

The Context of Conversations

— Texts and Communities

HansWeigand^{*}, Stijn Hoppenbrouwers[§], Aldo de Moor^{*}

^{*} Infolab, Tilburg University

[§] ID Research B.V., Gouda, the Netherlands

H.Weigand@kub.nl, S.Hoppenbrouwers@idr.nl, AdeMoor@kub.nl

Abstract

Communicative action concerns achieving and maintaining mutual understanding among all those who are involved in a coordinated (organizational) situation. Communicative action is embedded in a social context. Current approaches in the Language/Action Perspective have focused mainly on the dynamic aspect of language, that is, on (business) conversation. In this paper, we develop a new model of context in which particular attention is given to the role of explicit context, that is, the role of texts. A text is not a (communicative) action, but it shapes the world in which communicative actions can take place. The model is illustrated by three cases: the design of a communication tool, the management of Organizational Memory, and the support of professional communities. The three cases exemplify three abstraction levels at which texts can be analyzed.

1 Introduction

In their influential book *Understanding Computers and Cognition*, Winograd and Flores not only introduce the use of speech acts in modelling communication, but also present a fundamental critique of the rationalist view of cognition that underlies most work in Artificial Intelligence and Computer Science. Drawing on the philosophy of Heidegger, they attack the idea that we are able to make complete models of (parts of) the world that can subsequently be used for building intelligent systems incorporating these models. For example, they argue that (p34-35):

- our implicit beliefs and assumptions cannot all be made explicit
- practical understanding is more fundamental than detached theoretical understanding
- we do not relate to things primarily through having representations of them
- meaning is fundamentally social and cannot be reduced to the meaning-giving activity of individual subjects

The world is encountered as something already lived in, worked in and acted upon before we start thinking and speaking about it. World as the background of obviousness is manifest in our everyday dealings and every possible utterance presupposes it. That which is not obvious is made manifest through language. That which is obvious is left unspoken, but is as much a part of the meaning as what is spoken (p58).

In this paper, we want to arrive at a deeper understanding of what is meant by *explicit* and *implicit*. Roughly speaking, “explicit” versus “implicit” is mostly taken to mean “represented by symbols” versus “being present in peoples minds”. Vague as this may seem, it is often the bottom line in most non-theoretical discussions which involve these notions.

Mostly taking an individual perspective, cognitive science deals with concepts as part of the question how our mind/brain works. We avoid this field as such, taking the neutral stance that it is still unclear whether or not symbolism is fundamental to cognition, or whether symbol manipulation is just one of the many tricks up the sleeve of the human mind. The most prominent alternative to symbolism is connectionism (Bechtel and Abrahamsen, 1991). However, communication being a social activity, it requires not so much an individual but a social, *communal* approach to knowledge. *Communal knowledge* refers to knowledge people share, or have in common.

Communication is based on the existence of some common ground between communicating parties, but this common ground may include quite a few different kinds of knowledge. The most basic common ground we share is that we are all humans, with some shared experience as a result of our inherently similar biology and basic psychology. At some level, this includes an innate capacity for acquiring language (Pinker, 1995). On top of all this, we have many similar experiences as we grow up, which however may differ to a considerable extent (culture and personal history). Beyond this, experiences can be shared right up to highly specific domain level (e.g. specific work) and even situations, for example a task. All these experiences result in the acquisition of knowledge, of many kinds and at many levels.

In and through social interaction, humans create communal norms, values, and rules. Norms and rules are not necessarily explicit; in fact, many are not. Communal signs, i.e. “agreed on” (shared) combinations of forms and meanings (most typically words) could be seen as akin to other kinds of agreement (e.g. implicit rules, norms, values). Agreement on sign systems enables people to communicate, provided that they are part of the same community and thus partake in a certain number of communal agreements. When used in conversations, signs are interpreted with in the background a wealth of contextual agreements in addition to some information provided in the actual conversation.

In some cases, people “fix” agreement by means of *texts*. This may be done either deliberately or coincidentally. In any case, it involves symbolic constructs (usually complex): independent objects (artifacts) which stand more or less apart from their author and the situation in which they were created. Texts may carry information revealing their origin, but how this information is interpreted and whether it is deemed important depends on the interpreting party just as all other information contained in the text. Note that not every explicit conceptualisation represents agreement which holds outside the conversation. Whether or not it does depends on the authority assigned to the text and whether it is valid at the time of interpretation.

Very domain-specific shared meaning can still be communal knowledge, even though the community may be very small. And very general knowledge may exist in a very explicit form, for example in the case of formalised international laws. So we cannot say that implicit knowledge is general and explicit knowledge specific.

Essentially, humans always act upon *internalised* knowledge. Text is not knowledge, it is of course just a representation. However, text may well have a strong control over internalised knowledge: depending on the authority (incl. trustworthiness) of a text, it may actually shape internal knowledge.

The structure of this paper is as follows. In section 2, we introduce a new model of context in which implicit and explicit context are distinguished and related to the notions of text and community. Context is viewed as existing on three levels of abstraction. We also discuss the ambivalence of texts as being both instrumental but sometimes also obstructing. After this theoretical basis, we illustrate the role of explicit context on the three levels of abstraction. In section 3, we show how the context model can influence the design of the communication system (first level). In section 4, the model is applied to the recently emerged field of Knowledge Management: how can contextual knowledge be managed (second level)? In section 5, an application is discussed which deals with the problem of how communities create and maintain their texts in a legitimate way (third level). This application is based on a Ph.D. project on virtual professional communities.

2 A new Model of Conversational Context

In this paper, we want to develop a new model of context in line with the Language/Action Perspective. We draw on the three levels of communication used by Dietz (Dietz, 1990). Dietz distinguishes between a documentary level, an informational level, and an essential level. At the documentary level, we see messages going around represented in some medium. The medium can be speech, writing, computer records, or whatever. At the informational level, we operate at a level of abstraction higher than the medium and focus on the content of communication: the information that is exchanged. At the essential level, we focus on what communication brings about, on what has effect in the intersubjective world¹. Note that the documentary level and informational level together make up what is traditionally called the “sign” as a combination of a token and a signifier (de Saussure, 1916). Linguistic acts (referring, predicating) are performed by means of utterance acts (Searle, 1969). Linguistic acts themselves are instrumental in the execution of speech acts (Weigand, Hoppenbrouwers, 1998).

We claim that context can also be viewed on these three levels:

- *Context at the documentary level.* The location in which the message is represented. For example, a message “the price of product X is Y” can occur in a price catalogue, but also in a strategic scenario. The context of the message plays a role in the interpretation.
- *Context at informational level.* A message contains information, but this information is most often only partial. As Dik (1989) states, messages are more like Delta’s: they specify a kind of operation that the Hearer must apply to the shared knowledge. At the informational level, the context is the total of background knowledge relevant to the message that the communicative agents share. The less context they share, the more explicit the message has to be.
- *Context at the social level.* The effect of communicative acts is dependent on social institutions or conventions. As Derrida has spelled out, a communicative act would have no effect if it would be unique, without precedence or possible succession. The speech act of marrying is rooted in the age-old institution

¹ As can be seen from its description, a better, more neutral label would perhaps be “social”. To call the social level essential is a choice that can be made independently from the three-level distinction itself. In the following, we will therefore use the term “social”.

of marriage and national or religious laws surrounding it. But also simpler speech acts like asking a secretary to copy something depend on role definitions of a secretary and sometimes on task descriptions that an organization has fixed.

In other words, context has a locational character (the physical or virtual space), a symbolic character (the knowledge space), and a social character (the social world).

Secondly, we distinguish implicit and explicit contexts. There is always a context, but it need not be explicit: it is only explicit when it uses symbolic signification (“language”), and when its boundaries are clear and serve some identifiable purpose. On each of the three levels, this means the following:

- *Documentary level.* The context is explicit when it is separated. A whiteboard offers an explicit context in group discussions. A company product database is also an explicit context. An explicit context is finite. In contrast, everyday language sometimes draws on the whole body of world literature, for example, when a reference is made to a biblical image. This context is implicit.
- *Informational level.* The shared knowledge of the agents can be left unexpressed, or it can be made explicit, typically in the form of a text. In the case of an organization, the explicit informational context consists of all kinds of shared knowledge resources, also called “organizational memory” (Reimer 1998).
- *Social level.* The context is explicit when it is institutionalized, for example, in the form of an organization (Taylor 1993). Other examples are an auction, or an EDIFACT agreement. The context is implicit for example in family life or in business as long as it is based on notions which are usually not further defined, like “fairness” and “trust”.

We should keep in mind that the three levels are levels of abstraction, not separate domains. This is also the case when they are applied to explicit context. For example, an EDIFACT agreement is a contract (agreed-upon procedure) at the social level, but it takes the form of a text (informational level) written on some identified piece of paper (documentary level).

Note further that at all the three levels of abstraction, we can distinguish atomic and complex objects. For example, on the documentary level of message, we not only have characters, but also words, sentences and message sequences. On the social level, we not only have elementary speech acts, but also larger conversational units. In (Weigand & vd Heuvel, 1998), five aggregation levels are distinguished: speech act, transaction, workflow, contract/interaction, scenario. In some sense, larger units function as a context for smaller units. For example, the speech act of requesting cannot be seen without taking the transaction (request/commit) into account, since only in this context, the request becomes meaningful. In another context (another transaction), the meaning of the request may be quite different. Although differing in details, a similar aggregational notion of context was already introduced in the Milan Conversation Model (De Michelis et al, 1994) that stresses the situatedness of conversations in “work practices”. In this paper, we will not focus on this aggregational notion of context.

2.1 Implicit vs Explicit Context

The distinction that we have made between implicit and explicit context needs some further discussion, since in our view, the distinction has not been sufficiently recognized in current LAP approaches. A notable exception is Taylor (1993;1996a;1996b), to whom we will come back shortly. Also in the work of Carla Simone, context has been taken very seriously; we will compare our approach with her results in section 2.4.

Winograd & Flores have stressed that language is performed in a given world, which is an open-ended, not explicitly given life world. This life world constitutes the context for speech acts. In other words, W&F make a distinction between implicit context and explicit language. However, this binary distinction hides the fact that explicated context (explicated by means of language) can itself function as context for new language acts, and in this case, we have to do with an explicit context. It might be argued that the implicit context is never exhaustively explicated, and hence the explicit context does not totally replace the implicit context. This we can agree with. But what is important at this point is that there are explicit contexts and that they do play a role in the interpretation of communicative acts.

There is one important condition for the existence of explicit context, and that is that the language acts that explicate the (implicit) context have effect over a longer period of time. Therefore, the possibility of explicit context is closely bound to the existence of writing as opposed to oral speech. It is not absolutely necessary, since a language act can for example also be kept in memory and retold in the form of myths, be recorded, or whatever, but obviously, the existence of a writing system greatly facilitates the representation of explicit context.

Although the possibility of explicit context is not worked out in Winograd & Flores, it is completely in line with the phenomenological approach that they take. They show how language creates a new world against the background of an implicit lifeworld. Much in the spirit of Derrida, we can say that in writing the power of

language is even more clearly visible than in dialogue. Therefore, writing should not necessarily be seen as opposed or orthogonal to communicative action, but rather as capitalizing on one of its main functions, that is, the disclosure of the world.

Before taking a closer look at “text” and “community”, and the relationship between them, let us give some concrete examples of “text” in an organizational environment.

- An interface specification between two systems
- A database with personnel records
- An EDIFACT agreement between two companies
- An employee’s contract
- A mission statement
- A map of the company building
- The budget overview for the next period
- A management report
- A product catalogue
- A web-site with articles on some topic
- A name sign on the door of the room
- A (paper) telephone guide

In all these cases, we are dealing with persistent linguistic objects. They describe or explicate the implicit context. In some cases, the focus is on social relationships: for example, an EDIFACT agreement that describes explicitly the obligations and rights of two parties. In other cases, the focus is on shared knowledge of the world: for example, the telephone guide. Also subjective views and intentions can be explicated: the mission statement is an example. But apart from describing the implicit context, the texts provide a context for the conversations in the organization as well. For example, in the presence of a product catalogue, an order procedure can refer to a product number instead of describing the product completely every time. So texts have two faces: a face towards the implicit context on the one hand and a face towards the conversations on the other.

Most of the examples are texts that are closely related to a certain organization. The example of the web-site with articles is a bit different from the others. Articles are typically positioned at a more generic level. Even if they describe a certain situation, e.g. a case study, the underlying goal is to come to a more general understanding. In other words, articles do not describe the contexts of a particular organization, but are part of a body of communal knowledge (see section 2.3). The other examples are all related to a particular organization, or its environment.

2.2 Texts

James Taylor (1993;1996a;1996b) has described, in a series of articles and books, the importance of text for organizations. Conversations and texts are dependent on each other in an entangled hierarchy. On the one hand, conversations can be viewed as occurring in an organization, in other words, against the background of texts; on the other hand, the organization and its texts are created and maintained over time by means of conversations. In this way, text and conversation are complementary but irreducible worldviews. Not surprisingly, many of Taylor’s arguments are derived from the philosophy of Jacques Derrida, which can be called a philosophy of writing (Derrida, 1988).

According to Derrida, there has historically been a prejudice to see dialogue and speech as primary, and writing (text) as derived and secondary. The rationale for this assumption is that with writing something has been added (a “supplement”): marks or traces that persist beyond the moment of their production. However, what is special about writing, according to Derrida, is that, unlike speech, it can be detached from the speech situation, author, and original recipients and hence has a certain degree of autonomy. It can be used and reused in many situations just because it is not bound to any specific situation. In this way, writing is not an add-on to, but the enabler of communication.

The only way communication can occur is if an utterance can exist independently of its speech situation. It is only because the speaker draws on this potential that she can put it to use in the specific speech situation. Signs that do not rely on a given autonomous potential are limited to indexical or iconic behaviour: pointings and grumblings. So in this way, even utterances of ordinary conversation are themselves a form of writing (Écriture).

Taylor also builds on the work of Ricoeur (1986) on “distanciation”. According to Ricoeur, discourse (language) differs according to whether we think of it as an event or as work. Discourse as an event is seen as realized temporally, in the present, referring to a specific situation, addressed to someone. Intention and meaning overlap. Discourse as work is seen as fixed (in writing), referring to a general world, accessible to everyone.

Intention and meaning are dissociated. By writing the discourse down, it is "objectified". Taylor identifies no less than six degrees of *separation* when a conversation is translated into a text (space does not allow us to repeat them here).

Whereas most LAP approaches today focus on the description of communication structures, we suggest to take another dimension into account as well: the line that goes from conversations to explicit context (texts) to implicit context. Once we do take texts into account, besides conversations, several interesting questions can be posed. For example:

- What exactly is the role of texts (for example, a signed contract) in relation to conversations? How do we formalize that subsequent communicative acts derive their effect from such a text, and are not effective otherwise?
- How is an implicit context disclosed in explicit context, and why?
- In which cases are explicit contexts needed, and in which cases not?
- What is the effect of an explicit context on the form of conversation? It seems that there is a trade-off: the more explicit context is given, the shorter and more direct the conversation can be. On the other hand, in the absence of explicit context, the conversation should be more elaborate and explicit itself.

2.3 Communities

In (Clark 1996), the psycholinguist Herbert Clark works out a functional approach to language in which language is viewed as a form of joint action. It is the joint action that emerges when speakers and listeners perform their individual actions in coordination, as ensembles. Clark argues persuasively that conversations as joint activities presuppose certain common ground, and also add to it. Common ground is defined roughly as "a shared basis of propositions where every member of the community has information that the shared basis holds, and knows that the every member has this information".

According to Clark, common ground is essential for coordination. In fact, he asserts that for something to be a coordination device, it must be a shared basis for a piece of common ground. An explicit agreement made by two parties can be a coordination device, but also a precedent can take this function; as long as it is part of common ground.

If common ground is an essential condition for coordination, how do language users find the shared bases for common ground? To answer this question, Clark comes to communities. We categorize people according to what they know or believe. The main categories are cultural: for example, nationality, profession, hobbies, language, religion or politics. A cultural community is a set of people with a shared experience that other communities lack. Ophthalmologists do not all live in one place or know each other, but what makes them a community is a shared system of beliefs, practices, nomenclature, conventions, values, skills and know-how about eyes, their diseases, and their treatment. It must be stressed that a community is more than a collectivity of people. It is only a community when there is a certain "consensus" on "common ground" (cf. Weigand & vd Heuvel, 1997).

What are the contents of communal common ground? Clark gives a tentative answer. First, it contains information about human nature. This is grounded in the fact that we share similar bodies. Secondly, there are what Clark calls communal lexicons. Different languages offer different lexicons, but within one language, sublanguages can be distinguished that have their own "slang" or "terminology". In some sense, it can be argued that any lexicon is communal: conventional word meanings hold only for words in particular communities (Lewis, 1969). Thirdly, there is a body of cultural facts, norms and procedures. This includes knowledge about social roles, such as those of husband, or neighbour, and what can be expected from them. It is the kind of knowledge that AI has tried to capture in the form of scripts.

The question of how to find the shared basis for common ground can partly be rephrased as a question of how people recognize someone to belong to a community. Different kinds of evidence can be used to determine community membership. The simplest is natural evidence: for example, if someone starts talking in French, you know he is a speaker of French, and probably a member of the French culture. But there is also deliberate display of community membership. Examples are dress codes (male/female, white-collar/blue-collar etc), or badges used by certain personnel.

The relationship between text and community is two-directional. On the one hand, communities "ground" texts. A text that is not supported (anymore) by a community, for example, an old Babylonian constitution, can be interesting for a group of scientists as an object of research, but it is not effective anymore. The community is also a reference point when interpretation conflicts arise. On the other hand, texts can help to "identify" communities. The common knowledge of a community need not be explicit, but especially in the confrontation with other communities, or in the face of internal conflict, texts can help to establish the identity of the

community. Examples are national laws, a canon of scientific journals of a certain field, but also mission statements in companies, or just a company logo.

2.4 The ambivalence of texts

Fixing shared meaning and making it explicit helps by limiting the options people have in interpreting communication. Since communication creates overhead, action with a minimal need for prescriptive and supporting communication is optimally efficient. (Note: if computerised systems consisting of symbols count as texts, computerisation is the ultimate form of text: it makes procedures explicit to the level of actually automating them).

In this respect, texts are a mixed blessing at best: they help fix agreements of interpretation (and hence help decrease uncertainty), but unless they are completely formal, they themselves allow relatively much room for interpretation compared to most verbal, situated conversation. To make the context of a text clearer (thereby reducing uncertainty of interpretation), other texts may be used. Thus, texts may help delimit the “interpretation space” of other texts. For example, national law requires jurisprudence to be interpreted.

Many actions (including linguistic and communicative actions) require a degree of flexibility in behaviour which cannot be statically prescribed. Therefore, fixing agreements (i.e. making shared knowledge explicit) can easily have unfortunate consequences, especially if such agreements concern procedures of a mandatory nature, and then become an impractical burden. They may cause situationally inappropriate actions to be carried out even if the agent responsible is aware of this (i.e. the authority of the text overrules individual, situationally determined action). This is the well-known problem of bureaucracy.

The use of texts may have some further unintended consequences. In the first place, texts are entities with an existence of their own. This means that they are owned by somebody or some group. The owner may restrict access to the text, consciously or unconsciously, and thereby make it difficult for outsiders to understand the conversations that occur. The owner may also restrict the updating of the text.

If communicative agents have put their way of working in a text, and then want to change their way of working, the text must be changed. In that case, the owner of the text must cooperate. As a practical example, consider the case of a piece of software (based on some design document). If users want to change the software or the design document, they have to apply for these changes at the IT department, since this is the subject by which the software is owned, or at least maintained. There is abundant evidence of cases in which this process is cumbersome and time-consuming. It might be suggested that in some cases the owner can be eliminated, or that the group itself is the owner. However, even if this is possible, there needs to be some procedure by which the group can come to updates of the texts, since it must be a coordinated action (cf. section 5).

Simone (1997) discusses a related problem that she calls *linguistic opacity*. This involves “the uneasiness people may experience in moving in a linguistic framework which is perceived as unfamiliar”. Opacity arises from the discrepancy between the knowledge possessed by the cooperating actors and the knowledge needed for the interpretation of the messages they exchange. There is a negative feedback cycle here. A group tries to harmonize its knowledge, and deliberately builds up a common set of references. This leads to the people simplifying communications by exchanging ambiguous and incomplete information (without the intended common ground) whose interpretation can become difficult in the presence of different sets of references. The CHAOS system (De Cindio et al, 1986) tries to address the linguistic opacity problem by supporting the “creation and maintenance of contexts”. This includes not only “linguistic” knowledge about concepts and actions, but also organizational context and public commitments.

From a more philosophical point of view, the text can be seen as a world in-between our conversations and the social lifeworld. As such, it may take on a life of its own, and hence obscure the lifeworld. Both Foucault and Derrida have described the process in which we become dominated by the texts that we create ourselves. Language imposes a map of the world, and we only have access to the world through this map. In Cooper’s words (cited in Taylor), “writing is the process by which human agents inscribe organization and order on their environments”. Having done so, they also have to live in this environment.

In the work of Habermas, the distinction between lifeworld and system world plays a central role (Habermas, 1994). The lifeworld is what communicative action is grounded in. But system worlds have arisen because of increasing rationalization, and these system worlds sometimes support, sometimes take over (colonize) the life world. It would be interesting to see in what respect Habermas’ system world corresponds with our text level. In both cases, we have something that exists between communicative action and the lifeworld. Habermas does not discuss text (or writing) as such. He focuses on communicative *action*, and although this does not exclude the use of written media, it is usually interpreted as being a synonym of conversation. Note, however, that the definition of communicative action as coordination of action on the basis of a shared understanding

(Verständigung), i.e the idea of a “common world model”, a shared context, not as a given but as something jointly created, already contains what we call the textual character of communication.

If system world and text have a common denominator, this would be quite interesting since it would show that system worlds are not necessarily opposed to communicative action, but intimately connected to them, as texts and conversations are. It would also show that system worlds are not necessarily as huge as "the market", but also can take the form of a simple text that is agreed upon by a group of subjects. Of course, there is a quantitative difference in impact, but there would be no qualitative difference.

Finally, Ricoeur (*op cit*) has pointed to the the fact that text as “work” may give rise to styles and professional customs concerning this “work”. This is a natural development. But sometimes the professionalism may hinder the transparency and accessibility of the texts. In former days, writing was a skill that only a few possessed, and for that reason, access to writing automatically divided the more powerful and the powerless. Nowadays, computerized information systems, for example web-sites, may become the exclusive domain of technically or perhaps artistically skilled people. Although we abstract from this aspect when we consider the text on the informational or social level, it has to be taken into account when we try to represent or design a real-life communicative situation.

In the next three sections, the explicit context model developed so far is illustrated by three cases: the design of a communication tool, the management of Organizational Memory, and the support of professional communities. The three cases exemplify the three abstraction levels at which context can be studied.

3 Communication: Channels and Shared Workspaces

In section 2, we have distinguished three levels of abstraction. The lowest level is the documentary level, at which we look at communicative objects as represented in some physical medium. When communication is viewed as a conversation (the process view), the documentary level is concerned with messages in the form of speech, office forms, e-mails, letters, etc.. This is sometimes referred to as the channel view on communication (see fig. 1). Messages are exchanged (physically moved or copied) from A’s mailbox to B’s mailbox. Media Richness Theory (Daft & Lengel, 1986) is built on a channel view, as are current LAP approaches.




Figure 1: channel view

Following the context model developed in the previous section, we propose to extend this channel view with a “shared workspace” view




Figure 2: Shared workspace

(Fig. 2). A shared workspace is a place into which subjects can put information, where the information is stored over time, and that is accessible to all subjects or a subset of them. Writing and reading are appropriate metaphors for access to such a workspace. Note that there is no direct interaction between the subjects: the workspace acts as a mediator. The [O] in Fig. 2 refers to the owner of this workspace. This could be A and B themselves, or a third party. An example of the latter case is where A and B are users of a company DBMS.

B may have a shared workspace that they own themselves, which is connected to a shared workspace of the company.

At the moment, communication systems exist that follow the channel view - for example e-mail. We also have systems that support the workspace view - for example, DBMSs or the BSCW system for cooperative work (BSCW). In principle, a combination of the two types is also possible (see fig. 3).

In the context of a large research project on Electronic Commerce, we are currently designing such a system, especially aimed at supporting business negotiation. The problem with using a channel view on business negotiation is that a lot of discussion concerns the contract (from its draft form until it is signed). Sending messages back and forth can easily create confusion about which version is the valid one. It seems much better to put the contract in a shared workspace and let the messages exchanged refer to this

Since in this section we focus on the documentary level, we are not concerned with the information content of the workspace. What is interesting is how different workspaces are connected. For example, A and




Figure 3: Combining channel and workspace

workspace. Not only contracts, but also other background information, like information about the companies involved, or product descriptions, can be put in the workspace so the partners know from each other what they can assume to be mutually known. A similar “two-level” communication system is sketched in (Sengupta & Zhao, 1998). The problem addressed there is the support of virtual teams. The *ad hoc* and transient nature of virtual teams means that they lack built-in mechanisms for the coordination and division of labor. The teams are further handicapped by a lack of continuity of membership. Thus a considerable amount of shared knowledge is lost every time the team is disbanded or reconstituted. The developed prototype system supports both a simple workflow with message templates and common data structures, e.g. a room allocation tool..

In recent years, the Media Richness Theory has been criticized on several points. Ngwenyuma and Lee (1997) have argued for a communicative action approach that takes the institutional context of messages into account. Others have pointed to social factors that influence media choice. From our context model, we would propose to extend the MRT with a model of explicit context. Communication not only takes place in the form of messages, but also by sharing workspaces. The two media levels should be considered in combination.

4 Managing Contextual Knowledge

Contextual knowledge is needed for interpretation of all but the most rigidly formalised propositions, and therefore any attempt at controlled informal communication will have to include ways of dealing with context. Context is a complex notion, heavily dependent on particular situations (*situational*) and thus changing over time (*dynamic*). Therefore, it seems both impossible and undesirable to try and work out and comprehensively store knowledge concerning every concept which might occur in conversation, even within a limited domain. The more context (or rather: *contextual knowledge/information*) is fixed, the less it will be useful for support of dynamic communication, especially if such communication has a (semi)-informal and/or *ad hoc* character. Weigand and Dignum (1997) have argued that indeed, formalisation of communication is a mixed blessing and should be applied only in situations where it actually does some good. Previous sections have elaborated on this idea.

Still, in many cases we do seem to need explicit contextual knowledge, and often it does need to be formal. So far, we have mainly discussed theoretical underpinnings of context (both implicit and explicit) and communication. We will now try and work towards more applied principles and approaches. If we want to address applied matters of contextual knowledge and the explicit-implicit discussion, especially in relation to communication and information systems, a central question is *when* we want to get explicit (creating or using text). Two questions directly follow from the central one: i. *How* we can derive relevant knowledge from existing, explicit knowledge available. This involves such subjects as information retrieval, knowledge- and expert systems, organisational lexicons, and organisational memory. These subjects will not be discussed in depth here; we merely hope to place them in the context of context. ii. *How* can communicating parties reach consensus about explicit concepts, rules, norms. etc. (i.e. their use in context) when the context is not provided by existing text.

In this section, we are mainly concerned with matters at *information level*. We will suggest an approach to contextual knowledge which in general complies with the LAP, also indicating some possible directions for further research. Finally, we will discuss some relevant ongoing research in the area of workflow.

4.1 Contextual Knowledge

The distinction that we have made between explicit and implicit context corresponds roughly with the difference that is often made in Knowledge Management literature between *explicit* and *tacit* knowledge. Although tacit knowledge is often connected with personal experiential knowledge, Nonaka & Takeuchi (1995) draw a further distinction between *individual* and *organizational* knowledge. The latter kind is shared by the members of the organization. The process of knowledge generation is described as a never-ending spiral that involves four modes of converting tacit and explicit knowledge: *socialization* (tacit to tacit), *externalization* (tacit to explicit), *internalization* (explicit to tacit), and *combination* (explicit to explicit). They also stress the communal basis of organizational knowledge (“BA”) (Nonaka & Konno, 1998).

The ambivalence of texts should be taken into account when organizations promote knowledge externalization and knowledge sharing. Following the previous discussion of textual ambivalence, the following remarks can be made:

Writing down knowledge is a creative process that requires some effort. This effort should be recognized as such, and not considered as just “typing in what is already there” (say, in someone’s head). Even an apparently simple task like making minutes of a meeting requires more than a mere transcription of the conversation. It involves active selection, translation from speech style to written style, highlighting decisions and action points, etc. If the knowledge thus made explicit is intended for potential use in other contexts, care must be taken that ambiguities and potential interpretation problems are avoided.

Knowledge is only relevant when it is connected to the communication processes, i.e. ongoing conversations. Knowledge is not an individual affair. For example, when a senior consultant is asked to frame his expertise in a knowledge base for the benefit of junior consultants, this is a reasonable request only if there exists a mode of communication between the seniors and juniors in which the former have the role of coach and the latter the role of student. Alternatively, one could ask the senior to write a report after every finished job, and use this as a basis for evaluation. In this case, the knowledge creation is embedded in a conversation between the senior and her manager, and its use as a part of organisational memory would be secondary. A third possibility is to develop knowledge through a group communication process, e.g. by means of a workshop or Electronic Meetingroom session. Any method can be chosen, but the attempt to share knowledge should be connected to some communication process.

Explicit knowledge (texts) and community are complementary. Consider the following real world example. In a large bank, a Knowledge Management program in a certain department was rather unsuccessful in spite of good will of the people involved and the availability of tools. An electronic forum had been set up in which people could bring in topics or questions, but often there was no response, which of course was rather discouraging. After some time, it was concluded that the forum did not contribute to knowledge sharing. One reason for this failure turned out to be that the people of the department were all located in one and the same room. For urgent (important) questions, they simply walked to the desk of the person in question. For less urgent questions, things felt to be of more general interest, they (sometimes) used the forum. Hence, the forum never achieved much functionality. There was no real need for a textual knowledge level between conversation and community.

Explicit knowledge (texts) and conversation are complementary. This means that as knowledge is put into texts, conversations may be reduced in volume or size. The extreme case is where all knowledge is put in the text, and the conversation has become superfluous: then also the actor has become superfluous, and full automation has been realised. But also in more general cases, some knowledge is taken away from an actor, and comes under the authority of a third party, or at least becomes public property. This may threaten the identity of the actor as a communicative agent.

4.2 Dealing with Text in Context

Explicating context is useful when this is done in a controlled, deliberate, and sufficiently situation-specific way. Whether context should be explicated or not depends on the *goal* of the communicative situation in question. From this functional perspective, the context is not one huge reservoir of pre-existing meaning, but the result of selection. The agent interpreting a conversation or text determines, consciously or not, what background knowledge applies to the act of interpretation.

However, the system world of organisations often depends on explicit symbolic representations for technical if for no other reasons, and therefore what constitutes context is driven by both communicational needs and when/how one is capable of making them explicit. Taking this into account, we can distinguish a sequence of questions which may help us deal with contextual knowledge:

- Which knowledge needs to be shared among participants to achieve successful communication?
- Which subset of this knowledge can be safely assumed to be known by all participants? (and therefore does not need to be made explicit)
- Of the knowledge thus left (i.e. knowledge that must be shared but is not), which subset is already “fixed” in some sufficiently explicit text which can be used to effectively share it?
- If some knowledge is not shared nor captured in some text, how can we produce new text which enables us to effectively share it?
- What subset of this new text can be assumed to be useful for future communication and should therefore be added to the permanent, communal body of text?

Clearly, many more questions arise in the process of walking through the five basic ones given above. What we attempt to make clear by the questions is that sharing knowledge and explicating knowledge are two different but also related activities. Problems can be expected when only one dimension is considered.

4.3 Application Area: Workflow and Organisational Memory

As an illustration of the discussion above, let us consider workflow, which plays a prominent role in many LAP publications (e.g. Medina-Mora et al. 1993). Outside the LAP context, attempts are in progress to actively link knowledge bases to workflow management systems, supporting business processes by providing appropriate, focused information for execution of a specific business process (Reimer et al. 1998, Reimer 1998, Kaathoven et al. 1998). Recently, researchers and practitioners started to promote the idea to explicitly represent knowledge of

the enterprise as so-called *Organisational Memory*. Such knowledge includes the goals of the enterprise, its tasks, its rules, and its resources. (Kaathoven et al. 1998 p2).

Within the LAP, a similar approach can be envisaged, though it would be expected to centre more on communicative transactions (speech acts). Delivering the appropriate information for support of a particular transaction (e.g. a workflow transaction in the ActionWorkflow paradigm) would then boil down to providing information directly related to the context of that transaction, drawn from organisational memory. Depending on the goal and type of the transaction, organisational memory could be applied to give customised access to “rules, tasks, and resources” (also called *business rules*, cf. the GUIDE project), but also, most typically in the case of informative transactions, to more elementary conceptual items, e.g. concepts and terminology.

With respect to implementation, we foresee the use of technology not unlike that supporting the application of Organisational Memory to the execution of tasks in workflow. Primary use for a system supporting semi-formal or even informal communication in a workflow context would typically concern *discussion about work*. In most current workflow management systems, workflow is modelled explicitly, comprehensively, and formally, which leaves little or no room for support of so-called *ad-hoc workflow* (Weigand & Dignum 1997). In the event of an *ad-hoc* workflow transaction, the parties involved will typically have to discuss a job or task since it is not part of a standard, well-defined operation within the workflow management system.

Such a discussion could take place via e-mail. Schwartz (1998) discusses the use of Organisational Memory to improve the quality of e-mail communications. This includes the use of “heuristics for determining common knowledge in communications”. However, Schwartz’s approach lacks a clear mechanism to capture the context-specific essence of transactions, using e-mail messages as a central concept. The LAP framework, with its well defined meta-concepts concerning communication and transactions, seems a much better basis for providing, dealing with, and possibly *negotiating* context-sensitive organisational knowledge. It would enable us to work towards solutions to the problems faced, for example, in the integration of well-defined workflow and ad-hoc workflow.

If concepts as such are being discussed, defined, and managed, it is recommendable to use specific tools supporting the controlled use and storage of terminology or *lexis*. However, in such a case it is of vital importance to allow for sufficiently domain-specific lexical management and to link the *lexical management system* (J. Hoppenbrouwers, 1997) to texts (either expressed in natural language or otherwise). Furthermore, we will have to develop substantial and practical procedures for context-dependent *definition* (Viskil 1994).

In the next section, we discuss a concrete workflow modelling approach that makes a strong connection between communities, text and conversation.

5 Conversation and Text in a professional community

It is not only the case that linguistic acts depend on the knowledge context (informational level), but as social acts, their effectiveness also depends on the social context. When we consider the text (explicit context) not just as an information item but as something that shapes the social world (e.g. the ontologies to which the community is committed, the norms and values that are agreed upon), we arrive at the third level of abstraction, the social level. The RENISYS method described in this section aims at this level. We first give a brief overview of the method, and then we show the role that context plays in specification processes of Virtual Professional Communities.

5.1 The RENISYS Method

De Moor (forthcoming) has developed a method to facilitate the legitimate user-driven specification of network information systems of *virtual professional communities (VPCs)*. These communities are networks of professionals whose collaboration on activities required to realize common goals is mostly or completely enabled by information technology. An important class of VPCs are research networks making use of the Internet to facilitate their work. In order for the members of such a virtual professional community to collaborate productively, it is not sufficient to merely make available a set of standard Internet information tools. Instead, an integrated network information system needs to be constructed out of these standard technological components. The selection and configuration of such tools should be determined by the collaborative context in which the tools are used. In this continuously evolving socio-technical system, active user participation in the specification process is needed. The specification process should be *user-driven* in the sense that it is started the moment one of the community members faces a breakdown, in the sense of (Winograd-Flores, 1986). This is so, because the tasks, organizational structure and network information system of a VPC typically evolve strongly, and new requirements emerge continuously when users encounter new work-related problems.

The members of VPCs are guided in their behaviour by social norms. As these networks are basically egalitarian in nature, specification changes cannot be imposed. The specification process should therefore also be

legitimate in the sense that any specification is both meaningful and acceptable to all members of the community. It is *meaningful* if a common understanding about the meaning of a specification has been reached. A specification change is *acceptable* if the users for whom the particular change is relevant (as determined by social norms) have been adequately involved in the specification process. In this way, all members can see the particular interests they represent sufficiently reflected in the specified goals, activities, organizational structures, and information technologies of the network.

The RENISYS (**RE**search Network **IN**formation **SY**stem Specification) method facilitates this legitimate user-driven specification process. It represents specifications as knowledge definitions belonging to different categories. Knowledge definition changes take place in *conversations for specification*. A Specification Process Model has been constructed in which the conversational moves of the various conversational roles are described. This model is an adaptation of Van Reijswoud's (1996) Transaction Process Model.

5.2 Conversation in Context

Context plays a very important role in situating conversations for specification. Both text and community are important constituents of this context. The philosophy underlying RENISYS can be summarized as: less representation, more interpretation. Explicitly representing specification knowledge in the form of knowledge definitions remains necessary, of course. However, only key specification entities are represented, leaving the lifeworld implicit as much as possible. Thus, contrasted with more traditional specification approaches, there is a shift in context from text to community. This has two advantages: (1) users are not burdened by unnecessary specification efforts, and (2) text ages, but a community is always up-to-date. This is a prerequisite for handling breakdowns, which often occur and need to be resolved long after explicit specification knowledge was produced.

As mentioned above, for a knowledge definition to be legitimate it must be both meaningful and acceptable. The roles that text and community play in conversations for specification are illustrated for both qualities.

5.2.1 Producing Meaningful Specifications

One important category of knowledge definition distinguished in the specification method is that of ontological definitions. These are represented in the form of simple type definitions that express the essential properties that specification entities have. An example could be the following type definition of an edit-process (in conceptual graph notation), which says that an edit process is an activity that (at least) requires a reviewed paper as an input object, and produces an edited paper as an output object:

```
[Type: [Edit:*x] -> (Def) -> [Activity:?x] -
  (Matr) -> [Reviewed_Paper]
  (Rslt) -> [Edited_Paper]].
```

This explicit knowledge definition is an example of a text, which is used in a conversation for specification to produce new ontological and other explicit definitions. However, the conversation not only uses explicit ontological definitions to establish shared meaning. In addition, it makes use of a significant amount of communal (i.e. implicit) knowledge. For example, the users who define the edit process and related specification entities know of many additional things that influence their specification decisions: the role of the edit process in the overall publication process, a basic understanding of editing conventions and rules, political sensitivities, etc. None of these shared complex, subtle, and volatile constraints is or needs to be made explicit. However, communal knowledge both influences the conversation, and in turn is generated by it. Schematically:

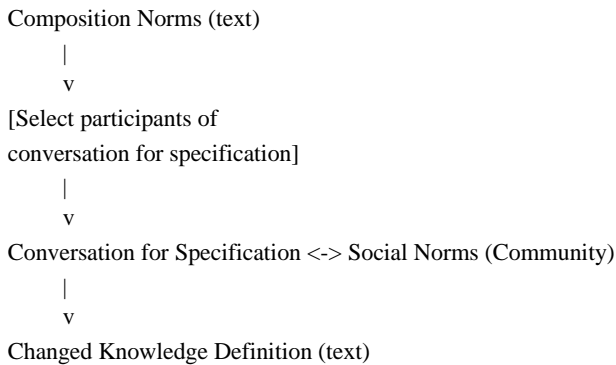
```
Ontological Definitions (text)
  |
  v
Conversation for Specification <-> Shared Meanings (community)
  |
  v
Changed Knowledge Definition (text)
```

5.2.2 Producing Acceptable Specifications

Not only the operational behaviour of networks, but also specification behaviour is guided by social norms (community). In RENISYS, only some of these social norms are made explicit. These explicit norms are called *composition norms* (text). A norm which says that an editor may control the creation of new kinds of review processes would look like this:

```
[Perm_Comp: [Editor] <- (Agnt) <- [Control] -> (Obj) -
[Create_Type] -> (Rslt) -> [Type: [Review]]].
```

Composition norms are used to select the individuals who are to participate in a conversation for specification on a particular knowledge definition change. Composition norms themselves are also defined in conversations for specification, as they are explicit knowledge definitions themselves.



In this way, rational communicative action is supported by the use of texts (normative texts, in the form of specifications, define the space for legitimate conversation processes in the community), and the use of texts is supported by rational communicative action (the texts are created and adapted in a process in which all relevant members can participate).

6 Conclusion

We have introduced a model of context in line with the Language/Action Perspective. Whereas most work within the LAP has concentrated on modelling communicative actions (as processes) against a given social background, we have argued that this social background should be more clearly described as composed of an explicit part -- text -- and an implicit part -- the community --. We have shown how explicit and implicit context play a role in practical applications. In the discussion, it has also become clear that the three elements -- conversation, text and community -- are closely interrelated. We could summarize the relations in the following triangle:

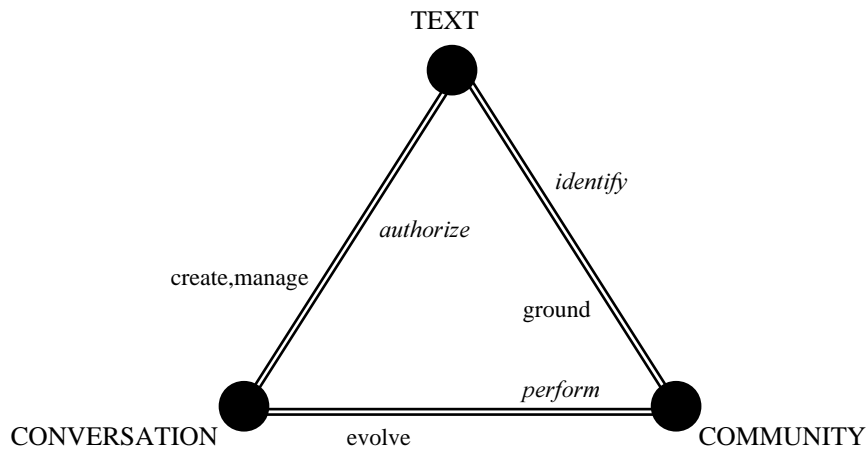


Figure 4: Communication triangle

Communities *ground* texts. Texts are only effective as long as they are *supported* by the community. The community can strengthen its *identity* by means of texts. Communities *perform* conversations, but at the same

time, the community *evolves* and persists in the conversations. Finally, texts are *created, adapted* and *managed* in conversations. However, texts *authorize* conversational actions and can reduce the need to be explicit in the conversation all the time.

The current model should be seen as a starting-point. Though many models of communicative action as a process exist, including ones with formal logical semantics, very little research has been done on modelling texts, and its interdependencies with conversation. Therefore, one of the first goals is to develop formal models of text and on text-creating conversations. We also intend to do more case studies on the balance between conversation and text in practical situations.

References

- Bechtel W. & A. Abrahamsen (1991): *Connectionism and the Mind*. Cambridge, Mass.: Basil Blackwell.
- BSCW (Basic Support for Cooperative Work): <http://bscw.gmd.de>.
- Clark H. (1996) *Using Language* Cambridge Univ Press
- Daft, R., R. Lengel (1986): Organizational information requirements, media richness and structural design. *Management Science* 32, pp. 554-571
- De Cindio, F et al (1986): CHAOS as a coordination technology. *Proc. CSCW'86*, pp.325-342.
- De Michelis, G., M.Grasso (1994). Situating conversatins within the Language/Action Perspective: The Milan Conversation Model. *Proc. CSCW'94*.
- Derrida, J. (1988). *Limited Inc*. Northwestern Univ Press.
- Dietz, J. (1990): A communication-oriented approach to conceptual modelling of information systems. In: *Lecture Notes in Computer Science 436*, Springer Verlag, pp 85-110.
- Dik, S. (1989) *The theory of Functional Grammar, part I*. Foris, Dordrecht.
- Habermas, J. (1984): *The theory of communicative action*, I. Beacon Press.
- Hoppenbrouwers, J. (1997): *Conceptual Modelling and the Lexicon*. PhD Thesis, Tilburg University.
- Kaathoven R., M.A. Jeusfeld, M. Staudt, U. Reimer (1998): *Organizational Memory Supported Workflow Management*. Technical report, Swiss Life Information Systems Research, Zurich, Switzerland.
- Lewis, D. (1969): *Convention: a philosophical study*. Harvard Univ Press.
- Medina-Mora R., T. Winograd, R. Flores, F. Flores (1993): "The Action Workflow Approach to Workflow Management Technology". In: *The Information Society*, vol. 9, pp.391-404.
- Moor, A. de, forthcoming (October 1999), *Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems*. Ph.D. Thesis, Tilburg University.
- Nonaka I.& H.Takeuchi (1995): *The knowledge-creating company*. Oxford Univ Press.
- Nonaka, I., N.Konno (1998): The concept of BA: building a foundation for knowledge creation. *California Management Review* 40,3, pp40-54
- Ngwenyuma, O & A.Lee (1997): Communication Richness in Electronic Mail: Critical Social Theory and the contextuality of meaning. *MIS Quarterly*, 21(2), pp.145-167.
- Pinker, S. (1994): *The Language Instinct*. London: Penguin.
- Reimer, U. (1998): "Knowledge Integration for Building Organizational Memories". In: *proceedings of the 11th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, Alberta, Canada.
- Reimer U., A. Margelisch, B. Novotny, T. Vetterli (1998): "EULE2: Knowledge-Based System for Supporting Office Work." In: *ACM SIGGROUP Bulletin*, Vol. 19, Nr. 1, 1998, pp56-61.
- Reijswoud, V. van (1996), *The Structure of Business Communication: Theory, Model, and Application*. Ph.D. Thesis, Delft University.

- Ricoeur, P. (1986): *From text to action: essays in hermeneutics II* (K. Blamey and J. Thompson, Trans). Northwestern Univ Press.
- Saussure, F. de (1916): *Cours de linguistique gÈnÈrale*. Paris: Payot.
- Schwartz, D.G. (1998): "Towards the Use of User-centric Meta-knowledge in Applying Organizational Memory to E-mail Communications. In: *Proceedings of the 1st Workshop on Building, Maintaining, and Using Organizational Memories (OM-98)*, Brighton, UK –August '98).
- Sengupta & Zhao (1998): Improving the communicational effectiveness of virtual organizations through workflow automation. *Int Journal of Electronic Commerce*, 3,1, pp49-69.
- Simone, C. (1997): LAP based mechanisms for maintaining the contexts of cooperation. *Proc. LAP'97*. TU Eindhoven
- Taylor J.R. (1993): *Rethinking the Theory of Organizational Communication*. New Jersey: Ablex.
- Taylor, J.R. (1996a): Apples and Orangutangs: the worldviews of organizational communication (ms, Univ of Montreal)
- Taylor, J.R., et al (1996b): The communicational basis of organization: between the conversation and the text. *Communication Theory* 6,1, pp1-39.
- Viskil E. (1994): *Definieren: een bijdrage tot de theorievorming over het opstellen van definities*. PhD Thesis, University of Amsterdam (in Dutch)
- Weigand H, F. Dignum (1997): "Formalization and Rationalization of Communication". *Proc of LAP'97*.
- Weigand H. & W.J. van de Heuvel (1997): Ensuring the validity of EC communication. *Proc. of LAP'97*.
- Weigand H. & W.J. van de Heuvel (1998): Meta-patterns for Electronic Commerce based on FLBC. *Proc. HICSS'98*, IEEE Press.
- Weigand H. & S.Hoppenbrouwers (1998): "The Dynamic Lexicon from a Functional Perspective". In: H. Olbertz, K. Hengeveld, and J.M. Sanchez Garcia, *The Structure of the Lexicon in Functional Grammar*. Amsterdam: John Benjamins.
- Winograd, T. & Flores, F. (1986), *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing Corporation.