model as outlined above can be fulfilled by the proposed architecture.

1. The Mentioned Object Memory

When an object is mentioned during the discourse, then mutual knowledge about its existence is usually established. Thus a representation of the object can be entered into MB (implying, for instance, that S can now use definite NPs to refer to these objects). Finer distinctions in what is now known can also be expressed: to take Chin's example, if S mentions a name unknown to U (i.e., the relationship between the name and its bearer is contained in SB only), then the existence of a person with this name can be entered into MB. In MB(UB) it can be represented that U does not know, and in MB(SB) that S does know the bearer of this name.

All of the mentioned partitions are part of the user model. One might argue that the above architecture may completely cover the mentioned object memory function in an NL dialog system. However, three kinds of information are lost thereby (this defect also seems to apply partly to Schuster's model).

a. The information that objects have been explicitly mentioned in the discourse:
MB contains not only those objects which were explicitly mentioned in the discourse (and which are therefore mutually known by all dialog participants), but also representations of those objects whose existence is mutually known due to *stereotypes* (Rich 1988) or to *inferences* from the discourse. Sometimes, however, the system should possess information about whether or not some object had been explicitly mentioned, for example in order to increase coherency in its own dialog contributions ("As I/you said before. . . ") or to point out inconsistencies in dialog contributions of the user ("But previously you said. . .").

b. Information about the sequence (and thus recency) of objects' mention:
This information is very important in NL systems, since the choice of various forms of anaphora depends on the degree of recency.

c. Information about the linguistic structure of dialog contributions:
Sometimes the system should also possess information about the wording or the syntactic structure of the user's and the system's previous dialog contributions (i.e., information on *how* objects have been mentioned). This information can be exploited by the system for reiterating a description in its own dialog contributions or for avoiding reiteration, for instance.

In XTRA, two additional knowledge bases have been introduced which serve the above-mentioned functions, among others: the FSS knowledge base and the so-called *Linguistic Dialog Memory*. The FSS (Allgayer and Reddig 1986) represents the functional semantic structure of both the user's and the system's dialog contributions. FSS contents can also be linked to the linguistic surface forms (NPs, PPs, etc.) which caused their creation (in the case of user input) or became their linguistic realizations (in the case of system dialog contributions). The dialog memory, among other things, records the objects that have been mentioned during the discourse, and, in its dialog sequence part, the sequence of the objects' mention.

In general, all system knowledge about what objects have been mentioned in the on-going discourse, in what order they were mentioned, and how they were described (i.e. all parts of the dialog memory) are regarded as being part of MB, and hence part of the user model. This is necessarily so, since, by definition, MB contains all knowledge that is shared by system and user. And only if knowledge about the previous discourse is shared between both participants can it be safely employed in the generation of dialog contributions. (For example, an anaphor generated by the system will probably fail to fulfill its referential function if only the system, but not the user, believes that its intended referent has been mentioned just recently. Hence the system should check whether its records in the dialog sequence memory are shared by the user.)

In cases of communicative failure, however, there exists system discourse knowledge that is not shared by the user, and thus not part of MB and the user model. In such cases, entries are made into SB instead, and thereby form part of the system's knowledge only. For instance, when S assumes that U does not remember what has been said (see Wahlster's hastily-presented-names example), the FSS descriptions and the representations of their referents in the dialog memory can be entered in SB instead of MB, and thus do not form part of the user model. In addition, however, all sorts of uncertain assumptions about what the user has or has not, in fact, kept in mind can be expressed in SBUB or MB(UB), i.e., in the user model. (For example, the system can note in the user model that the user probably remembered the first two but not the subsequent names.) For simplicity, however, neither of the described cases of communicative failure is dealt with in XTRA, and for implementational reasons the FSS part of MB forms a separate partition.

2A. THE DIALOG PURPOSE

Here the problem arises in research on NL dialog systems as to whether or not one should regard the intentional structure of individual utterances (Cohen, this issue) or dialog segments (Grosz Sidner 1986) as being independent of the dialog setting and the dialog participants. To put it in a more provocative way, do dialog constituents or dialog participants have an intentional structure? In my view, the essence of problem-solving dialog lies in the recognition of the dialog participants' goals and plans, and in the construction of mutually known goals and plans. Dialog contributions of the dialog partner serve as a more-or-less helpful aid in this process (Pollack et al. (1982) present transcripts in which dialog contributions of clients with misconceptions even impair the recognition of their actual goals). Apart from that, no intentional character pertains to dialog constituents that is independent of the dialog and situational context and of the current (beliefs about) goals and plans.

In the BGP-MS philosophy, (beliefs about) goals and plans are contained in those partitions whose labels include a *W*. A user's dialog contribution is first represented by FSS structures in MB. If user plans or goals can be inferred by S, they are represented in MB(UW), or in MW if there is mutual knowledge that they have been accepted by S. Conversely, when S gradually communicates its plans or goals to U, the corresponding representation structures are transferred from SW to MB(SW), and finally hopefully to MW. Thus any sort of intentional structure is part of the user model.

2B. THE ATTENTIONAL STRUCTURE

I agree with Chin (this issue) that the attentional structure is also not a context-independent characteristic of discourse (although Chin's notion of attentional structure seems to be broader than mine). Only mutually known objects can be in focus. In XTRA, focus is expressed by focus values in the dialog memory which, logically, can only be applied to representation structures in MB and MW and therefore form part of the user model.

Summary

The above discussion demonstrates that the function of a UM and of all mentioned DM components can be completely fulfilled by the outlined belief, goal, and plan maintenance system. I cannot deal here in detail with the question of whether this is also the case for other components that have been proposed for a DM, for example, the structuring of the dialog into dialog segments (Grosz Sidner 1986), a context space grammar (Reichman 1981), or rhetorical predicates (schemata; McKeown 1982). With respect to the analyzed components, we have seen that the discourse model almost completely overlaps with the user model at the level of content. Only if the user does not fully catch the system's dialog contributions are entries in the DM created which do not form part of the UM (see, for instance, Wahlster's example, this issue). But at a procedural level as well, only a few processes can be found which operate exclusively on that part of the user model that is identical with the discourse model, or upon the remaining parts of the user model.

This large degree to which the DM is included in the UM, however, is not surprising: Discourse models are ultimately based on linguistic conventions. In order for the linguistic, intentional, attentional, etc., structure of the previous discourse to be exploited for future dialog contributions, conventions about what the structure of a particular ongoing dialog actually is must exist. Knowledge about convention is mutual knowledge, however, (Lewis 1969, Schiffer 1972), and thus part of MB. The same holds true for the above-mentioned additional components of the DM that could not be dealt with in this paper. And, by the way, it also holds true for the grammar the system employs (but see the opposing views of Morik and Wahlster, this issue). If the system did not assume that its assumptions about the syntactic structure of language (as expressed in its grammar) be shared by the user, it could not justifiably use it in the analysis and generation of dialog contributions without risking miscommunication. And there definitely exists work in user modeling (e.g., Schuster 1985, Kilbury 1986, Lehman Carbonell 1988), which is concerned with the recognition of those parts of a user's idiosyncratic grammar that deviate from the mutually shared kernel grammar. Of course, an entry in MB never means that the system assumes that the user "has the same structure in his/her mind" (e.g., ATNs, KL-ONE, or LISP), but only that these structures are functionally equivalent reconstructions of the user's competence.

Does the large degree of inclusion of discourse models in user models at the level of content imply that the notion of discourse model is superflous? As was pointed out by Morik (this issue), extensionally overlapping notions may still prove useful if their intension highlights different aspects of a system. For example, in the above architecture, such a concept might characterize an orthogonal substructure and denote, for instance, entries in different partitions with specific origin or function. The above as well as Morik's and partly Wahlster's discussions demonstrate, however, that it is very hard to find such differential criteria for DMs. I therefore suspect that a happy fate of that kind will more probably apply to notions such as mentioned object memory or discourse sequence memory than to the vague notion of discourse model.

Acknowledgement

This research was supported by the German Science Foundation in its Special Collaborative Programme on AI and Knowledge-Based Systems (SFB 314). I am indebted to Carola Reddig and Norbert Reithinger for their comments on an earlier version of this paper.

References

Allgayer, J. and Reddig, C. 1986 Processing Descriptions Containing Words and Gestures: A System Architecture. In Rollinger, C. R. and Horn, W. (eds.) *GWAI-86 und 2. Österreichische Artificial-Intelligence-Tagung*. Springer, Verlag, Berlin—New York.

Allgayer, J.; Harbusch, K.; Kobsa, A.; Reddig, C.; Reithinger, N.; Schmauks, D. 1988 *XTRA: A Natural-Language Access System to Expert Systems*. Technical Report, SFB 314: AI-Knowledge-Based Systems, Department of Computer Science, University of Saarbrücken, W. Germany.

Kilbury, J. 1986 Language Variation, Parsing, and the Modelling of User's Language Variations. In *Proceedings of the 7th European Conference on Artificial Intelligence*, Brighton, England: 29–32.

Kobsa, A. 1985a *Benutzermodellierung in Dialogsystemen*. Springer-Verlag, Berlin—New York.

Kobsa, A. 1985b Using Situation Descriptions and Russellian Attitudes for Representing Beliefs and Wants. In Proceedings of the International Joint Conference on Artificial Intelligence, Los Angeles, CA: 513–515.

Kobsa, A. 1988 A Taxonomy of Beliefs and Goals for User Models in Dialog Systems. In Kobsa, A. and Wahlster, W. (eds.), *User Models in Dialog Systems*. Springer-Verlag, Berlin—New York.

Lehman, J. F. and Carbonell, J. G. 1988 Learning the User's Language: A Step Towards Automated Creation of User Models. In Kobsa, A. and Wahlster, W. (eds.), *User Models in Dialog Systems*. Springer-Verlag, Berlin—New York.

Lewis, D. K. 1969 *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA.

McKeown, K. R. 1982 *Generating Natural Language Responses to Questions about Database Structure*. TR MS-CIS-82-5, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.

Pollack, M.E.; Hirschberg, J.; and Webber, B. 1982 User Participation in the Reasoning Process of Expert Systems. MS CIS-82-9, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.

Reichman, R. 1981 *Plain Speaking: A Theory and Grammar of Spontaneous Discourse*. Report No. 4681, Bolt, Beranek and Newman, Cambridge, MA.

Rich, E. 1988 Stereotypes and User Modeling. In Kobsa, A. and Wahlster, W. (eds.), *User Models in Dialog Systems*. Springer-Verlag, Berlin—New York.

Schiffer, S. R. 1972 *Meaning*. Clarendon Press, Oxford, England.

Schuster, E. 1985 Grammars as User Models. In Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-85, Los Angeles, CA: 20–22.

Note

1. The abbreviations are mnemonic: read "system believes" for "SB", "system wants" for "SW", "user believes" for "UB", "mutual belief" for "MB", etc.