CHAPTER 4

Contextual knowledge management in discourse production*

A CDA perspective

Teun A. van Dijk
Universitat Pompeu Fabra

Introduction

One of the major contributions of psychology and A1 to the theory of discourse has been the fundamental insight that discourse production and comprehension require vast amounts of shared knowledge of the participants. Against the background of various theories about the nature of knowledge representation, it has been proposed that relevant portions of knowledge are being activated and applied in the understanding of words and sentences, the establishment of local coherence, the formation of oyeran topics or semantic macrostructures, and more generally the generation of any kind of inference, among many other aspects of discourse understanding. Since the production and comprehension of discourse about events and actions, such as stories and news reports, basically involves mental models in episodic memory, and the construction of these models also requires the application of (more or less) knowledge, we may conclude that knowledge in discourse processing is pervasive. This insight has become so obvious that it is sometimes forgotten that until the 1970s this was not a standard part of the theory language processing at all (for details, and among many other studies in the last decades, see, e.g., Britton & Graesser 1996; Clark 1996; Graesser & Bower 1990; Johnson-Laird 1983; Markman 1999; Schank & Abelson 1977; Van Dijk & Kintseh 1983; Van Oostendorp & Zwaan 1994; Wilkes 1997).

Another well-known insight in the theory of discourse is that discourse production and comprehension is context-dependent. Although in many areas of discourse studies this is nearly as trivial an observation as emphasizing the

role of knowledge, cognitive psychology has largely ignored this aspect of discourse processing. In linguistica, discourse analysis and the social sciences, the role of context is extensively discussed, but without much explicit theorizing, and thus far without a single monograph on the theory of context (see however Duranti & Goodwin 1992; Leckie-Tarry 1995). In an earlier paper (Van Dijk 1999) I proposed that the role of context in discourse processing should be accounted for in terms of mental models of the relevant dimensions of the communicative event or situation, mental models I called "context models" or simply "contexts".

In the present paper, I shall examine the interface of these two fundamental aspects of discourse processing, namely the way knowledge in discourse production and comprehension is managed as a function of context. It will be argued that contexts, defined as mental models, need a special knowledge component that represents the relevant beliefs of speakers or hearers about the knowledge of their interlocutors. In other words, language users not only need to have general "knowledge of the world", and not only knowledge about the current communicative situation, but of course also mutual knowledge about each others' knowledge. These assumptions are relevant dimensions of the current communicative situation, and hence must be accounted for in a theory of context models. In other words, how do language users actually manage the common ground of knowledge they need in order to be able to be mutually comprehensible. This question of "common knowledge" or "common ground" is hardly new in the theory of language and discourse processing (see, e.g., Clark 1996; Keysar, Barr, Balin, & Paek 1998; Planalp & Garvin-Doxas 1994), but so far has not been explicitly related to a theory of context models. In this respect this paper is intended as a contribution to a new theory of the role of knowledge in discourse processing as well as a contribution to a new theory of context.

A theory of the way knowledge is managed in discourse and interaction is also relevant for critical discourse analysis. Indeed, many of the ways power abuse operates in communication, as is the case for manipulation, involve specific knowledge strategies in discourse. In this sense this theoretical paper is also intended as a contribution to CDA.

The definition of knowledge

The theory of knowledge has been the object for thousands of years of epistemology in various cultures, and of psychology and the social sciences for many decades, and it is therefore impossible to summarize the most important re-

sults of so much reflection, theory and research (of the thousands of books on knowledge, I shall cite only the recent reader of Bernecker & Dretske 2000).

I shall therefore merely state my own position in a very long and complex debate, and basically define knowledge in terms of shared beliefs satisfying the specific (epistemic) criteria of an (epistemic) community. This very succinct definition is rather pragmatic and socio-cognitive than philosophical and abstract, and does not feature, for instance, the notion of "truth", as it is used in the traditional definition of knowledge in epistemology as "justified true beliefs". I take truth as a notion that only applies to language use, discourse or speech acts, and not to beliefs. Each community, or historical moment of a community, has its own criteria that allow members to establish that some beliefs are treated and shared as knowledge, whereas others are not. Obviously, these criteria are different in for instance scientific communities and in the "common sense" community of the public at large. One of the empirical criteria is surprisingly simple, and directly relevant in the study of discourse: A belief is treated as knowledge in a community if it is presupposed in the public discourses of that community, for instance in storytelling, songs, or news reports.

Types of knowledge

Both in discourse studies and in the psychology of text processing, we usually deal with one type of more or less abstract and general "knowledge of the world", e.g., the kind of knowledge represented in scripts or similar knowledge structures, and usually assumed to be stored in "semantic" memory. Apart from speculations about the neurological or formal aspects of knowledge representations in memory (or in the brain), there is surprisingly little explicit theorizing about the various types of knowledge. And since it is likely that different kinds of knowledge also may affect discourse processing in different ways, it is crucial to devise an explicit theory of knowledge types. Summarizing a long discussion, I therefore propose that knowledge may be typologically variable along the following criteria, for instance:

- Scope: personal, interpersonal, group, organization, nation, culture.
- Specificity: more or less general or specific knowledge.
- Concreteness: more or less abstract or concrete knowledge.
  "Reality": More or less "fictional" or knowledge about the "real" world.
- Objects: The objects of knowledge: people, animals, things, nature, etc.
- Firmness: More or less "sure" knowledge.

These and other types mix in complex ways, such that we may have, for instance, knowledge shared by the members of an organization about specific, concrete events that actually might have happened, but that might also be a company myth, but which all members nevertheless treat as "real." Much "knowledge of the world" is general, abstract and shared by members of a whole culture. It is this knowledge that is presupposed in the public discourses of that culture.

It will also be assumed that such knowledge is represented in semantic or social memory, and that personal knowledge about specific (autobiographical) events — one's personal memories or experiences — are stored in mental models in episodic memory, and that these different kinds of memories mutually influence each other (Tulving 1983). Thus, our interpersonal knowledge about specific events (such as about the dinner we had last night) may be instances of general, abstract knowledge about dinners, and vice versa, we learn about general properties of eating and dinners by generalizing and abstracting from these more detailed, ad hoc and varied instances. Beyond these elementary notions, we have little idea about the representation formats of all these different kinds of knowledge, and about where and how they are stored in the brain. In fact, we have surprisingly little solid knowledge about knowledge in general!

Context as mental model

It is fairly generally agreed upon that a sound theory of discourse should comprise not only a theory of the structures of text and talk, but also a theory of context, of the relations between text/talk and context and of (re)contextualization processes in general (Auer 1992; Duranti & Goodwin 1992; Gumperz 1982). The notion of context used in most of these approaches in the humanities and social sciences is however quite vague and intuitive, and based on the concept of a social "environment" or "situation" of language use. Such situations would involve categories such as Setting (Time, Place, etc.), Participants in various roles, Actions, and Cognitions (aims, knowledge, opinions, etc.). A context would in that case more specifically be the structure consisting of the *relevant* categories of such a situation, that is, those categories that make a difference (in the production and comprehension) of discourse structures. In other words, contexts have to do with relevance (see also Sperber & Wilson 1995).

There is one major problem with this concept of context. It lacks the cognitive interface that is able to account for subjective "relevance" in the first place: Settings, participant roles or aims of communicative events are not relevant *as*

such, but are *defined* as such by the participant themselves. That is, both the definition of the communicative situation as well as the relevante of its properties for discourse production and understanding is not only interactionally but also mentally accomplished. What may be obvious for psychologists is less so for many discourse analysts interested in context, namely that such a social context cannot possibly be "causally" related to text or talk: Social structures, participant roles, actions, time or place, etc. simply have no way to influence discourse directly, and cannot be influenced directly by discourse either. Hence we need a cognitive interface between social situations and discourse. Mental models fit that role perfectly. I therefore define a context as the mental representation of the participants about the relevant properties of the social situation in which participants interact, and produce and comprehend text or talk. This mental representation is called a "context model". Such models are stored in episodic memory, just like any kind of mental model of ongoing events and actions (for details, see Van Dijk 1999). Indeed, context models are just a special case of the kind of mental models that define all our personal experiences and that control all the situations and interactions in which we participate.

Interpreting communicative situations (and "contexts" in the traditional sense) in terms of context models has many advantages. They account for the fact that the different participants may have different interpretations and hence different models of the current situation, and these different context models will also have different effects on what they say or write or on what they understand, possibly also leading to misunderstanding and conflict. Thus context models may be seen as the crucial interface between actual discourse and the surrounding communicative situation, including the way participants represent themselves and the others as speakers and hearers.

Just like more general experience context models are not static but *dynamic,* ongoing, interpretations and representations of the current situation, That is, context models change constantly — if only because of the ongoing discourse, which dynamically changes at least what the participants know (such as the things talked about), as well as the relations between participants in interaction, as a result of what is being said. Thus, context categories influence all structures of discourse that may vary, including speech acts, rhetoric, lexical and syntactic style, and so on.

The K-device

One important category of these context models is *knowledge.* That is, it is crucial for participants that they mutually represent the knowledge of the other

participants, because many aspects of discourse depend on what the speaker assumes the hearer to know or not to know. Indeed, whenever the speaker assumes that the hearer knows something, the speaker no longer needs to assert such knowledge, but may tacitly or explicitly presuppose it, or perhaps remind it when it might have been forgotten or when it is not easily accessible, such as information about recent events. It is this knowledge component of context models that will be dealt with in the rest of this paper.

Because knowledge is such a crucial component of context models, I shall assume it has a specific status as a cognitive device, which I shall call the *K-device.* This device is permanently active "calculating" what the recipients know at each moment of a communication or interaction. This device adapts the structure of talk or text to the dynamically changing common ground of knowledge, for instance by selecting the appropriate speech act (assertions or questions), definite or indefinite articles, presupposed that-clauses, conversational markers such as "You know", reminding markers such as "as I told you yesterday" or "as we reported last week", providing explanatory details, giving accounts, and so on. In other words, if we assume, in line with current theorizing in cognitive psychology, that what participants know about an event is represented in a subjective mental model of that event (Van Dijk & Kintsch 1983; Van Oostendorp & Goldman 1999), the K-device of their context model tells participants which of such event knowledge must be asserted, which knowledge should be reminded and which knowledge can be presupposed because it is irrelevant or can be inferred by the recipients themselves.

Of course, these and other features of discourse also depend on other characteristics of the context model, such as one's intentions, the kind of people one is addressing, the nature of the interaction, the institutional setting and so on. That is, one presupposes and expresses different kinds of knowledge when speaking to children, students, one's colleagues on the job or one's spouse or friends. Indeed, expressing or presupposing knowledge not only depends on what we know that the recipients already know, but also on what we know they may *want to* know, e.g., because it is *interesting* or *relevant* for them.

That is, the K-device is related to the other characteristics of the context model. It does so reflexively in the sense that the other categories of the context model are themselves produced by the K-device. For instance, if we have a conversation with a friend, we have an ongoing context model with a participant in the social role of a friend, that is, a person we know. In order to be able to represent such information in the context model, the very K-device needs to activate and make available the now relevant knowledge about that person. Thus, what information is included in the current context model necessarily

means that participants now *know* in which setting, with whom and why they are communicating or interacting. Whereas context models are the controller of all interaction and discourse, the K-device is itself the controller of the context model. Indeed, it should even represent our knowledge of self — who we are, and as what we are now participating in the current ongoing interaction.

Since the K-device must manage a vast amount of permanently changing contextualized knowledge, it is plausible that it operates strategically, that is, by fast but imperfect operations, as we know from the strategic processing of discourse more generally (Van Dijk & Kintsch 1983). For instance, if speakers or writers must take into account what (they think) recipients know already, it would be impossible to feed all such knowledge to the K-device. Rather, there must be a fast decision strategy that says something like: *Recipient knows all I know, except X and Y,* in which case we would have one of the contextual conditions of asserting *X* and Y. Of course, several other contextual conditions must be satisfied before *X* or Y can be appropriately asserted, such as our beliefs about whether such information is relevant for the recipients or meets social conditions of politeness, but these conditions will not further be detailed here. It only needs to be emphasized that also these other conditions are of course controlled by the K-device, as argued aboye.

Since knowledge is generally defined as the shared beliefs of an epistemic community, such a strategy also has an empirical basis: If two people are members of the same epistemic community, they share, by definition, all the general knowledge of that community. On the other hand, this is not the case for all personal knowledge, much of which is not shared by others and hence must be asserted first before it can be presupposed later in the same communicative event or in next communicative events with the same recipients. Now, let us examine these and other strategies more systematically, assuming that each type of knowledge, as defined aboye, may need its own management strategies.

K-strategies

Let us examine some hypothetical strategies for different kinds of knowledge.

Personal knowledge

Personal knowledge is autobiographical knowledge about personal experiences (Neisser & Fivush 1994). The K-device assumption for all personal knowledge is that it is "private" and hence not shared by.others who did not particípate

in the relevant experiences, unless communicated. This means that if personal knowledge is presupposed, speakers need to remember that they told their interlocutor about the experience before. This means that speakers must activate a previous context model, featuring the relevant information. If they have access to such a context model, that is, if they remember they told the recipient before, then this personal knowledge need not be expressed and asserted, but it may be reminded if the corresponding context model of the recipient is probably difficult to access. It would be anomalous if the speaker were to repeat the same assertions or remind the interlocutora several times in the same communication event, unless repetition is necessary for didactic, rhetorical or other reasons. On the other hand, if the communicative event took place years ago, and the information communicated is not very relevant for the recipient, it is likely that it has been forgotten, in which case a reminder will be necessary.

Interpersonal knowledge

Interpersonal knowledge is personal knowledge that is shared by two or more individuals on the basis of previous interpersonal communication or common experiences. The strategy here is already explained aboye for personal knowledge: If speakers have access to an experience model or a context model of a communicative event in which the relevant information was shared, they may presuppose that the recipients know such information. In that case an assertion would be anomalous and a reminder necessary if the context model is old and the event model not very relevant for the recipient. The relevant context model would typically be one of storytelling in which personal experiences are told.

Group knowledge

Group knowledge is socially shared knowledge, either of group experiences, or of general, abstract knowledge acquired by the members of a group, such as a professional group, a social movement or a sect. In the first case; such experiences may be told to new members of the group (children, apprentices, novices, rookies, etc.) in various forms of "collective" stories, which may be oral or told in various kinds of literature (legends, histories, etc.) or movies reproduced by the group. In the second case, the socially shared knowledge of the group is being taught as general, abstract knowledge, for instance the kind of knowledge shared by linguists or physicians. Intragroup communication presupposes group knowledge, as is the case in scholarly articles or technical conversations. Note that our concept of knowledge is by definition relative: What is called

"knowledge" within a group, may well be called mere "beliefs" or "superstition" by members of other groups.

Institutional or organizational knowledge

Institutional knowledge is social knowledge shared by the members of an institution or organization, and in general satisfies the strategic criteria of group knowledge and discourse. Competent members of institutions or organizations may presuppose all knowledge acquired as members in the process of socialization, for instance during training or "telling the code" to newcomers. Shared organizational or institutional knowledge may itself vary between more or less official or unofficial knowledge, where the official knowledge is not only partially known to competent members but typically also recorded in institutional documents of various kinds, ranging from the minutes of a meeting to the trade secrets of a company or the "morgue" of a newspaper.

National knowledge

Natíonal knowledge is knowledge shared by the cítízens of a country. It is typically acquired at school and through the mass medía, and presupposed by all public discourse ín the country. Sínce most everyday communication for most people ís with members of the same country, most national knowledge will be presupposed in most conversations as well as in most publíc discourse. Indeed, in such communicative events it is not necessary to recall the narre of the country, the capital, the current president or . the great historical heroes of a country, or a host of other national knowledge we Iearn as citizens of a country. Little of the shared national knowledge will be personal or interpersonal. The same is true for smaller political units, such as villages and cities, although then the shared social knowledge may also feature (inter)personal knowledge, for instance about leaders or specific events.

Cultural knowledge

Cultural knowledge is the general knowledge shared by the members of the same "culture". Although the notion of "culture" is fuzzy, we shall nevertheless assume that people may identify with a culture, for instance on the (possibly combined) basis of language, religion, history, habits, origin or appearance. All discourse of competent cultural members presupposes cultural knowledge,

which is in turn acquired by all discourses of the culture, first in the family, then through schools and the media and in interaction with friends.

Cultural knowledge is the fundamental Common Ground for all other discourses and for all other kinds of knowledge, and hence presupposed by all discourses — except the didactic ones — of the culture. Most of what is traditionally called "knowledge of the world" is cultural knowledge. Cultural knowledge is usually general and abstract, and hence not about concrete social or historical events, as is the case for much national knowledge.

The general strategy for cultural knowledge is that for the large majority of intra-cultural interactions, such knowledge is supposed to be shared by the recipients. In other words, in most situations the K-device assumes that what I know is also known by the recipients, and vice versa. All other types of knowledge are supposed to include cultural knowledge.

We see that although the amount and diversity of knowledge presupposed in interaction and discourse is huge, the K-strategies of context models are fairly simple. They may be summarized as follows:

- If the recipients are believed to be members of my own epistemic community (culture, country, group, etc.), presuppose all socially shared knowledge of this epistemic community to be known by the recipient(s).
  If the recipients are believed to be members of another epistemic community, then activate knowledge about that other community. If such knowledge fails, assume that knowledge may be the same or similar to that of your own community. When in doubt, ask or otherwise show ignorance.
- If I have just acquired new knowledge, e.g., about specific events, it is probably not socially shared throughout the community, and hence not to be presupposed to be known to the recipients unless these recipients are known to have used the same source of information (e.g. the media).
- Interpersonal knowledge by definition may be presupposed to be known by the recipients with whom it was shared. In doubt, it should be reminded.
- Personal knowledge is not assumed to be shared by recipients, and should hence not be presupposed.

We see that there is a gradual transition between general cultural knowledge and specific personal knowledge, the first being virtually always presupposed to be known, the latter virtually always presupposed to be unknown to the recipients.

This means that language users may focus on a relatively simple set of alternative condition for special cases.

For personal information this means activation of previous contexts in which some personal experience may have been talked about. If such a context can be retrieved, the knowledge should be presupposed if the context was of a recent date and/or the knowledge item was interesting or relevant for the recipients; otherwise it should be reminded. Since searches in episodic memory are difficult, it will frequently happen that speakers "tell the same story twice". Such repeated partial contexts are recalled better, so that repeating the same story many times becomes less and less likely, also because recipients will comment on such repetition, and such a comment will become part of the (recalled) context model.

For new social information, i.e., any general knowledge that is not part of the common ground, or any specific knowledge about recent events, speakers or writers will not presuppose such knowledge to be known to the recipients, unless these are believed to share the same sources of information as the speaker (such as reading the newspaper, watching TV, reading the same books, etc.).

For recipients who are known to the speaker, this is rather easy, because in that case also much of their information sources may be known. In any case, when in doubt, speakers in conversation will preface any assertions with questions such as "Did you read (hear, see...) this about....?". In written communication, such doubts may be expressed in reminders, such as "as we previously reported" or "as was reported in the press last week".

Although, as we see in many communicative events, the K-strategies are rather straightforward, in the sense that the speakers or writers know what knowledge to presuppose or not, we see that the most difficult situation is one of written communication about unknown general social knowledge or unknown specific social knowledge about events.

The first is typical for all situations of learning, that is, in educational contexts, in schools and universities, or when communicating science through the mass media or other popularization situations. In both cases, if speakers (and their K-device) assume ignorance of the recipients, it needs to figure out how much the recipients (already) do know, and build the new knowledge from there by various strategies of explanation (definitions, descriptions, metaphors, comparisons, etc.) (Calsamiglia & Van Dijk 2003). Similar strategies are applied in communicative situations in which recipients are from another epistemic community (culture, country, etc.).

The second situation, however, is also quite typical of communicative events within the same epistemic community, for instance for all forms of public information, such as that of the mass media, books, and so on.

Indeed, how do, for instance, journalists or writers know what knowledge to presuppose in their writings? We have seen that for all general cultural and national knowledge, also in this situation, such knowledge will be presupposed, and that new social knowledge (e.g., about the human genome) may be partly explained. For new specific knowledge, for instance about recent news events, writers only need to know whether the recipients are likely to have used information sources that may have communicated the same knowledge. If such information is very recent (typically of the same day) and not yet reported in the media, journalists may assume that recipients do not yet know. If the information has been reported by the mass media, the journalists may assume that many people already know, but probably not all, in which case the knowledge will typically be reminded.

If we want to summarize these various conditions in one general meta-strategy, we may formulate them as follows:

> In communication with members of the same epistemic community, presuppose all shared general knowledge, as well as all specific new knowledge that has recently been communicated before.

There are other ways to formulate the same or similar general (meta-) strategies for the K-Device, such as:

> Presuppose all I knew for a long time and do not presuppose what I have just learned.

Of course, the latter strategy does not apply to didactic contexts, e.g., in education or science communication, but it will apply to most everyday situations, conversations as well as the mass media.

The point is that with the vast amounts of knowledge that must be managed, language users need to have fairly simple but efficient strategies when calculating what their recipients know already or do not yet know. This is easy for general cultural common ground, which is more or less stable and only changes gradually. It is also relatively easy for everyday interaction and personal knowledge, which only requires reactivating context models, and where various forms of checking knowledge are possible. For written communication about new public events, usually by the mass media, the easiest strategy is:

> What the media have not reported before, the recipients don't know.

Now we have a first informal idea about plausible knowledge strategies in discourse and communication, we need to be more specific about their more

precise cognitive basis. How does all this work in actual discourse processing? How does the K-device actually work in context models?

Processing assumptions

We have argued that in each communicative event participants construct and ongoingly update a mental model of the communicative situation, that is, a context model. We have also assumed that such context model construction does not occur from scratch, since context models are specific cases of ongoing experience models. This means that significant parts of the categories and contents of context models are already in place: Setting (place, time, etc.), Self (who we are, what current social identity is relevant, etc.), other participants and their roles, ongoing social actions, and so on. The same is true for the knowledge that is relevant in such experience models, from general cultural knowledge to personal knowledge shared with the other participants. In other words, the K-device of a context model is often overlapping with the K-device of the experience model(s) that precede it.

We only have rather general insights into the strategic construction of experience models in episodic memory, but it may be assumed that they are ongoingly built from combinations of (a) new perception data, (b) previous experience models, and (c) various types of socially shared knowledge. Thus, we now "know" that we are having breakfast with our partner on the basis of previous similar events, as represented in previous experience models in autobiographical (episodic) memory, on the basis of instantiated general, cultural knowledge about breakfast, on the basis of generalized personal knowledge about our partner, and finally the now relevant (self) perceptions such as the time of day (typically morning), place (say the kitchen) and other setting characteristics (table, food, etc.), the perception of my partner being present , and so on. These and other properties of the ongoing event will both trigger similar previous events in episodic memory, that is, previous experience models, as well as general cultural knowledge about breakfasts, all of which will result in the current definition of the situation as "We are having breakfast". Similar processes are at work for the definition of all social events in which we participate.

Since context models are specific forms of experience models, they basically are constructed in the same strategic way. Thus, in the breakfast example, the context model for the conversation we may have with our partner will be construed by the same general processes and constraints, that is, of ongoing

perceptions of the social situation (settings, props, participants, etc.), previous context models (earlier conversations with our partner) and general cultural knowledge about breakfast, interaction and conversation — as well as about the topics of conversation.

Note that experience models are not some kind of mental representation of all properties of a physical setting, participants or a social interaction, *but a construction of what is relevant in the ongoing situation for the (inter) actions of the participants.* Thus, in the breakfast scene it will most likely be relevant that we are in a specific place where one can have breakfast (e.g., the kitchen), that there is something to eat, and so on, but not for instance the color of the table or the precise size of a package of cereals, among a host of other "objective" properties of a breakfast scene. That is, in order to be able to successfully "do" breakfast, we only need a much reduced construct of the — for us — relevant aspects of a social situation. Since experience models are by definition personal and subjective, also these definitions of what counts as such an event (what having breakfast means for us) may be largely subjective. However, strong cultural constraints exercise control over such definitions, such that having a walk or a steak at night is not usually interpreted and described as having breakfast. In other words, experience models are subjective, but culturally speaking not arbitrary, and hence formed also by instantiations of cultural knowledge.

All this also is true for context models. Thus, a conversation at breakfast needs the construction of context models of the participants in which a large part of the experience model is already present, such as setting, participants and ongoing social action. The same is true for the relevant K-device of the context model: we already know who our recipient is, and have a pretty good idea what he or she already knows about. In other words, context models are seldom built from scratch, especially not when they are part of a more comprehensive experience model, such as a conversation during breakfast, a testimony in court or a question in a classroom. Only when major changes in ongoing events and interaction occur, context models may overlap very little with ongoing or previous experience models, as when reporters, professors or politicians engage in public discourses (a TV report, a lecture or a speech in parliament) following informal conversational or other interactions with friends, colleagues or family members. In such a case we need to construct a new setting, new participants and their relevant properties, including their relevant knowledge, and a host of other relevant social dimensions. It is also likely that such context models are already partly constructed during anticipation: public speeches, lectures or TV shows are typically engaged in by professional participants who also mentally plan and prepare such events, that is, design future context models. In a sense,

this is probably also true for more mundane situations, such as going to a shop to buy something, to ask a colleague on the job for help, or to phone our parents or children: we usually already "plan" many of these events before, that is, we already construct part of the context model before: where we go, with whom (or with what kind of person) we will speak, and so on.

In sum, context models are seldom built from scratch at the moment the communicative event begins, especially in situations in which we initiate such events. And when initiated by others, and when unexpected, then it will need fast strategic comprehension to construct the relevant model of the situation, as when a stranger asks us directions in the street. In other words, most context models are parts, or further developments of, ongoing experience models, very similar to previous context models in similar social situations, partly previously planned, or instantiations of well-known cultural knowledge, such as scripts for buying things in shops or asking a question in class. Strategically this means that at the beginning of the communicative event and ongoingly during the communicative event, a large part of the context model is already in place and only needs to be attended to or changed marginally. Thus, most of the mental resources can be dedicated to modeling the properties of the communicative situation that are changing constantly: what has (not) been said before, what recipients (now) know, how the social relationship with the recipients is developing, and so on.

One of the crucial differences between experience models and context models, and one of the reasons to distinguish them theoretically even if they have many properties in common, is that in communicative events participants not only need to monitor each others' knowledge about the current situation and the ongoing interaction, but also about (a) discourse and discourse structures, and (b) about what the discourse is *about.* We shall not further deal with the first kind of knowledge, which is (a very specific) part of more general national or cultural knowledge, and hence presupposed by the speech participants, and a general condition on verbal communication in the first place.

It is the second kind of knowledge that is the object of this paper, and that was traditionally called "knowledge of the world", and which we proposed to subject to further analysis, typology and processing assumptions. Thus, we have assumed a difference between, e.g., personal, interpersonal group, institutional, national and cultural knowledges, of which the first tend to be represented as specific, autobiographical event knowledge, that is, as mental models, in episodic memory, and the latter as more general knowledge in "social" memory. If the K-device in the context model neds to manage the regulation of

these kinds of knowledge, we must assume that together with the rest of the context model, it controls the verbal expression of event models and general knowledge and their formulation in discourse. For instance, if we want to tell an everyday story about a personal experience to our partner or a friend, as part of a conversation, the context model and the K-device will regulate which information of our mental model of that event will be selected for inclusion as propositions in the semantic representation of the story. Hypothetically, this process may be thought of as featuring the following strategic mental operations in the K-device of the context model for telling a conversational story about event E:

- Assume that the recipient has more or less the same social and cultural knowledge as I have.
- Hence assume that the recipient has more or less the same general knowledge about events such as E.
- Assume that the recipient is interested in knowing things like E.
- I have not told the same recipient about E before.

In these hypothetical conditions we recognize, not surprisingly, some of the usual appropriateness conditions of the speech act of assertions. When these conditions apply, that is, when the speaker or writer makes these assumptions, then only the relevant new information of the mental model of the event will be selected for inclusion in the semantic representation.

We see that in the process of discourse production the context model will always be constructed first, as may be expected from the theory that takes context models as special cases of ongoing experience models. It is the context model that is the result of the analysis of the communicative and interactional situation, including the conclusion that it is now possible or necessary to tell a story about E. This involves the previous or current activation of M(E), the model of E, and then a process of context-dependent selection of information in M(E) that satisfies the contextual criteria represented in the context model, including the previous knowledge assumed to be known by the recipients.

The knowledge of the recipients about E is, however, not the same during the whole communicative event: by interpreting and understanding the story about E, the recipients construct their own mental model of E, implying gradually increasing knowledge about E. The context model of the speaker, including a model of the text, strategically keeps track of what has been asserted, so that the K-device may be constantly updated with the assumed new knowledge of the recipients about E. Although these K-strategies of the speakers are fast, they are not perfect, and speakers therefore may occasionally forget what they have

told already, and hence repeat what they already have said, especially in long discourses. The theoretical ideal-case strategy is, though, that each proposition asserted by the speaker is added to the model the speaker has about the knowledge of the recipient.

Knowledge management in CDA

Let me finally sketch some of the possible application of the theoretical proposals made aboye in the area of the critical study of discourse. These suggestions are made within the general framework of an approach to CDA that does not only analyze the social conditions and consequences of discourse, but also the sociocognitive ones. The main arguments for this orientation are, firstly, that cognition is a necessary interface between society and discourse, and secondly that the cognitive structures we deal with are ***at the same time*** social, as is the case for knowledge, attitudes, ideologies, norms and values. Indeed, these social cognitions are primarily defined in terms of the beliefs shared by members of groups and communities. It is also within this perspective that we defined knowledge not as personal beliefs, but as social beliefs certified, shared and hence discursively presupposed by the members of epistemic communities. The conditions formulated aboye for the management of knowledge of discourse are in that sense social in a double sense, namely as beliefs that are socially shared, on the one hand, and because they are managed by context models that are representations of communicative situations, on the other hand. We have seen at the same time, however, that these mental models are at the same time personal, even when also socially based, because they must of course integrate the individual personal experiences, aims and interests of language users. It is also in this sense that language is inherently social, but of course used with individual variation in concrete social situation. Only in this way can we ***both*** explain the social, cultural and political dimensions of discourse, ***and*** the unique, individual variation of each specific instance of text and talk. That is, a sociocognitive approach to discourse offers a unique and necessary interface between the macro aspects of society, and the micro aspects of discourse and interaction.

Such a theoretical framework is also crucial for a more critical perspective. CDA specifically deals with the study of the discursive reproduction of power abuse, with forms of domination and social inequality. This also means that CDA needs to make explicit the way socially shared beliefs are discursively reproduced and how such beliefs are abused in the maintenance and legitimation

of domination. This is not only true for ideologies, as I have shown elsewhere (Van Dijk 1998), but also needs to be examined for knowledge. If knowledge is defined as socially certified, shared beliefs of a community, it is obvious that those groups or institutions who have preferential access to public discourse, such as that of the media, or other forms of power and authority, such as politicians, professors or priests, are in an excellent position to influence people's knowledge formation. The contextual strategies of knowledge management discussed aboye need therefore also to be examined in such a more critical social perspective. Indeed, it is quite common in CDA to talk about the discursive manipulation of the audience, but as is the case for many of such critical notions, they are hardly defined and analyzed in a rigorous way. Manipulation, thus, means manipulation of the mind, and hence needs to be examined (also) in sociocognitive terms.

CDA is specifically interested in the power and dominante of the symbolic elites, those who have special access to public discourse. Let us therefore limit this brief discussion to the contextual knowledge strategies for such discourse, and ignore personal and interpersonal knowledge management and control.

We have seen that the K-strategies for public discourse are based on the assumption that knowledge may also be defined as the beliefs that are presupposed by such public knowledge. Indeed, politicians, journalists and professors assume to be known by their audience what they do not assert. Such an assumption, however, has important social conditions, consequences and biases. First of all, it is based on ideas about the knowledge of "average", more or less well-educated citizens, and ignores vast segments of the audience that do not have such knowledge, thus excluding them from adequate comprehension of public discourse. Secondly, the symbolic elites may presuppose as knowledge beliefs that are not — certified, accepted — knowledge at all, but opinions or prejudices. By not explicitly asserting such beliefs, but just pretending that such beliefs are generally accepted, they may manipulate many readers into accepting such biased, ideologically based beliefs as certified knowledge of the community. We might call *presumptions* such presupposed beliefs that are in fact ideological assumptions and not knowledge. Typical examples, in much current elite discourse, is the association of immigration and delinquency, and that terrorism is the major problem of today's world. Since people learn by the acquisition of new knowledge through public discourse, they may thus be manipulated into believing that such presumptions of authoritative sources are in fact forms of knowledge that no longer need to be certified (demonstrated, proved, etc.).

Also the opposite happens: Elite speakers may presuppose that their audience does *not* share in general knowledge, and thus in fact treat them as being ignorant. This may not only be generally a feature of many TV programs, but more specifically when white male elites (politicians, journalists, professors, etc.) address women or specific minorities. Indeed, infravaloration is a well-known form of sexism and racism, consistent with the general polarized pattern of positive self-images and negative other-images between ingroups and outgroups. In other words, discourse that is *too* explicit, may in some cases be a manifestation of class, gender or race domination. It is therefore crucial in all forms of communication that the contextual K-strategies are finely tuned to the actual knowledge shared by the audience.

Third, another form of what could be called the abuse of the contextual K-device is the assumption in much elite discourse that knowledge is *only* conveyed by elite discourse. We have seen that the media generally use the overall strategy that what they have reported is assumed to be true, and what has not been reported is not shared knowledge. Due to the vast influence of the media in contemporary society, such a presumption may not be entirely false in actual fact, but obviously people also learn from everyday interaction, experiences, and forms of alternative discourse and communication that are not controlled by the media — such as the internet today. In this case presuming ignorance of the audience is not just a question of infravaloration, but rather an example of the corresponding process of supravaloration of the self as the only instance of "truth". A characteristic example is that if the media do not widely report about racism in European society, including of course in the media itself, then racism does not exist. If mentioned at all, it will then typically appear between quotes, that is, as a mere belief, as an accusation of an outgroup (NGO, antiracists, minorities, an international institution, etc.). Indeed, it is in this way that certain kinds of facts are *not* routinely presupposed in public discourse, typically those that have to do with the power abuse of the elites themselves. We see that in this way, all the details of the strategies of the K-device discussed aboye may be examined further for specific social and critical implications. The same is true for the manifold complications in the way knowledge is managed in multicultural communication, interaction and discourse — where "western" beliefs are not only taken to be knowledge in the speaker's own cultural community, but in fact as universal knowledge. Much more work along these lines, combining sociocognitive strategies of knowledge management in discourse, and critical analyses of the social and cultural conditions of such strategies, is still on the agenda.

An example

To illustrate these contextual K-strategies, let us examine a few examples. Consider for instance the following beginning of a routine news report in the *New York Times* (June 16, 2003; see the Appendix for the full text):

> Deal Seems Near on Israeli Pullout From North Gaza
> By Greg Myre
>
> JERUSALEM, June 15 — Israel and the Palestinians appeared today to be edging toward an agreement that would remove Israeli troops from the northern Gaza Strip, the scene of repeated confrontations, and replace them with Palestinian security forces.
>
> Visiting American and Egyptian delegations were trying to broker the deal with the larger goal of moving ahead on an international peace plan. However, violence persisted today, with Israeli troops shooting one Palestinian militant to death in northern Gaza, and Palestinians firing several rockets at Israeli towns.

In order to write this news article, the reporter Greg Myre in Jerusalem first constructs a context model for his writing a news report for the NYT, featuring the following schema categories, among others:

Macro-context:

—Global Domain: Media
—Global Action: Informing the public
—Global Participants: NYT, NYT readers

Micro-context

—Setting

— Date: June 16
—Location: Jerusalem, Israel

Participants

—Writer: Greg Myre: Role: Foreign Correspondent NYT in Jerusalem
—Reader of NYT

Note that several of the context categories are explicitly, and deictically, expressed at the beginning of the report, such as the correspondent's name, as well as his Location, whereas the date is presupposed to be known by the read-

ers of today's newspapers (and printed on top of the page). In other words, a relevant part of the context model for news reports (the current date) is already assumed to be part of the recipients' context model, but for instance not the identity and location of the current speaker/writer. We shall not further pursue these other contextual properties of news report production and reading, and focus on the management of knowledge. However, we see that knowledge is not only presupposed or expressed about the events talked about in the news, such as the Israeli pull-out from Gaza, but also about the communicative situation and hence about context models themselves. The correspondent knows that the readers of his article know that they are reading (a report in) in the NYT and that this is a newspaper. And he supplies the contextual information that most readers do not know: the narre and the location of the NYT correspondent in Jerusalem. Note that such understandings also presuppose the activation of relevant general (geographical) knowledge, namely that Jerusalem is a city and that it is located in Israel. Readers are also assumed to know of course that news reports are usually written by journalists, and that journalists located in foreign cities are usually the newspaper's foreign correspondents in that city or country.

Now, when writing this fragment, what knowledge does the NYT correspondent presuppose, remind or assert? The headline ("Deal seems near on Israeli pullout from North Gaza"), presupposes the following general (geographical) knowledge propositions, among others:

Israel is a country
Gaza is part of Palestine

— Gaza is a region.

Similarly, the headline presupposes part of a mental model of a historical event or situation, namely that Israel has occupied (North) Gaza. Presupposed but reminded for those who have forgotten or obliquely asserted for those who did not know is the fact that there have been plans for a pullout. That is, in that case the only new information is that the pullout is imminent. All other information that enable the correspondent to write this headline is presupposed general (geographical) and specific historical knowledge (the Israeli occupation), and perhaps recent specific event information (there have been plans for a pullout).

Let us now examine the first lines of the body of the news report, and list first the various kinds of knowledge presupposed by the K-device of the correspondent:

General, sociocultural knowledge

Israel is a country in the Middle East.
Palestinians are the people of Palestine.
Gaza is a region that is part of Palestine.
— Northern Gaza borders with Israel.
- Most countries have military troops.
- Countries or peoples can make agreements or a deal.
Difficult agreements or deals between countries may be witnessed or brokered by (delegations of) other countries.
- Egypt is a country in the Middle East, bordering with Gaza and Israel.
- (The United States of) America is a country.

Specific sociopolitical, historical knowledge

- Parts of Palestine are occupied by Israeli troops.
- Gaza is occupied by Israeli troops.
- Palestinians have been resisting the occupation (Intifada).
- Egypt and the US have been involved in early attempt to mediate between Israel and Palestine.
- There is an international peace plan.
- There is violence in the relationship between Israel and Palestine.

Explicitly reminded specific sociopolitical knowledge

— The Gaza strip is the scene of repeated confrontations.

New information/knowledge

- Israel and Palestinians appeared today to be edging toward an agreement.
The agreement would remove Israeli troops from the northern Gaza strip.
- The Israeli troops would be substituted by Palestinian security forces.
- Visiting delegations...broker the deal.
- They share the goal to move ahead with the international peace plan.
Violence persisted today.
- Israeli troops shot one Palestinian militant in Northern Gaza.
- Palestinians were firing several rockets at an Israeli town.

We see that the actual "news" consists of only a few propositions, and that much of the news report consists of expressions whose understanding presup-

poses various kinds of general (geographical, historical, sociopolitical) knowledge, as well as more specific knowledge about recent events. It is interesting that the news article also features a "reminded" piece of recent event knowledge, namely a description of Gaza as the scene of repeated confrontations (between Israelis and Palestinians). Note that even the new events are not totally new in the sense that they appear as referring to events that are instantiations of a fairly general script, of the following kind:

- Country A occupies country B.
  B resists
  — A reacts to B's resistance, etc.
- A and B want to make peace
  — Countries C and D are mediators.
- A and B stop hostilities
— A withdraws from B.

This general script may of course feature various sub-scripts. Thus, resisting occupation may be executed by various kinds of military or militant actions (such as firing rockets), and reactions to resistance may consist of killing militants.

The contextual knowledge management of the reporter thus consists in assuming that the readers of the NYT of course have very general, sociocultural information (for instante that countries have troops), specific political-historical knowledge (about the occupation of Palestine by Israel and about Palestinian resistance) as well as knowledge about "recent developments" of such historical events, such as the current peace plan. The correspondent is able to make these assumptions because he knows that the NYT has been regularly reporting on these events before, so that the writer may assume that readers have updated their (complex) mental model of the Israel-Palestinian conflict accordingly. On the other hand, more detailed knowledge about the conflict, namely that Gaza was the scene of earlier confrontations may not be known to the readers, and must hence be reminded, by an obligue assertion in a relative clause. Finally, the correspondent may assume that the readers cannot know about the most recent developments (e.g. of today), such as the contents of a new plan, and recent fighting, since such information has not be supplied by the newspaper, and hence is to be dealt with as "new", and hence to be asserted. We see that also the general metastrategy applies here: the reporter treats as new what he could not know himself a day earlier, or about which he had not reported before.

Apart from the different kinds of knowledge being expressed and presupposed by the reporter, note that also the "modality" of the knowledge may thus be managed. Indeed, the correspondent does not plainly assert that there is an agreement on withdrawal, but rather, quite cautiously, that Israel and the Palestinians "appear to be edging toward an agreement". As suggested before, one of the modalities of knowledge is that it may be more or less sure, and such a modality is typically expressed in various kinds of hedgings ("appear", "edging toward").

The rest of the article basically follows the same pattern of mixing presupposed knowledge of various kinds, with new information. The latter largely concerns the following information:

- Details of the preparation of the agreement, and actors involved.
- Details of the attacks mentioned (time, place, kind of attack, victims, etc.).
- Declarations of several people, from president Bush to local leaders.

There are also some typical forms or reminding of "old" knowledge, such as the recent summit meeting on June 4, and the year (2000) that the current violence started, that is, the beginning of the Intifada, and finally a brief summary of the most recent events (recent violence). Note that these various kinds of knowledge also are related to different schematic categories of news discourse, such as Current Events, Previous Events, Historical Background, Context, Verbal Reactions, and Expectations (Van Dijk 1988). Of the vast amounts of old knowledge presupposed by the news report, also observe the pragmatic criterion of the *relevance* for the readers in the US, such as the death of the cousin of the US ambassador to Israel, as an added detail about the previous violence (such as a Palestinian suicide attack on an Israeli bus). That much of the current news events are themselves communicative events (declarations of politicians involved) shows also how the news is routinely gathered, namely by seeking comments from various important people on current events.

Finally, as suggested in the theoretical section aboye, such analyses of news may take a more critical perspective. We already have seen that some K-assumptions of this article are not only biased towards the knowledge shared by US citizens — which would of course be normal for a US newspaper — but also by citizens who have a minimum of political and geographical knowledge. There are however some other, more interesting critical dimensions to the analysis. Thus when in the first paragraph the NYT describes the Gaza strip as "the scene of repeated confrontations", it presupposes or obliquely asserts (reminds) that it is true that what happened in the Gaza strip is "confrontations", namely between Israeli troops and Palestinians. However, from another

perspective the description of the events in terms of the nominalized expression "confrontation" might be seen as a euphemism in which especially the active role of the occupying Israeli troops — killing Palestinian citizens — is underemphasized, if not hidden. In other words, presupposed knowledge may not only take the form of (false) presumptions, as shown aboye, but also as presumptive forms of denying or hiding facts by euphemistic descriptions. Similarly, the new knowledge reported about "violence persisted today", in the second paragraph, mentions violence by both sides of the conflict, and is thus balanced. However, note that the events are described by one general category of "violence", whereas others may want to described the same "facts" as killings by occupying troops, on the one hand, and in terms of acts of resistance by those who are occupied, on the other hand. In other words, as we have seen, knowledge is essentially relative, and what is "known" by one (national) community, such as the USA or Israel on the one hand, may be a form of biased representation of the "facts" from the perspective of Palestinians or Arabs, on the other hand. In other words, an analysis of the contextual K-device strategies used by journalists, such as this one, should carefully and critically examine not only what beliefs are taken for granted as knowledge, but also how this is done. The "same" facts may be described as different "truths" from different perspectives, and by different communities. This again shows that knowledge and other forms of social cognition, such as attitudes and ideologies, are closely related.

Conclusions

A sound theory of text processing needs both a powerful knowledge component as well as a theory of context processing. Thus, standard theories of the role of "world knowledge" in discourse processing should be further refined for different kinds of knowledge. Unlike the definition of knowledge in epistemology as "justified true beliefs", I use a more pragmatic, socio-cognitive definition: the shared beliefs of a knowledge community that are based on the knowledge criteria of that community. Secondly, many aspects of discourse processing, such as speech acts, politeness phenomena, style, rhetoric, deictic expressions and many more, are controlled by context. However, there is no direct link between social context and text, and hence we need a cognitive interface in the form of subjective mental models ongoingly constructed by the participants of the current communicative events: context models.

One of the central properties of such context models is the knowledge of language users about the knowledge of the recipient. This K-device is crucial

for the control of many important aspects of discourse, such as what information is explicitly expressed and asserted, which information is reminded, and what information is presupposed. Since however it is impossible that such a K-device includes all the assumed knowledge of the recipients, speakers need to make use of fast and flexible strategies.

In this paper we discuss a number of these strategies, and show how for various kinds of knowledge (personal, interpersonal, group, institutional, national or cultural) different kinds of presuppositions are managed in discourse production and comprehension. Although the various strategies allow speakers to make more specific assumptions about what recipients know, the overall (meta) strategies are surprisingly simple, such as: When the recipients are members of my community, assume that they know all that I know, except the information about recent personal experiences or sources not (yet) used by the recipients. This will account both for everyday storytelling as well as for news in the press. For special cases, e.g. of new knowledge acquisition in learning and science communication, there are more specific strategies, which, however, also presuppose a common ground of shared general, sociocultural knowledge.

All these K-device strategies may also be critically examined from a CDA perspective, for instance in order to study various forms of manipulation. Thus, symbolic elites may impose their own beliefs as generally accepted knowledge, marginalize large audience segments by presupposing knowledge that is not generally known, or conversely by infravalorating non-dominant groups as ignorant.

These general theoretical assumptions about the role of the K-device in context models were finally illustrated by a partial analysis of a routine news report in *The New York Times* about a recent event in the Middle-East conflict. In that example we see how the structure of the discourse, and hence its production and comprehension is controlled by a context model in which the correspondent of the NYT in Jerusalem makes assumptions about what the readers of the NYT do know, could know, could have forgotten and do not yet know. These different ways of presupposing and asserting knowledge also relate to the standard categories of news reports, such as Headlines, Current Events, Previous Events, Historical Background, Context and Expectations. Thus, the information expressed in the Previous Events category is information earlier reported by the press, and hence in principle part of the event models of the readers, but possibly forgotten, and hence in need of a reminder. Besides this specific knowledge about specific historical events, most of the knowledge presupposed in the news report is general geogg@phical knowledge (about the

Middle-East), and specific historical-political knowledge (about the Middle-East conflict). New knowledge that provides the very basis of news pertains to current events, occurring today or yesterday, about which the newspaper has not yet reported, and which the readers are not likely to know, unless from earlier reports on the radio or TV.

A CDA perspective on such a K-analysis of news also shows that what are presupposed truths for one epistemic community, nation or newspaper, may be at best a euphemistic, incomplete or otherwise biased "version" of the facts from another perspective.

We thus see that our theory of contextual knowledge management not only accounts for many aspects of text processing, but also explains several structures of important genres such as news reports.

Note

This is a new version of a paper read at the conference of the Society for Text and Discourse, Madrid, 26-28 July, 2003.

References

Auer, P. (Ed.). (1992). *The contextualization of language.* Amsterdam et al.: John Benjamins.

Bernecker, S. & Dretske, F. I. (Eds.). (2000). *Knowledge: Readings in contemporary epistemology.* Oxford: Oxford University Press.

Britton, B. K. & Graesser, A. C. (Eds.). (1996). *Models of understanding text.* Mahwah, NJ: Erlbaum.

Calsamiglia, H. & Van Dijk, T. A. (2004). "Popularization discourse and knowledge about the genome". *Discourse Society,* 15(4), 369-389.

Clark, H. H. (1996). *Using Language.* Cambridge, UK: Cambridge University Press.

Duranti, A. & Goodwin, C. (Eds.). (1992). *Rethinking Context: Language as an Interactive Phenomenon.* Cambridge: Cambridge University Press.

Graesser, A. C. & Bower, G. H. (Eds.). (1990). *Inferences and text comprehension. The psychology of learning and motivation,* Vol. 25. New York: Academic Press.

Gumperz, J. J. (1982). *Language and social identity.* Cambridge, UK and New York: Cambridge University Press.

Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference and consciousness.* Cambridge: Cambridge University Press.

Keysar, B., Barr, D. J., Balin, J. A., & Paek, T. S. (1998). "Definite reference and mutual knowledge: Process models of common ground in comprehension". *Journal of Memory & Language,* 39(1), 1-20.

Leckie-Tarry, H. (1995). *Language & Context. A functional linguistic theory of register.* Edited by David Birch. London: Pinter.

Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Erlbaum.

Neisser, U. & Fivush, R. (Eds.). (1994). *The Remembering self. Construction and accuracy in the self-narrative.* Cambridge: Cambridge University Press.

Planalp, S. & Garvin-Doxas, K. (1994). "Using mutual knowledge in conversation: Friends as experts on each other". In S. Duck (Ed.), *Understanding relationship processes IV: The dynamics of relationships* (pp. 1-26). Newbury Park, CA: Sage.

Schank, R. C. & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Sperber, D. & Wilson, D. (1995). *Relevance: Communication and cognition.* Cambridge, MA: Blackwell Publishers.

Tulving, E. (1983). *Elements of episodic memory.* Oxford: Oxford University Press.

Van Dijk, T. A. (1988). *News as discourse.* Hillsdale, NJ: Erlbaum.

Van Dijk, T. A. (1998). *Ideology: A multidisciplinary approach.* London, UK: Sage Publications.

Van Dijk, T. A. (1999). "Context models in discourse processing" In H. van Oostendorp & S. R. Goldman (Eds.), *The construction of mental representations during reading* (pp. 123-148). Mahwah, NJ: Lawrence Erlbaum Associates.

Van Dijk, T. A. & Kintsch, W. (1983). *Strategies of discourse comprehension.* New York: Academic Press.

Van Oostendorp, H. & Goldman, S. R. (Eds.). (1999). *The construction of mental representations during reading.* Mahwah, NJ: Lawrence Erlbaum.

Van Oostendorp, H. & Zwaan, R. A. (1994). *Naturalistic text comprehension.* Norwood, NJ: Ablex.

Wilkes, A. L. (1997). *Knowledge in minds. Individual and collective processes in cognition.* Hove: Psychology Press.

Appendix

New York Times, June 16, 2003
Deal Seems Near on Israeli Pullout From North Gaza
By Greg Myre

JERUSALEM, June 15 — Israel and the Palestinians appeared today to be edging toward an agreement that would remove Israeli troops from the northern Gaza Strip, the scene of repeated confrontations, and replace them with Palestinian security forces.

Visiting American and Egyptian delegations were trying to broker the deal with the larger goal of moving ahead on an international peace plan. However, violence persisted today, with Israeli troops shooting one Palestinian militant to death in northern Gaza, and Palestinians firing several rockets at Israeli towns.

"I'm taking this as a serious proposal from the United States," said Ghassan Khatib, a Palestinian cabinet minister, refering to the Gaza security plan. "We believe the American administration can deliver when it wants to."

On June 4, the Israelis and Palestinians held a summit meeting, their most ambitious peace effort since the current period of attacks and counterattacks began in September 2000.

But last week brought a surge of violence that left more than 50 people dead, and was accompanied by angry promises from both sides of still greater violence to come. In the past two days, pressure from the United States and others has induced the warring parties to restart talks.

Earlier today, the Israeli and Palestinian cabinets held separate sessions in which they endorsed the basic principie of having the Palestinians police part of the Gaza Strip, while setting several conditions. Israelí news reports said the Palestinian security chief, Muhammad Dahlan, met for the second straight night with senior Israeli security officials about details of the plan.

The Israeli prime minister, Ariel Sharon, said his government would continue to pursue members of the Islamic group Hamas and other Palestinian militants if they were planning to strike Israel.

President Bush, speaking in Kennebunkport, Me., said, "The free world and those who love freedom and peace must deal harshly with Hamas and the killers."

Senator Richard G. Lugar, Republican of Indiana and chairman of the Senate Foreign Relations Committee, suggested that American forces might be needed to fight Hamas.

In a television interview on Sunday, he said, "Clearly, if force is required, ultimately to rout out terrorism, it is possible that there will be an American participation."

Hamas, which has always opposed peace talks with Israel, has rejected the Mideast peace plan.

The group carried out a suicide bombing on Wednesday that killed 17 people on a Jerusalem bus, including Anna Orgal, 55, identified as the cousin of Daniel C. Kurtzer, the American ambassador to Israel.

Hamas is facing pressure to suspend attacks, and group leaders today joined other Palestinian factions in discussions with an Egyptian delegation seeking to negotiate a truce.

Mr. Sharon said if the Palestinian leadership could persuade Hamas and other militant groups to agree to a truce, it would be welcomed.

"We will hold our fire, except in cases of self-defense against ticking bombs," Mr. Sharon was quoted as saying by a cabinet official.

The Palestinians said that they were prepared to take over security in northern Gaza, and that they wanted the Israeli forces to leave other areas in the coastal strip, as well as the West Bank town of Bethlehem. They also want American guarantees that Israel will not reoccupy arcas it vacates, and that it will stop the targeted killings of militants.

"We don't want a random Israeli withdrawal," said Nabil Amr, the Palestinian information minister. "It should be based on a political vision."

The head of the visiting American delegation, Assistant Secretary of State John S. Wolf, met with Israeli officials today and planned to see Mr. Sharon on Monday.

After that, he will meet with Palestinian leaders. The Palestinian prime minister, Mahmoud Abbas, was expected to travel Monday from his West Bank headquarters in Ramallah to Gaza, to renew the dialogue on a cease-fire.

Before dawn today, Israeli forces entered the northern Gaza town of Beit Hanun, the site of frequent Palestinian rocket launchings at Israel.

The troops carne under heavy fire and shot back, the military said. Palestinians said a militant was killed, identifying him as a member of Al Aksa Martyrs Brigades, a group with ties to the Fatah movement of Yasir Arafat, the Palestinian leader.

Later in the day, Palestinians fired several rockets from northern and southern Gaza at Israeli towns, but caused no damage or injuries.

The Israeli Army also said it arrested three wanted Palestinian militants in a cave outside the West Bank town of Bethlehem. One of the men, Essa Batat, was the local leader of Islamic Jihad and was linked to attacks that killed six Israelis, according to the army.

In another development, Peace Now, an Israeli group that monitors settlements, said Jewish settlers had established five outposts in the West Bank over the past week, though the peace plan calls for recently erected settlements to be dismantled.

"The settlers are nervous, and things are much more tense," said Dror Etkes, a Peace Now official. "But it's the same old story. The construction is still going on."

The Israeli military demolished 10 uninhabited settlements last week, and planned to remove five that had a small number of residents. The settlers have challenged the plan in court, and no action has been taken.

The army took down one additional outpost today, removing a bus that had been fashioned into living quarters on a hill south of Hebron, in the West Bank, witnesses said. No one was living there.