

Some observations on the role of knowledge in discourse processing¹

Algunas observaciones sobre el papel del conocimiento en el procesamiento del discurso

Teun A. van Dijk²

Abstract

In this paper, I discuss some aspects of one of the thorniest problems of the theory of discourse processing: The role of knowledge in the production and understanding of discourse. Although there are a large number of articles and books on this topic, some very fundamental questions have as yet hardly been asked. Firstly, I review the scant references to the role of knowledge in discourse processing and then I outline a tentative typology around four groups of knowledge. I then address the key question of how these different types of knowledge play a role in discourse processing, focusing on some hypothetical processes involved in the production of discourse. I finish with a call for an explicit theory of text processing that integrates context models.

Key words: discourse processing, knowledge and discourse, types of knowledge, mental models.

Resumen

En este artículo, discuto algunos aspectos de uno de los problemas más peliagudos de la teoría del procesamiento del discurso: el papel del conocimiento en la producción y la comprensión del discurso. Si bien hay una gran cantidad de artículos y libros sobre este tema, algunas preguntas fundamentales todavía no han sido formuladas. En primer lugar, reviso las escasas referencias al papel del conocimiento en el procesamiento del discurso y luego describo una tipología tentativa en torno a cuatro grupos de conocimiento. Luego abordo la pregunta clave de cómo estos diferentes tipos de conocimiento juegan un papel en el procesamiento del discurso, centrándome en algunos procesos hipotéticos involucrados en la producción del discurso. Terminó con una llamada a la formulación de una teoría explícita del procesamiento de textos que integre modelos de contexto.

Palabras clave: procesamiento del discurso, conocimiento y discurso, tipos de conocimiento, modelos mentales.

¹ Versión original en inglés, ofrecida como conferencia en Seúl, Corea; posteriormente publicada en castellano en Parodi, G. (ed.) (2002). *Lingüística e interdisciplinariedad: Desafíos del nuevo milenio*. Valparaíso: Ediciones Universitarias.

² Universidad Pompeu Fabra, Barcelona, España. E-mail: vandijk@discourses.org

1. The problem

At least since the Ph.D. thesis of Eugene Charniak in 1972 we know that people need large amounts of knowledge in order to produce or understand even the simplest of children stories (Charniak, 1972). Similarly, somewhat later in the 1970s Artificial Intelligence (AI) researchers and psychologists provided the first ideas about representation formats of knowledge, for instance in terms of scripts or similar schematic structures (Schank & Abelson, 1977).

In addition, the theory of mental models has taught us how specific event knowledge used to process discourse may be derived from, or generalized into, general knowledge, thus providing a firm theoretical bridge between knowledge and discourse, and insight into at least one aspect of the classical problem of 'learning from text' (Johnson-Laird, 1983; Van Dijk & Kintsch, 1983; Van Oostendorp & Goldman, 1999). True, we are not yet sure how much of such general knowledge must be inferred, 'activated' and 'applied' in the construction of models, just the bare minimum, or anything that might be relevant to 'deeply' understand a discourse (Britton & Graesser, 1996; Graesser & Bower, 1990).

Most of this work in AI, psychology and linguistics did not ask many questions about the very nature of knowledge. In psychology, one will seldom find a reference to, for instance, studies on the theory of knowledge in epistemology. Psychologists use the notion of 'belief', but they would be hard pushed to respond to the perennial question, at least since Plato, what the difference is between knowledge and belief (see, e.g., Hintikka, 1962; and for general introduction, e.g., Greco & Sosa, 1999; Pojman, 1999; Wilkes, 1997). Even in linguistics and discourse analysis, detailed accounts of the role of knowledge are scarce (but see the original study of the late Paul Werth, Werth, 1999).

Similarly, the 'world knowledge' referred to in work on text processing is typically socially shared, and hence somehow characterizes groups or communities rather than individuals. Yet, we barely find references to studies on knowledge in social psychology (e.g., Bar-Tal & Kruglanski, 1988; Fraser & Gaskell, 1990), sociology (e.g., Campbell & Manicom, 1995; Choo, 1998; Knorr-Cetina, 1999; Mannheim, 1952) and anthropology (Bateson, 1972; Fardon, 1995; Geertz, 1983).

The reverse is also true, however. In the philosophy, sociology and anthropology of knowledge, one seldom finds extensive references, if any, to work on the role of knowledge in the psychology of text processing. We may all agree that most learning is being done through the processing of discourse, but many aspects of the processes of knowledge acquisition are still unknown. And the work on knowledge and discourse by Foucault (e.g., Foucault, 1972) maybe referred to by many in the humanities and in the social sciences, but is as much ignored in the psychology of discourse as he himself ignored the psychological study of knowledge.

In AI much work has been done on the formats of knowledge representation (Markman, 1999), but if one compares what is known today to what was known 20 years ago, one of the conclusions is that we have progressed very little. We have seen vast interest and resources being deployed in the cognitive and neurosciences, but we still have hardly a clue on the way knowledge is related to the neurobiological structures of the brain.

Knowledge is obviously related to other beliefs, such as opinions, attitudes or ideologies, among others (van Dijk, 1998). Discourse processing also involves these other beliefs, as even a cursory glance in the newspaper or a brief moment in parliament or in a political meeting show. However, despite thousands of studies on knowledge and such other beliefs, a detailed analysis of their relationship seems to be still on the agenda. Or worse: It is not even on the agenda. There are so many questions about knowledge that are not even asked in the psychology of discourse processing, or elsewhere for that matter. Do we, for instance, even have an approximate idea how *much* knowledge an average competent member of an epistemic community has? Hundred thousand items (propositions, etc.)? A million? Ten million?

In sum, what is needed is a broad, multidisciplinary theory of knowledge, and based on that a similarly sophisticated theory of the role of knowledge in discourse production and understanding. Such a theory should define the very conceptual scope of the notion of knowledge, and specify the relations with other kinds of socially shared beliefs.

One of the products of such a theory, as in any incipient discipline, should be an adequate typology of sorts of knowledge. Secondly, it should of course detail the structures of the mental representations of knowledge and other beliefs, and the ways such belief structures are used in any kind of processing, and especially in discursive language use, interaction and communication. Thirdly, the theory should spell out how the various types of knowledge are acquired, related to existing types of knowledge and integrated in the knowledge system. And finally, the theory should explain which are the social, cultural and political conditions, structures and processes that are involved in the reproduction of knowledge in groups and cultures so that such knowledge can play a role in discourse processing of all group members.

Of this vast theory -- which would constitute its own discipline if epistemology would not claim that, although providing only a tiny, conceptual part of the theory -- I shall here very briefly deal with only some, rather neglected, aspects of the role of knowledge in theories of discourse processing.

2. Types of knowledge

Both in the linguistics as well as in the psychology of discourse, pretty vague reference is generally made to 'world knowledge'. To establish local and global coherence, to produce or understand overall topics or macrostructures, to yield 'bridging' inferences, or to build mental models, language users are said to have and to apply some kind of 'world knowledge' (for detail, see Van Dijk & Kintsch, 1983). The problem is that such 'world knowledge' is of many different types, and does not always apply in discourse processing in the same way. Let me, therefore, begin with some observations on the ways the vast amounts of knowledge used in discourse processing can be categorized into different types.

In psycholinguistics, reference is sometimes made to the notion of 'common ground', that is, the knowledge or other beliefs speech participants need to have in common in order to understand each other (Clark, 1996). Note though that also this 'common ground' may be of different sorts, ranging from knowledge about shared personal experiences, contextually present objects, to socially known properties of the social and the natural 'world'. Also, specific groups, such as scholars or other experts, may have knowledge they share with members of their group, but which outside the group is unknown or only partly known.

Similarly, we may have knowledge about specific events or about the general properties that characterize many events (e.g., car accidents or inflation), about historically prominent and complex events such as the holocaust, on the one hand, and more abstract social and political knowledge about genocide.

Even these brief informal observations already seem to suggest the following typological criteria for a categorization of knowledge:

- a. Personal vs. social knowledge
- b. Social/Group vs. Cultural knowledge
- c. Knowledge about specific events vs. general properties of events
- d. Knowledge about historical events vs. societal-political structures

3. The distribution of knowledge

Let me try to make this typology somewhat more explicit. The typological criteria a. and b. are based on the person(s) who have or share knowledge, whereas those mentioned in c. and d. tell us something about the objects of knowledge. Criterion a. merely indicates whether or not knowledge is shared with others, in the sense that personal knowledge is by definition private and, hence, not accessible to others unless a person communicates it to others. In criterion b. both forms of knowledge are shared, but the distribution is different.

It follows that personal knowledge is merely personal until it is expressed in discourse (or other types of communication and interaction that allow recipients to attribute specific knowledge to someone), and that personal knowledge needs to be expressed in assertions, and is never presupposed. Social knowledge, on the other hand, is typically presupposed in discourse, although such knowledge may only be presupposed in ingroup discourse and not in outgroup discourse. Cultural knowledge may then be defined as all knowledge that may be presupposed in all forms of public discourse within the same epistemic community. It serves as the general basis for the specific inferences that are part of the Common Ground knowledge shared by the participants in the same communicative event. In other words, the typology as proposed not only has social implications in the sense of characterizing the nature of knowledge distribution, but also has implications for some semantic and pragmatic properties of discourse, such as the nature of implications or presuppositions.

We might further analyze criterion a. by distinguishing between personal and interpersonal knowledge. Personal knowledge is then defined as above, but interpersonal knowledge is not simply 'social' knowledge, as further specified in b., but knowledge that is only shared by a few persons and typically in forms of interpersonal discourse such as conversations. And once shared it may be presupposed (under some further conditions, such as availability) or recalled in discourse among the same persons. In other words, direct assertions of personal or interpersonal knowledge constitute inappropriate speech acts.

Of course, in the real world speakers often repeat assertions of the 'same' knowledge, but will either somehow signal that they know that the recipients already share that knowledge, so that the assertion takes the form of a reminder. Or they may not signal this and then are open to the recipient's reaction 'Yes, I know', which may function as a challenge to the appropriateness of the assertion.

Group knowledge may be acquired by repeated instances of interpersonal communication, especially in small groups of which most or all members interact in face-to-face interaction. Under different conditions, this is also the case for the interpersonal acquisition of group knowledge in the family, in peer groups, or in professional situations, for instance among journalists at the newspaper or among nurses in a hospital.

In many situations, however, the acquisition and distribution of group knowledge is more or less public, and takes place in meetings, universities, churches or political rallies, or through publications, that is, when the knowledge is expressed for and acquired by more or less extensive groups of recipients, such as professionals, members of a church or political party, or students of a specific academic field. Note that among speakers of different groups, such group knowledge is typically not or only partly presupposed.

Cultural knowledge is shared by most or all competent members of a whole culture -- that is, except children or outsiders who still have to acquire such cultural knowledge. It is this kind of knowledge that is so widespread and already part of what is often also called 'common sense', that it is generally presupposed or recalled in public discourse. Cultural members initially acquire this knowledge through socialization discourses at home or at school, and later largely through the media.

I shall assume that this generally shared, cultural knowledge is the basis of all social cognition.

Any group knowledge and personal knowledge is ultimately rooted in this cultural knowledge, even highly specialized expert knowledge. We need detailed theories of social knowledge production and distribution in order to understand the precise mechanics of these processes of 'social learning'. Below I shall come back to the cognitive properties of these kinds of social knowledge and their role in discourse.

4. Knowledge about what?

The other criterion we have used for this knowledge typology is the nature of the 'objects' about which persons, groups and cultures have knowledge. Perhaps trivially, individuals tend to have personal knowledge about autobiographical events and actions, such as the myriad of episodes that define their everyday private life, and interpersonal knowledge of the events in which they interact with others. Such events are defined by specific parameters for time, location, occurrences and participants, and are, therefore, at least in one sense unique. If only one person participates in or witnesses such events, these typically give rise to personal knowledge, as is the case for many events in daily life from the moment we get up in the morning, until we go to sleep at night. When such events are interesting for some (personal or social) reason, they become the objects of everyday conversational storytelling.

Precisely because many events in everyday life re-occur several or many times, people may also infer generalizations and abstractions, that is, less specific knowledge, for instance my general experience of shopping on Saturdays, or my general knowledge about the properties of my partner, children, family members, friends or colleagues.

5. Mental models

This kind of personal and interpersonal event knowledge is typically represented in *mental models* stored in episodic memory. Indeed, episodic memory derives its name from the specific memory people have of lived and interpreted events that give rise to what we call their 'experiences'. Mental models embody constructions of some of the properties of the event they are about: They are unique and personal, and feature knowledge (and opinions) about one specific state of affairs or event with its own unique parameters of time, place, actions/events and participants. Mental models represent how people subjectively construe the events of the world through their experiences.

But we saw that people typically form many mental models about the same or very similar events, such as their daily routines, going to work, doing shopping or seeing friends. We need only to generalize and abstract from the unique space-time parameters and the unique properties of a specific event, and then form (still personal!) generalized models of such events. These generalized models are still personal, because I am still the main protagonist in such mental models. They represent generalized knowledge about my everyday life: How I use(d) to go to work, do my shopping or go on vacation, with which persons I am related, who (also in general) these persons are, and what kind of properties they have. In other words, episodic memory is unique and personal, but not only features mental models of unique events, but also generalized models of personal event types, as well as other schematic representations that include the knowledge and opinions about other people I know.

6. Knowledge of the world

The next step in the process of abstraction and generalization brings us from personal mental models of specific events to socially shared general knowledge and other beliefs. If actions and events are generalized such that not only time and place, or some action properties, but also the participants are abstracted from (most crucially Self), then the models may take the form of socially shared scripts or other knowledge schema, for instance about types of objects, animals, people, groups, or social structures. Together, such knowledge types were traditionally called 'knowledge of the world'.

Despite vast amounts of studies in many disciplines, this 'knowledge of the world' is still a vast field of confused insights. We know virtually nothing with certainty. Psychological work as well as intuitions suggest that the huge body of our knowledge is organized, and probably forms a big neuronally based network, of which some links are used and reinforced more often than others and thus facilitate access. We assume that knowledge may be further organized in schema-like packages, such as scripts that organize knowledge about conventional social episodes, such as meetings, shopping or eating in a restaurant. The same may be true for the structures of objects, animals and people.

But apart from some more speculations about the structures of such scripts or other knowledge schemas, that is about it. Even if social episodes are mentally represented in script-like schemas, we ignore what happens at a higher level. Do all social events cluster together? Do all schemas about all animals, or about people cluster together? Is it easier to go, across species boundaries, from knowledge about heads of people to those of animals than to go from heads to feet of humans? How are all these millions of schemas related, how do they cluster, how does one access and activate them, how does one de-activate them?

These and many more questions must be answered if we want to describe and explain the amazing feat of minds/brains to find specific information among millions of knowledge items within a few milliseconds! We now read the newspaper about a terrible earthquake in India (in January 2001), and we immediately are able to access an enormous amount of relevant knowledge of which large parts we have never explicitly learned.

The standard explanation is that by reading the papers and watching TV we have acquired an earthquake 'script' and that it is this script that organizes and manages the access to all this information. But it is doubtful that current (rather abstract) concepts of scripts can accommodate the vast and unruly cluster of knowledge items people know about when they read about earthquakes and similar events.

Indeed, is the knowledge that people can die in the rubble of their house specific to the earthquake script, or a more general feature of knowledge that may be inferred from more general knowledge about possible causes of death? Do we have to assume that vast amounts of more general knowledge (e.g. about death) are 'replicated' in the large number of scripts in which they are relevant, from car accidents to earthquakes?

Or are such general knowledge items stored and organized separately (e.g. in 'cause of death' clusters of some kind) and only called upon when needed in more specific scripts. And if we do have earthquake scripts, how are such scripts related to a higher level category such as 'natural disasters', which also controls scripts on flooding and volcano eruptions, and many others? And if people die in earthquakes, does that mean that earthquake scripts are somehow related to car accidents or being-eaten-by-a-shark accidents? Or do we relate earthquakes to all other knowledge we have about the earth -- relevant to make inferences about possible changes in the landscape, and to explain earthquakes in terms of movements of planetary plates. Or with all knowledge we have about houses?

If we have millions of knowledge items (from general categories to scripts to individual propositions or other basic knowledge items), and they are all relevantly connected among each other, then the resulting framework would need so many connections that even the trillions of neural

connections of our brain would not be able to 'represent' them. Not to speak of the cognitive task consisting in managing this huge amount of knowledge -- which routes to access and strengthen, which clusters to mark, and what changes to make when somewhere a few changes are made in that huge network.

I think that we simply cannot yet begin to imagine the complexity of the task of accounting for a cognitive theory of knowledge -- not to mention its neurobiological, linguistic, and sociocultural dimensions. The current theories of scripts and schemas are only tiny steps in getting some order in this huge network, and how to use it for the many tasks that make up our daily social practices, including discourse. Much, much more, research and multidisciplinary endeavors, including vast empirical research projects on what people actually know, will be needed to get a glimpse of that mental universe, compared to which the 'real' universe in some respects may have a simple structure.

7. Group knowledge vs. cultural knowledge

One of the questions we have formulated above pertains to one of the ways all this knowledge may be organized, namely as general, cultural common ground, on the one hand, or as specific group knowledge on the other hand. We have rather informally assumed that Cultural Knowledge is the basis of all cognition. It is presupposed by specific group knowledge, as well as by all other social and personal beliefs, including opinions and ideologies. We need to know what an earthquake and houses are before we can have social attitudes about those responsible for building houses that are not earthquake-proof, or before having a personal opinion about such people. Similarly, architects have specialized knowledge, e.g., about how to build earthquake-proof houses, but also such specialized knowledge is based on an enormous framework of everyday knowledge about houses, building materials, landscapes, and how people live in houses. Indeed, specialized knowledge can only be learned by members of groups through specialized learning and training, and such processes of acquisition (in schools, universities or organizations) again require initial formulation in everyday, common sense terms requiring only cultural knowledge.

Thus, if the huge set of cultural knowledge is the basis of all social cognition of the members of a culture, how do the large number of specialized knowledges relate to that? Are specialized knowledge items learned and stored separately, or simply integrated in the one and only huge network that is called our mind/brain? Or is specialized knowledge stored separately, with only necessary links to more general knowledge about the world?

At the moment, we can only speculate about the research program that is able to respond to such questions, thinking for instance of experiments that establish differential forms of access (response latitudes, etc.) for cultural knowledge and specialized knowledge in any domain. For instance, do linguists have equal access to their knowledge about syntax as to their knowledge about cars or kitchens? Do computer scientists just know a lot more about computers than we do, or is this knowledge stored, represented and used in a very different way than our everyday, commonsense knowledge? Detailed empirical studies of learning specialized knowledge may give part of the answers.

Other evidence might be derived from the sociology of experts and their institutions: Where is specialized knowledge acquired and used? If such knowledge is used and restricted to special forms of specialized interaction, e.g., on the job, also the mental organization of such knowledge may be different from the organization of our cultural knowledge. On the other hand, the same microsociological evidence may point out that specialized interaction is not fundamentally different from everyday interaction, and that also the institutional interaction of experts features many elements of non-specialized social practices and discourse. Indeed, further linguistic evidence would suggest that only some aspects of specialized discourse (technical terminology -- largely nouns and some verbs and

adjectives -- the overall schematic organization of specialized discourse genres, as well as contextual variations dependent on institutional situations) are different from non-technical discourse. The very tentative answer to the question raised above would then be that specialized knowledge may be organized differently and separately, but at the same time multiply connected to cultural knowledge.

8. General vs. Specific/Historical Social Knowledge

Social knowledge is often represented in rather general and abstract terms, such as the scriptal knowledge people have of shopping, eating in a restaurant or giving a lecture. Whatever the details of the mental representations that define such knowledge, we may suppose that the final nodes represent variables rather than constants. A restaurant script features for instance a waitress W, but not Leila or Jenny, and not the name of the restaurant. Indeed, socially shared knowledge is useful precisely because it can be applied in many situations, by many people and in many discourse and social practices. The activation and 'application' of knowledge in such specific situations thus usually involves the substitution of variables by constants: The restaurant and the waiter become concrete and get names and specific appearances. It is this transition that defines the relations between socially shared knowledge on the one hand, and personal models of specific situations or events. It may be assumed that most of our cultural knowledge as well as specialized knowledge, is of this general and abstract kind.

However, we not only have shared social knowledge about generalities, but also about specific, historical facts. Whereas the first are the bread and butter of textbooks and much scholarly discourse, the latter fill the newspapers, TV news shows, novels and history books. Such knowledge is specific in the sense of being organized by specific, known time, place, participants and actions or events. This specific, historical knowledge may be both cultural as well as specialized. Citizens of a country usually need no further explanation when newspapers refer to the last war, to the last terrible earthquake in their country, or to the assassination of a beloved president. Although such 'facts' may be better known in restricted epistemic communities (in a city, or country, and today, this month or this year), they are sometimes nearly universally known, such as the Second World War or the Holocaust, partly because of history and geography lessons in school, as well as being presupposed in the very news media, and thus 'accidentally' learned by inference.

Interestingly, such generally known historical facts may have a structure that is quite similar to that of personal mental models, with the difference that they are broadly shared, and of course without the individual properties of personal mental models, such as personal opinions and autobiographical memories. Indeed, as is the case for general and abstract knowledge, knowledge of specific historical facts is in a sense 'normalized', and may of course represent an equally shared social opinion, as we may have about the (horrors of the) Holocaust.

Historical fact knowledge in principle applies in discourse processing as is the case for other social knowledge: It can be presupposed, need not be asserted, may be recalled if necessary, and allows applications in models as well as many inferences. Thus, our shared knowledge about the Holocaust allows us to make inferences about the horrific consequences of ethnic hatred, and to make comparisons with other horrible acts of genocide current discourses may be about. That is, whereas general and abstract knowledge may be simply 'applied' by specifying variables, historical facts often give rise to comparisons with new facts.

9. What does it mean, to 'share' knowledge?

Another important problem to solve in a multidisciplinary theory of knowledge and its applications in a

theory of discourse processing is the question what exactly it means to 'share' knowledge. We may more or less informally speak of common knowledge, or to 'socially shared' knowledge, but does that mean that all members of an epistemic community have exactly the 'same' knowledge, personally variable versions of the 'same' knowledge, or merely different but overlapping knowledge sets?

For personal knowledge the answer is trivial, because personal knowledge is by definition not (yet) shared: mental models are private until we express parts of them in talk or text. That is, personal knowledge may become interpersonal knowledge when others share part of our mental models. Of course, the mental models of the recipient are also personal, and hence never identical to that of the speaker. Indeed, our recipients may have their own interpretation and evaluation of the things we tell them about. Recipients may say they 'know' a story we already told them about, but such knowledge may only refer to their main topic or upshot, not to the details as construed by the recipient. Stories, just like other forms of discourse, are by definition incomplete, and it is the recipient who needs to construct the vast amount of inferences needed to construct her or his own 'version' of the facts. In other words, interpersonal knowledge is only partially shared among individuals.

Social knowledge on the other hand is defined, among other things, as knowledge that is shared by the members of a community. In a sense this knowledge must be the 'same' throughout the community, much in the same way as speakers know and speak the 'same' language. And yes, most or all competent members of (say) European communities know what cars, cows and computers are. At least they know enough to understand most forms of public discourse, such as the news or textbooks. On the other hand, little empirical research is necessary to know that there must be individual differences as well. Some of us know more about cows, cars or computers than others, even without being experts, simply because of different social contexts, different interests, readings and personal experiences. Indeed, we may all know and speak the 'same' language, but there are idiolects and individuals may have their own personal style, have a different vocabulary, master different discourse genres better than others, and so on. The very variability of personal experiences, autobiographies, and hence of mental models, also may lead to differences of knowledge, which people may or may not share with one or more other members of the community. The personal differences among discourses are one of the many expressions (and hence forms of evidence) of this variation in language use and knowledge.

However, despite such personal differences, we also apparently seem to have some social knowledge in common, and it is this socially shared knowledge that is presupposed by many public discourses. Can we say that this social knowledge, whether general and abstract (such as our knowledge about wars) or specific and historical (like our knowledge about World War II), is indeed 'the same' for all members, in the sense of the same file or program, that runs the same in the same situations? Or do we have to assume that because all members have learned about their world and their language in very different contexts, and through different contents, we also must have, at least slightly different, social knowledge?

It is however not easy to conceptualize this notion of something that on the one hand must somehow be 'the same' and yet personally variable. The easiest would simply be one of elementary set inclusion or overlapping sets: all members have their own personal knowledge sets, but because of similar education, experiences and readings, and especially because of the same presuppositions in public discourse, these sets overlap, and the intersection would then be socially 'shared' knowledge.

Another way is to conceptualize 'shared' knowledge not so much in quantity, but in quality: What we socially share is not merely an overlapping subset of knowledge, but the 'most important information' of each concept. That is, we may have in common the basic, prototypical knowledge about cows, cars and computers, without all of us having the same details. Whereas set-overlap is quantitative and flat, this kind of sharing is rather hierarchical: it only applies to the higher level nodes of a

selection of concepts.

Other conceptualizations of 'sharing' are possible, but I provisionally would vote for the 'hierarchical' option: We have basically the same knowledge base (like knowing many words in our language), but this means only relatively superficial knowledge.

10. Knowledge and discourse processing

How do these different types of knowledge play a role in the production and comprehension of discourse? Traditionally 'world knowledge' was supposed to be activated for instance in the understanding of words, the construction of propositions, the derivation of inferences, the establishment of local and global coherence, and the construction or retrieval of mental models. Slightly more differentiated is an account in which such knowledge is construed in terms of scripts, scenarios or other schematic formats for the representation of 'packages' of knowledge, thus facilitating script-based activation as well as updating of knowledge. Unfortunately, this work has drawn more interest from psychologists and AI scholars than in sociology and anthropology: the sociocultural acquisition, uses and changes of knowledge and knowledge scripts is usually ignored in such studies.

Current work on the management of mental models in processing specifically focuses on the amount and the nature of inferences derived from such 'world knowledge' in the construction of mental models. Indeed, how much do we need of our scriptal knowledge about shopping in a supermarket in order to understand a story in which protagonists are doing such shopping? How many and how 'deep' must or may inferences be drawn from such knowledge? Some scholars opt for a bare minimum (so as to construe necessary bridging inferences), others for a flexible amount necessary to establish local and global coherence, and others again for much more comprehensive activation of knowledge spaces, which may explain heightened accessibility even of those knowledge items that were not really necessary in order to make necessary inferences or establish coherence.

We think that the answer may integrate aspects of all three responses: Depending on readers and context, as well as the level of description of the text or text fragment (highly detailed description, or high level, global description of events and objects), all three options may on occasion be relevant. In some situations and for some people reading and comprehension is fast and superficial, leading to fragmentary models; sometimes it just does activate the knowledge that is necessary for a minimally connected or coherence mental model; and in some cases readers may 'drift' into detailed associations of the events portrayed in a text, whether they read a newspaper article or a novel, activating vast amounts of knowledge and other beliefs, leading to very detailed, and often highly idiosyncratic mental models we know as the daily experience of 'dreaming' or 'fantasizing'.

11. Knowledge and context

It is within this well-known framework of the psychology of discourse processing that we need to explore some of the implications of the knowledge typology discussed above. However, at the same time, we have signaled the theoretical incompleteness of the standard theory. One major lack has been discussed above: a very limited concept of the knowledge(s) involved in discourse processing. The other theoretical deficiency is the virtual ignorance of *context*. Apart from taking into account some independent variables for different language users (age, gender) and special experimental conditions, the psychological theory of discourse processing, is largely context-free. Whereas discourse structures and their mental correlates are being accounted for in increasingly sophisticated detail, the complex context structures that define the communicative situation are either ignored or reduced to the abstract

minimum of different experimental conditions. For the rest, the artificial 'context' of the lab is considered to be the 'same' for all subjects.

In order to understand the role of knowledge in comprehension and production, however, communicative situations need to be modeled in terms of context models, just as the referents of discourse are accounted for in situation models or event models. In other words, we not only need semantic models, but also 'pragmatic' ones. To wit, language users activate, select, and express knowledge items primarily as a function of their beliefs about the knowledge of the recipients. These beliefs are part of the representation language users construct in their models of the recipients of the communicative situation. Secondly, also the aims of the communicative event will sometimes activate more or less knowledge, and depending on the type of situation a speaker may want to activate more or less general but also specialized knowledge. Speakers during a scholarly conference obviously construe mental models in which their social identity as scholars, as well as their audience and their knowledge, require ongoing activation of specialized knowledge.

In sum, many of the relevant properties of the communicative event that together define the context models of a speaker or recipient, may be involved in the subtle management of knowledge: its activation, its uses in word or sentence interpretation, or its role in the construction of mental models. We manage such knowledge differently when we speak to close friends than when we speak to strangers, when we speak to children or adults, and differently in various social roles and institutions. That is, the role of knowledge in text processing must be contextualized, even when some aspects of knowledge activation may be 'automatic' and context-free. In more technical terms this means that the activation and uses of knowledge is controlled by various parameters of context models. This also explains how and why assertions and presuppositions of knowledge may vary according to whether what kind of knowledge is being supposed to be shared by the recipients, as discussed above.

12. Knowledge processing

We now are in a position to be a bit more specific about the role of knowledge in discourse comprehension and production. Within the confines of this paper, this means that at present we can offer no more than a summary of the hypothetical processes involved. Let us do so for the production process, because that is still the side of processing we know least about:

- a. **Situation.** Discourse production, like other social practices, takes place in a social situation or episode. Social actors who participate in the actions of such situations have personal, interpersonal as well as shared social knowledge about such situations. These knowledges allow them to understand and evaluate such situations by ongoingly constructing mental models of them. These mental models are subjective episodic representations that define our everyday experiences. We may therefore call them 'experience models'. In other words, even before engaging in discourse, participants already are activating and applying knowledge that defines the current situation.
- b. **Context.** In some social situations, interactions take a communicative nature, e.g., when speakers want to convey some of their personal or social knowledge or other beliefs to other social participants or engage in some discursive social action. In order to be able to do that appropriately, they need to establish which of the properties of the social situation are *relevant* for the communicative event. The mental representation of these relevant properties is called a context model, or simply a 'context'. In other words, contexts are not 'out there', but (inter)subjective mental constructs of participants. Such contexts models are construed on the basis of personal and social knowledge about educating, schooling and students. Context models feature such categories as overall social domain, overall (macro) action -- e.g. education or legislation -- local actions (like

writing an editorial or teaching a class), participants in various social roles, a Setting (Time, Space, etc) and some other environment properties. Context models control the whole production and comprehension process.

- c. **Contextual knowledge.** An important part of the context model in processing of production are the representations of speakers/writers about the knowledge of the recipients. These representations may be vague beliefs or assumptions, or more or less firm forms of knowledge, e.g., interpersonal knowledge when personal knowledge has been communicated before to the recipient, group knowledge when the recipient is a competent member of the same group, or cultural knowledge when the recipient is a member of the same culture. These knowledge and beliefs about the knowledge of the recipients may apply to the topics the communicative event will be about, or about the properties of the social situation (knowledge about the social domain, the overall and local actions, the various social roles of the participants, etc). During the communicative event speakers ongoingly adapt these various knowledge assumptions to what has been said already, and to what the recipients express or presuppose about their own knowledge. In other words, context models feature the complex control mechanism for the activation of the relevant communicative knowledge of the speaker/writer, as well as the management of such knowledge as a function of the assumed knowledge of the speaker. Formulating the detailed interactive mental strategies involved is by itself already a major task for an explicit theory of discourse processing.
- d. **Event Models.** Once context models have been (partially) construed, speakers are able to decide which personal or social knowledge and beliefs may be selected for expression in discourse, which knowledge need only be signaled and presupposed to be known to the recipient, and which knowledge can simply be left implicit. Of course, speakers' knowledge (mental models of events witnessed, participated in or read about) usually 'exists' already when they want to engage in discourse about it, but it is only under the epistemic constraints of the context models that the *relevant* knowledge is being selected or focused upon. Speakers have vast amounts of knowledge, and only a small part of this knowledge is interesting enough so as to interest recipients. Thus, only small fragments of mental models will eventually serve as input for the discourse formation rules. The relevant mental model will however remain active during discourse production in order to be able to provide further 'information' when needed.
- e. **Meaning production.** Event model information selected by context model constraints will form the input of the various processes of semantic discourse production: topic formation, local coherence relations, consequences, etc. As a general rule, only those propositions are construed in the semantic representations that are not already (well) known to the recipients, although under special conditions, discourse repetition of known representations may be allowed. Apart from readymade fragments of social psychology, we'll find there many types of expertise.
- f. **Expression/Formulation.** Finally, relevant elements of event models (personal knowledge, etc), context models (situation knowledge) and semantic representations form the input for the various levels of expression or formulation: (i) lexicalization, (ii) syntactic structures, (iii) phonological and graphical/visual expression, and (iv) overall discourse schemata for overall ordering of text or talk. Thus, lexicalization will depend, e.g., on the knowledge about the assumed lexical knowledge of the recipients (represented in the context model), on the assumed object knowledge of the recipients, on the assumed context knowledge of the recipients (in formal contexts, more formal words will be used), and so on. Even the production of syntactic structures may depend on the knowledge of the speaker about the linguistic knowledge of the recipients, as well as about recipient's knowledge about the communicative situation. In principle, any structure that can vary in discourse may be a function of contextual constraints, including knowledge about the knowledge of the recipients about the communicative event or the genre. Most important for our discussion

however is the complex signaling of asserted or presupposed knowledge in

- clause hierarchy and position; e.g., initial that-clauses often express presupposed knowledge;
 - sentence initial position (topicalization); often signal known information
 - text initial position (headlines, summaries): tend to express important, new information;
 - text final position (summaries, conclusions): tend to express (con)textually known by important information;
 - definite articles: tend to express known information;
 - raising intonation (in questions): express wanted (unknown) information
- (for other discourse structures that express or signal knowledge, see van Dijk, 1997).

13. Concluding remark

In other words, we see that at all levels of text production, there is an ongoing process of complex control of variable structures that are a function of the various kinds of knowledge of the participants, including knowledge about each other's knowledge. An explicit theory of text processing needs to take into account this kind of knowledge management, in order to be able to describe and explain more adequately how language users adapt their discourses to their own knowledge as well as to that of the other participants. Crucial in such a theory is the integration of context models that represent such (mutual) knowledge states, as well as a sophisticated theory of the various types of personal, interpersonal, social (group) and cultural knowledge involved.

References

- Bar-Tal, D., & Kruglanski, A. W. (Eds.). (1988). *The social psychology of knowledge*. Cambridge: Cambridge University Press.
- Bateson, G. (1972). *Steps to an ecology of mind: Collected essays in anthropology, psychiatry, evolution, and epistemology*. San Francisco, CA: Chandler Co.
- Britton, B. K., & Graesser, A. C. (Eds.). (1996). *Models of understanding text*. Mahwah, NJ: Erlbaum.
- Campbell, M. L., & Manicom, A. (Eds.). (1995). *Knowledge, experience, and ruling relations: Studies in the social organization of knowledge*. Toronto: University of Toronto Press.
- Charniak, E. (1972). *Toward a model of children's story comprehension*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Choo, C. W. (1998). *The knowing organization: How organizations use information to construct meaning, create knowledge, and make decisions*. New York: Oxford University Press.
- Clark, H. H. (1996). *Using Language*. Cambridge: Cambridge University Press.
- Fardon, R. (Ed.). (1995). *Counterworks: Managing the diversity of knowledge*. London New York: Routledge.
- Foucault, M. (1972). *The archaeology of knowledge and the discourse on language*. New York: Harper & Row (Harper Colophon).
- Fraser, C., & Gaskell, G. (Eds.). (1990). *The social psychological study of widespread beliefs*. Oxford New York: Clarendon Press Oxford University Press.
- Geertz, C. (1983). *Local Knowledge: Further Essays in Interpretative Anthropology*. New York: Basic Books.
- Graesser, A. C., & Bower, G. H. (Eds.). (1990). *Inferences and text comprehension. The psychology of learning and motivation*, Vol. 25. New York: Academic Press.

- Greco, J., & Sosa, E. (Eds.). (1999). *The Blackwell guide to epistemology*. Malden, Mass.: Blackwell Publishers.
- Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca, N.Y.: Cornell University Press.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference and consciousness*. Cambridge: Cambridge University Press.
- Knorr-Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Cambridge, Mass.: Harvard University Press.
- Mannheim, K. (1952). *Essays on the sociology of knowledge*. London: Routledge & Kegan Paul.
- Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Erlbaum.
- Pojman, L. P. (1999). *The theory of knowledge: Classical and contemporary readings*. Belmont, CA: Wadsworth.
- Schank, R. C., & Abelson, R. P. (1976). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, N.J. New York: L. Erlbaum Associates distributed by the Halsted Press Division of John Wiley and Sons.
- Van Dijk, T. A. (1998). *Ideology. A multidisciplinary approach*. London: Sage.
- Van Dijk, T. A. (Ed.). (1997). *Discourse studies: A multidisciplinary introduction*. 2 vols. London: Sage Publications. (New, one volume edition, 2011).
- 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: Erlbaum.
- Werth, P. (1999). *Text worlds: Representing conceptual space in discourse*. Harlow, England: Longman.
- Wilkes, A. L. (1997). *Knowledge in minds. Individual and collective processes in cognition*. Hove: Psychology Press.

Note november 2018. For details of the multidisciplinary theory of knowledge, and its relations to the structures and processing of discourse, see Teun A. van Dijk, *Discourse and Knowledge. A Sociocognitive Approach*. (Cambridge: Cambridge University Press, 2014). The Spanish translation of this book was published by Gedisa in 2016. For my other publications on knowledge and discourse, see <http://www.discourses.org/projects/knowledge/>.