

# A generic medium model for new media

Lars-Erik Janlert  
University of Umeå  
April 2005

New information and communication technology brings with it new information and communication media. To understand and appraise the latest developments and the emerging new possibilities it is useful to have some conceptual framework for analyzing and discussing important properties of media, and similarities and differences between different media—whether they are old or new, or still in the future to come. This text provides a generic medium model, a simple model of communication, and generalized notions of text and context. The focus is on structural and contextual issues.

## Text

*Text* normally refers to words and wordings, written or spoken (or sung). It is convenient to extend the use of *text* to make it a generic term for any type of simple or complex symbol belonging to some particular symbol system. *Symbol* and *symbol system* are here used in a neutral, “technical” sense (e.g. Goodman, 1968). Some examples of texts in this generalized sense are pictures of various types, diagrams, musical scores, sign language expressions, chemical notation, Labanotation (for dance). When occasionally the everyday, narrow sense of *text* is intended, the term *linguistic text* will be used.

## A generic medium model

The following is a model of information and communication media in general, not restricted to mass media. Roughly, a *medium* is the implemented symbol system of a text (in the generalized sense), including its physical implementation as well as rules and conventions for its use. The abstraction level of a medium may range from very general to quite specialized, including the kind of subspecies that we sometimes call *genres*. Books, novels, crime novels, haiku, action movies, opera, hypertext, TV-commercials, instructions for household appliances, are some examples of media, at varying levels of abstraction. The *domain* of a medium is what the texts in that particular medium are about, the possible topics. The *topic* is what a *particular* text is about (entities, relations, course of events, or whatever may apply for the kind of medium the text occurs in). As terms for a particular text, *document*, *work*, *performance*, etc. will be used; it is difficult to find a plausible generic term. As the most general term for someone who

accesses or relates to a medium, *user* will be used. The domain is more or less determined by the particular medium. Thus, a novel, we know from experience, is a narrative involving a number of people (or at least “characters”) who have desires, character traits, live in certain circumstances, do things, think about each other, have goals and plans and interact with each other in various ways, etc. Sometimes it is also useful to separate from the medium an *interface*, which gives the user access to the text, often a limited view at a time. In a book, for instance, we see an opening at a time; there is a specific layout, choice of fonts, paper quality and binding, etc.

*Medium* is sometimes used in a broad sense that involves the whole culture of its use, as when we say that TV, newspapers, SMS, or blogs are media—sometimes it is used in a more restricted sense, focused on the medium as carrier of information and mechanism for communication, as when we say that gouache, watercolor, oil on canvas, al fresco, are examples of different artistic media.

If we take that more restricted view of a medium, e.g. the book, we will find that the medium in the restricted sense typically has a rather complex structure in which different layers may be discerned. In the case of a book the bottom layer would be the physical book pages and ink marks on the paper (and if you hold the book, you can feel the weight of the book in your hand, you can see and feel the thickness of the pages to the left and to the right of the opening, feel the smell of the paper). On top of this ink-on-paper-codex layer there is a sequence of letters running from left to right, jumping to the beginning of next row down at the edge of the line, jumping to next page at the bottom of the page. These letters form words and sentences in some language with a particular grammar. Larger structures will depend on the type of text. The text can form descriptions, arguments, questions, develop dialogues or whatever; complex structures spanning several pages or the whole book.

Clearly, some of the medium structure is closely related to the physical properties of the medium (the book page, e.g., is limited to two space dimensions), some of the structure has more to do with conventions established for the medium. The structure of the higher layers of the medium has some degree of independence of the structure of the physical substrate at the bottom. They are not wholly independent, however. Lower-level structures make various higher-level structures more or less easy and natural to implement. For example, a basically one-dimensional substrate, like a ticker tape, cannot accommodate a two-dimensional text such as a diagram. (It can, of course, carry a one-dimensional text that *encodes* the two-dimensional diagram.) The physical order of pages in a book, which is linear, although it does not rule out texts that disregard linear order, in practice strongly favors linear structures (on a large scale, spanning several pages) and tends to make readers interpret order as significant even where no order-dependent significance was intended.

There are similar considerations regarding the relation between the structure of the domain and the structure of the medium. For example, a development in time, a journey, a hunt, a fight or game, especially if it is a single thread of events, causes and effects, is eminently well suited to sing about or represent in a book using a linear language like Greek. If the development is more tangled, with parallel actions going on simultaneously, the (normal) book cannot completely *mimic* the time and cause-effect structure. It has to resort to other means, serializing in the medium what is parallel in the domain. In a hypertext medium by contrast, there is no pressure to linearize, and parallel proceedings pose no problem—except, that is, for the reader. Reading in parallel is not something we can do now, and most likely never will be. What is achieved by such a shift from a linear media to a hypermedia is typically a *displacement of responsibility* from the writer to the reader, in this case for serializing parallel developments. Another example of constraints in the relations between domain and medium structure is the impossibility to write an ordinary geography book of Europe where chapters on neighboring countries always are neighbors in the book. Again, in a hypertext medium you could.

One should note, however, that isomorphism between domain (topic) and medium (work) is not necessarily sought for. A straightforward, single-thread development in time can be represented in a book or a film with the text mirroring that order. Yet, what we often in fact find is a partly reversed, broken time-order, as in a detective story that begins with the end. *The very point* of a medium can be to impose a different order, a different perspective on the domain than the “natural,” bringing together elements and events that are not so obviously related in the domain in itself. “Transparency” of the medium is not generally sought for (and it is a mistake to confuse illusion with delusion). But we should be careful here, because the structures of the medium may have important effects on how we perceive the structures of the domain. How natural and absolute are the supposedly domain-natural structures? Can we even think of a domain without using symbols and symbol systems?

To sum up, a particular medium in the broad sense, can be viewed as consisting of a *medium space*, the space spanned by all the possible texts, and a *domain space*, the space spanned by the meanings of all the possible texts. An *interface* gives users access to the medium space (and thus indirectly also to the domain space). A particular *work* in the medium space represents a particular *topic* in the domain space. The structure of medium space is often complex and many-layered. Physical properties have a role in determining the structure, but higher-level layers may be structurally relatively independent. The relation between the structures of medium space and domain space is varied and often not straightforward.

### **Some complications**

Theater is a medium in which the medium space consists of stage, settings, actors, props, etc., and the domain space the vicissitudes of human (or divine) relations, actions and interactions, not unlike the domain of the

novel. Thus, for example, a stage performance of *Hamlet* is (at least at some level) about a Danish prince, his father's suspicious death, Hamlet's anguished thoughts of revenge, suicide, and so on. The drama also exists as a linguistic text, written by Shakespeare, and there is obviously a medium, a genre, of dramatic literature, of which *Hamlet* is an example, which is a *different* medium. The medium space is that of a book, the domain space is that of theater performances: the written drama is about the parts in the play, what the actors say, how they move, when they enter and leave, what the setting should be, and so on. That is how theater directors and actors will use it. But a book reader may very well read *Hamlet* (the book) as a kind of novel rather than as instructions for (or a description of) a performance—or possibly both at the same time. This is an example of how different media may overlap, be nested and interwoven.

Some media do not seem to *have* an obvious domain, for example abstract painting and “pure” music (program music and music with lyrics excepted). Is *The Art of the Fugue* about something? If it is, it certainly is in a sense far removed from the most basic sense in which *Crime and Punishment* is about something (Raskolnikov, a murder, a theft, Ilya Petrovitch and his growing suspicions, etc). There is, however, more to representation than *denotation*, and the domain might be non-vacuous even if it lacks entities *denoted* in the medium. Two other types of symbolic relations are *exemplification* and *expression* (Goodman, 1968). A swatch of cloth may exemplify a certain pattern, color, and texture, e.g.. Expression is metaphorical exemplification. E.g., a passage in a linguistic text can express speed and excitement (without *describing* fastness or excitement) by metaphorically exemplifying speed and excitement: it is an “excited” piece of text with “high tempo,” and this is an intended and indicated function of the text. It certainly appears as if a piece of music can express something like aggressiveness, without needing to refer to an aggressor; the pure expression detached from a concrete origin or carrier.

Within domain space there are also commonly levels. In a motion picture, if a man and a woman are walking through a park, discussing something with each other, that is clearly something taking place in the domain. If the camera in that scene pans from left to right, that is clearly taking place in the medium space. If the camera does not move at all, that is still a matter of using the medium to tell the story in a particular way: for example, the placement of the camera may determine if the couple walk from left to right or from right to left, if they approach or move away. But what if the camera is so-called subjective, moving about as if it were a person looking at the man and the woman, as if we are viewing the scene through this person's eyes? Is that person part of the story, man and woman walking and talking, is he part of a framing story, or is the subjective camera yet to be considered part of the medium (the interface)? It is not uncommon to have stories within stories within stories, and some amount of confusion as to what the levels are and at which level the focus is at a particular point (Calvino, 1989). This is an example of how domain and topic may be

many-layered, nested, and how the border between medium space and domain space may sometimes be difficult to uphold.

### **A simple model of communication**

A rather conventional and naive model of communication assumes one or more, let us say *producers* (speaker, writer, user, etc.), and one or more, let us say, *consumers* (listener, reader, viewer, user, etc). In the communicative process, *meaning* is supposedly conveyed. The producer's task is to choose the meaning to convey, and to construct a text (in the generalized sense) with the designated meaning. The consumer's task is to discover the designated meaning by interpreting the text. This terminology is easy to understand but may give wrong signals about the degree of activity (passivity), constructiveness (destructiveness), etc. on the part of the producer and consumer, respectively. Let us try to disregard such connotations.

One should be aware of the limitations of the *conduit metaphor* that this model builds on (see Reddy, 1979): language as a vehicle carrying thoughts from one person to another, linguistic expressions as containers or packages of ideas that are assembled and unloaded. The more modern approach, rather than taking the text to *contain* its meaning, insists that the text should be viewed as giving *cues* for *constructing* an understanding based on some *common ground*, etc. (Clark, 1966). For our purpose here, this is more a matter of differences of emphasis and perspective, and the simpler model will still be useful.

Linguistics has been criticized for being too “syntactocentric,” for focusing too much on the *product*, the text. But linguists are not the only ones paying much attention to the communication products. The communicative act has been in the process of loosening its strong ties to “the ground,” becoming a self-sufficient, reified product, ever since sounds were turned into clay tablets. The modern information technology is bringing this to its extreme. Communication products are certain to prevail. Rather, we may have to take action to keep *some* communication ephemeral and intangible—“unplugged.”

### **Context**

*Context*, as it is commonly understood, is (1) the parts of a written or spoken statement that precede or follow a specific word or passage. It can also be used to mean (2) the set of circumstances or facts that surround and have relevance for a particular communicative act (*context of situation*). An additional and more general sense is (3) the set of circumstances or facts that surround and have relevance for a particular event or situation of *any* kind. It will be convenient to have a term for that which the context is the context *of*. Let us call it *target*.

First, let us generalize *context* with regard to the first two, textually related, senses to match our generalized meaning of *text*—allowing any type of simple or complex symbols belonging to some symbol system. The target is some particular smaller piece of the text. It has to be a syntactically proper part of the text, that is, it should be a part that is syntactically *well formed*

and a syntactical *constituent* of the text (the first requirement is really redundant, of course.) The strict linear order (“before” and “after”) of the ordinary, linguistic text, needs to be generalized too. Let us just say that the context is the part of a text that *surrounds* the target. The exact meaning of “surround” calls for further analysis and discussion related to the various types of texts and media, which we will get back to. In many cases our intuitions about what belongs to the context, are shared. For example, we can all agree that in *Mona Lisa*, (the depiction of) the face, the eyes, the nose, and the chin, are all within the context of (the depiction of) the mouth.

### **Textual-target context**

It will be useful to make a clear distinction between context that has text as its target, which we can call *textual-target context*, and context that has some non-textual target, which we can call *non-textual-target context*. This third sense of context can be seen as a metaphorical extension of the dominating textual interpretation of *context*. In practice, a basically linguistic, rather detailed and articulated framework is impressed on a wider, pragmatic and considerably less elaborated notion of influence. Note that non-textual-target context is the type of context primarily involved in Context-Aware Computing.

A textual-target context has two parts. One is the surrounding text of which the target is a syntactical part, the *context of text*, sometimes labeled *co-text*. The other part consists of surrounding extra-textual features and circumstances, the *context of situation* (or *situational context*). The paradigmatic example is oral text, like spoken English, and face-to-face communication. Here, the context of situation contains at least one speaker and one listener, a specific time and place, and various objects, events, other persons, etc. in the vicinity of the speaker(s). The co-text of a particular word or a phrase uttered by the speaker is other utterances, before and after the target, by this speaker or by some other speaker in the context of situation.

Textual-target context is a matter of *tokens* (in distinction to types). We are dealing with specific utterances, inscriptions—instances of symbols spatiotemporally located in the world. Spatiotemporal location is a normal property of symbolic tokens in the ordinary, physical world, but what about the digital, virtual worlds of information, like texts on the web, or information in databases? At first it may seem that the difficulty of pinpointing physical location of this or that piece of data, or even coming to a conclusion whether spatiotemporal coordinates can be meaningfully attributed to such virtual stuff (evasive and ephemeral like a thought in your mind), would make digital communication very different. In some respects it certainly does, for instance in the relation between production and consumption (see further below) but the basic consumer (reader, listener, etc.) point of view remains very much the same. It is the spatiotemporal location of the tokens on your computer screen, for instance, that we are talking about—not the myriad of tokens coming and going in this or that computer or connection, which together work to keep a piece of informa-

tion alive in the system. Those tokens on the screen are certainly as well defined in space and time as any old writing on paper or talk in the air.

It is the spatiotemporal relations between the token(s) of the target and the tokens and objects of the context that make the context the context of the target.

### **Contextual influence**

The reason there is an interest in the notion of context, indeed why there *is* such a notion, is the expectation that context (potentially) has a certain *influence* on its target. This is a somewhat simplified and tendentious way of putting it, because it is debatable whether the context changes the target or rather changes its interpretation, or constrains the admissible alternatives for target, etc.

In the case of a linguistic context the usual expectation would be that the context influences the *meaning* or the *effect* of the word or passage it is the context of. In other words, the influence is semantic and/or pragmatic. Some simple illustrations of semantic influence: “white” in the contexts of “white chalk,” “white coffee,” “white wine,” “white people,” and “a government white paper.” Some simple illustrations of pragmatic influence: “can you open the door?” in the context of a speaker carrying bags in both hands, and in the context of burglars in front of a bank vault (the context could be either co-text or situational context).

What about *syntactic* influence? The context, you might want to say, puts a pressure on the target according to the “laws” of grammar, determining what is syntactically permissible at that point. Linguistics, it must be said, would not normally count syntactical constraints as a case of contextual influence.

By definition, non-textual-target context is targeted on non-text, on non-symbolic targets, so any textual notion of syntactic, semantic or pragmatic influence will fail to immediately apply. Also, this context can only be situational; there is nothing comparable to co-text. Naively, the typical physical everyday objects and events are not symbols, they have no semantics and their arrangements do not form expressions in some “language of things.”

### **Anaphora**

Natural language contains some constructs that let the target more or less explicitly *draw on* the context, shifting the perspective somewhat from the context as an environment factor *working on* the target.

One such type of construct is *anaphora*. In the typical case it is the relation between a pronoun and some preceding word or passage in the text, which supplies the referent of the pronoun. In “John threw himself into the water. The poor devil was desperate” the pronoun ‘himself’ is anaphoric to ‘John’, which has a person in the domain as its referent, and this person is carried over as the domain reference of ‘himself’. At least, that is one way of understanding what happens. The target word or passage refers ana-

phorically—horizontally, so to speak, within the medium—to some word or passage. The domain reference of that word or passage becomes the domain reference of the target. Certainly there are some cues in ‘himself’ that help us locate some appropriate part of the context, but in the more general case it may not be enough. Anaphoric reference in the general case would seem to rely on the part of the *topic* that the preceding text refers to.

From the reader perspective we can see this as a gradual construction—and revision—of a mental model of the topic, helped by the text, from the beginning up to the target. The writer has rather the opposite task: that of putting in words a pre-existing real or mental model of the topic. We know of course that the writing process is often much more complex than that. Writing, and reading what has been written so far, are intermixed, and in the middle of the process an author of fiction is reportedly often as curious to know what will happen with the characters in the story as any future reader will be.

Instead of referring to the preceding text anaphoric reference can be anticipatory and relate to some word or passage in the text following. E.g. “Himself an expert swimmer, John jumped into the water.” In this case, some prefer to speak of *cataphora*. (The reason is that, etymologically speaking, *anaphora* means ‘carry up’ (in the text); *cataphora* means ‘carry down’.) The basic mechanism is supposed to be the same, but the idea of treating anaphora and cataphora as completely symmetrical is debatable.

The notion of anaphoric reference extends beyond the simple case of pronoun targets; for example ‘the poor devil’ in the first example is a somewhat more complex anaphoric reference that eventually ends up with John, again. Or so it would seem. Given that the example reveals the *whole* context (an unusual piece of prose, to be sure), it certainly does. Consider, however, this alternative context: “Bill was having serious trouble keeping Albert’s head above water. John threw himself into the water. The poor devil was desperate.” Now, it seems there are three about equally uncertain choices of who ‘the poor devil’ is. More context still, might (or might not) resolve the matter.

Let us not forget, however, that in an ordinary situation of communication, the producer of the text (speaker, writer, author, or whatever) is supposed to follow H. P. Grice’s *cooperative principle*, which includes being precisely as informative as is required, being authentic (not misleading the other), and relevant (Grice 1975). That obscure texts may have obscure contexts, is obvious and not the issue here. Then again, a certain amount of obscurity and ambiguity is sometimes actually intended. Besides producing certain aesthetic or pragmatic effects, one may argue that linguistic precision and clarity should not exceed epistemic.

In the above example, we also get to recognize the *non-monotonicity* of contextual influence. An outer context may change the meanings and effects conferred on the target by an inner context. It is easy to come up with (rather contrived) examples where the ultimate meaning of an arbitrarily



large text is abruptly upended by some additional phrase like “The following is completely wrong,” or “Thus ended Alice’s dream,” or by the fact that the whole performance is taking place in a virtual reality game, etc. This may have important implications not least for machine context acquisition in issues such as continuous versus event-driven context pickup. If we start in the middle of things, when dare we stop context search and declare that the meaning is reliably established?

### **Deixis**

*Deixis* is another common type of construct that works to add or specify meaning by pulling it in from the context. Deictic words and passages like *me, you, them, your house, there, this, the former, now, tomorrow morning*, point to elements and circumstances in the context of situation (or sometimes in the co-text) to fixate their meaning. These constructs are also referred to as *indexicals* (particularly in philosophical contexts). What person ‘I’ refers to depends on who is the speaker. What time ‘now’ refers to depends on when the word is uttered.

Other context effects do not depend on any explicit or half-hidden syntactically manifest reference to the context. The co-text can be used to put the reader in a certain mood and so view the target through specially colored glasses; the text can be designed to subtly exploit the interests and prejudice of the listeners in the situational context; etc. Some context dependencies of this kind are studied in *rhetoric*. Syntactically manifest but non-referential use of context is also a possibility. Consider e.g. a text in which a sequence of page-long sentences is terminated by a three-word sentence.

### **Push or pull**

The context can be viewed as an external power exerting influence on the target; alternatively as a resource which the target can draw on. Which perspective is the more appropriate? Not surprisingly, it largely depends on from which end you are viewing the communication: the production or the consumption.

If you are producing a text, the pull perspective is appropriate. The use of anaphora and indexicals can be seen as convenient shortcuts in the production process. By drawing on context, the situation as well as the parts of the text already constructed (and also, to some extent, planned but not yet constructed text), the production process as well as the product are simplified—probably also to the benefit of the consumer and the consumption process. At the consumer end, the push perspective is the more appropriate. The consumer is not at liberty to *choose* to make use of context, rather the context *forces* shifts or additions of meaning that the consumer must heed to (or suffer the penalty of missing the point). With a more relaxed view of the power of producer to determine meaning, the sharp division of responsibilities loosens up. The consumer finds/constructs meaning the producer didn’t know was there, wasn’t aware of, did not intend, etc.

With regard to the *syntactic influence* of context (granted this little terminological transgression), the situation is different, almost reversed. Context

cannot change the syntactical properties of a given word (or so it would seem), making push irrelevant to the consumer (pull remaining as irrelevant as before). For the producer, the push perspective is definitely less out of place than the pull perspective. The text obviously needs to be constructed under the constraint that syntactical properties of words and constituents are grammatically consistent with their contexts. Normally, one would expect this to happen largely (or completely) as an effect of generative rules that simply have the property that text generated by them complies with the grammatical restrictions.

But, of course, anything that constrains what can be produced is potentially a useful guide to effective consumption. For instance, syntactic context may help disambiguate homonyms. And why isn't this push? The received view is that *mouse* (noun) and *mouse* (verb, "to use a computer mouse to click, drag, etc.") are two different words. Alternatively we might say there is just a single word *mouse* (noun perhaps, or indeterminate); the context may change or further specify the syntactic properties.

It is easy to get the feeling that context is a kind of ragbag category where you put things that do not fit neatly into syntax (or semantics). Syntax is supposed to be *neat*. For example, why is the fact that we talk of "school of fish," but "flock of birds," and "herd of antelopes" a matter of context and not syntax proper? Presumably because it would make grammar too messy and/or the flora of semantic markers of lexical entities too wild and unwieldy. That would also explain why "syntactic influence" is not considered a matter of context: syntax takes care of this by itself very adequately.

With regard to *social context*, which is part of the situational context, it seems to be uncontroversial that there is a push on the producer. For example, social context determines whether to use *vous* or *tu*, say *Shut up!* or *Would you please be quiet!*, the proper way to address a Japanese emperor, the proper voice level for a theater audience or for a noisy machine hall, etc. The pull perspective also remains relevant for the producer—as with situational context in general. This is an example of production constraint that *is* considered part of linguistic context theory.

In a novel we expect the characters to follow the same rules of socially appropriate linguistic behavior under the particular circumstances of the situation. The social context of a book is, first, the social situation of the author, second, the social situation of the readers. Where then, is the social context of the persons in a book of fiction? It seems that the *situational* context in the domain in which the characters of the fiction talk and write (and *think*, because unlike in real life we can have access to people's thoughts—and they look very much like sentences in natural language!) has been transformed into the *textual* context of the novel.) The issue of context constraining production may have interesting parallels with non-textual-target context. Things and actions may be out of place.

### Textual surroundings and discourse types

You are not supposed to begin reading a novel at some random page, read a few sentences, pick another location at random, read a few sentences, and so on—and if you do, it is at your own peril. The novel is produced on certain assumptions about proper consumer behavior. And the novel is consumed, obviously, on certain assumptions about proper production procedures. These conventions are part and parcel of the medium. As noted earlier, they are not completely independent of physical medium properties. Speech, for example, does not leave the listener much choice as to the order of consumption—until, that is, the text is recorded and you can play the recording at your own discretion. Then it has become a different medium.

In the novel as well as in any ordinary narrative linguistic text, the particular convention is of course that you start at the beginning and proceed forward without skips towards the end. That means that the (textual) context at each point is the (entire) text preceding the current target. What about *cataphora*, anticipatory anaphora? Normally they will have quite short range, often within the same sentence, or the next. This may generally (statistically) be the case also with pronoun ‘upward’ references, whereas anaphoric reference in a wider sense can reach arbitrarily far back in the text. Also, anaphora and cataphora are not symmetric with regard to production and consumption conventions. Suppose we read a detective story, and the first chapter describes the preparations or evil deeds of some perpetrator, always using the pronoun “she,” no names named. In one sense we understand the meaning of this chapter. In another sense, we will eagerly read the rest of the book expecting to be filled in about who this person is. In the middle of the book it might be Irene or Agatha or some other character of the story up to this point, named or not, or it might be some other person, not yet introduced (depending a little on the exact subgenre). When finally we are able to make the match, of course it is true that the first chapter gets additional meaning, but—and this is the point—*it is not* the meaning that the author intended us to have when first reading the first chapter.

*Linear discourse* is the dominating discourse type. Narratives (whether in book form, as movies, oral traditions, theater performances, computer games, etc.), arguments, how-to-do instructions, mathematical proofs, and more. In linear discourse then, the context is *the section of the text that precedes the target*—in time or in space, depending on the medium. Proximity also plays a role. Although parts of the context far away from the target can have profound influence on the final interpretation, without taking the more local context into account we are sure to miss the point completely.

There are other common and well-known discourse types very different from linear discourse. One example is what we could call *pictorial discourse*, typical of pictures, or “pictorial texts” in general. The (textual) context of the target is 2-dimensional (or 3-dimensional), and there is *no* linearity (until we put the pictures in motion, and get a linear time dimension); we are not supposed to “read” a painting, a photograph or a diagram in any

prescribed order. There may be syntactically relevant relations, like “enclose,” “touch,” “above,” etc.—some of which may generate at least partial orders, like “enclose”—but there is no designated linear order directing “reading” and production conventions. All we can say in the general case is that proximity remains of some importance.

Another common discourse type we could call *hierarchical discourse*. Hierarchical discourse in linguistic text is sometimes signaled by decimal-numbered headings and subheadings (1.2.1, 1.2.2, 1.2.2.1, etc). Typically the hierarchy is a specialization hierarchy: subsection 1.2.2.1 is a more specialized elaboration of (the top-level text of) section 1.2.2, which in its turn is a specialization of 1.2, etc. (*Within* a certain subsection, the discourse will often be linear.) In hierarchical discourse, the context is the chain of texts on a path from the root down to the target; in other words *the part of the text that precedes the target in the hierarchical order*. Hierarchical discourse may be appropriate for analytical, non-narrative, documents, such as organization descriptions, complex systems descriptions, etc. It has the advantage of limiting the amount of context you need to consider as a reader interested in some particular topic contained in some certain subsection of the whole text. An extreme special case of the hierarchical discourse is the completely *flat discourse*, as in a telephone directory. Each line is self-contained and unrelated to the preceding and following lines; the only context is the area code.

If linear discourse puts severe restrictions on both consumer and producer which in effect make contextual relations from the consumer perspective quite reliably track contextual relations from the producer perspective, and vice versa, the *hypermedial discourse* permitted by the new hypermedia represents an opposite extreme. In the general case, the target may be preceded (in the hyperlink sense) and followed by any number of different texts. The consumer is not expected to walk all possible paths through the maze. Neither is the producer supposed to divine where the reader of the target will come from and where he is heading. What kind of context is that? You might say that the context is non-deterministic.

### **Proximity and continuity**

We note that *proximity*, closeness in space and/or time, plays a key role in determining what may at least potentially belong to the textual context in the most common media of today, and at least to some degree also as a rough, initial ranking of contextual importance. The closer it is, the potentially more important it is. “Space” is here not simply physical space as such, not even the subspace defined by the physical bottom-layer of the medium. In a novel, for instance, it is not a matter of measuring the physical distance between one word and another, because that would make words occurring just behind this word, on the back of the page, or on the following pages, very close. Neither is textual proximity measured as distance within the page: a word directly above in the preceding lines is usually not textually very close to this one. The relevant distance and proximity measure is in the *textual space*, and must be defined in terms of the basic

structural rules and constraints of the medium in question. In considering other varieties of textual space, we can maintain proximity as the general principle for determining scope and ranking also for new media like hypertext.

The above discussion has assumed that contextual effects work primarily through the medium, and that consequently proximity is proximity in the medium. But suppose the domain rather than the medium is the fulcrum of contextual influence. Then the relevant sense of proximity would rather be proximity in the domain: referents *close in the domain* could be quite distant in the medium. Compare the earlier comment on the possible mental mechanism behind anaphoric reference.

It seems likely that the domain has a contextual role to play that cannot efficiently be explained in terms of the medium. That opens up a range of possibly relevant proximity relations. For abstract domains, domain proximity may turn into conceptual closeness (relatedness)—one aspect of which should be similarity or likeness, other aspects might be contrariness, constituency, specialization, etc. For non-abstract domains, when proximity does not simply mean spatial and temporal closeness, it may translate into all kinds of relational connections: cause and effect, father and son, goal and plan, etc.

Proximity is more or less *part of the notion* of context. There seems to be a quite strong assumption of locality at work, the proximity principle:

*Things that matter are close.*

*Things that are close matter.*

It is a simple enough principle. It may not be entirely correct, but we seem to operate under some such assumption, generally. But it makes you wonder if there could be principles other than that of proximity that might be used to define contextual scope? Probably not: the locality assumption and the notion of proximity are so strongly interdependent that a non-proximal concept of context is a contradiction in terms. *Metaphor* and *metonymy*, that is, likeness (similarity) and closeness, are perhaps the two most generic ordering principles known to us. The obvious choice in trying ‘the opposite’ with regard to context and proximity would be to substitute closeness with likeness. The result would be some sort of “magic” correlate to context, that would be based on the principle that what matters is similar, and that which is similar is that which matters. We would, it seems, get something like the mystical correspondences of microcosm and macrocosm, between the life of humans and the life of stars, between resembling parts of nature, etc.—as found in voodoo, hermeticism, paracelsism, alchemy, sympathetic magic, etc.

One consequence of the proximity principle is that we may take for granted a certain amount of *continuity of context*: what is potentially relevant context for a particular target, will largely also be of potential relevance for other targets close to this one. If we choose to view the context of a target (the potentially relevant context) as a mathematical function of the

target, this means that the context function is continuous. But note that this is just the *scope* of the context. When we go from one sentence to the following in an English text, we must be prepared that for instance the pronoun references may jump around, although within largely overlapping ranges.

### **New contexts of situation**

Also with regard to situational context, proximity appears to be the general guiding principle for delimiting scope. Again, we cannot simply equate proximity with physical closeness. One factor, for instance, would seem to be visibility. The moon may well be in the situational context although it is physically very distant, whereas something in an adjacent room might not be.

New media make the notion of context of situation more complex, however. Already the situational context of a written text introduces complications compared to the primal speaker–listener situation. There are on the one hand the circumstances of the production of the text. There are on the other hand the circumstances of the use of the text. Production and consumption can be widely separated in space and time. The circumstances of consumption are less predictable for the producer, and the circumstances of production less knowable for the consumer. Multiple consumers in *different* contexts of situation is also a common enough case.

It is the movability of the text (hinging on its reification) that creates these complications. The text is no longer stuck in one unique situational context.

How can the consumer know, how can the producer ensure, that a written text like “Push the handle” is consumed in the appropriate circumstances? One type of solution is to nail it in place. Another type of solution is to provide instruction how to move to and/or set up and generate the right setting for proper consumption. For some texts an imagined or virtual context is as good as a real; then the situational context requirements can be converted into textual context requirements.

This separation and multiplication of situations is just another factor leading in the direction of *desituating* text, of decoupling it from situation, making it more a product, weakening the *control purpose* of the communicative act (which is so prominent in *speech act theory*).

When technology develops, the locality assumption becomes more and more tenuous, or rather, our concept of locality is transformed: the phone brings distant people close. The real problem facing us then is that when everything gets close, context becomes universal and proximity a useless tool for selecting a manageable slice. The dramatic expansion of our awareness and closeness horizon means that there is shorter notice to prepare for oncoming changes, there are more factors to count with, and the problem of “satisficing” is getting harder and harder. Add to that the increased mobility of objects and persons. Contexts will be in a constant flux. Not just the apparent and rather innocuous flux generated by self-

movement, but flux generated by the constant rearrangement of objects, people and information.

The common-sense definitions of *context* make a somewhat vague provision for “relevance” in delimiting what constitutes the situational context. In some authors’ view context can be anything, anywhere, anytime, that the consumer uses as a help to understand the text. That is a lot. Taken literally, it means either that context becomes a rather useless subjective concept, or even more useless, equal to the entire universe, every time. Our epistemic space is brought to deviate more and more from physical space by the invention and use of new information technology; our “natural” awareness and activity spaces do not match these new demands.

### **Adding meaning to the world of things and events**

The world is not primarily a message. The world may contain lots of purpose, but purpose is not equivalent to message. One may argue, however, that our everyday world is actually in the process of *becoming* messages. Increasingly, we do something (concrete) and we articulate a message in one single act; we make some (functional) artifact and create symbolic meaning in one and the same object. We should probably help this development on (but take care to control it and design it well), because it may save us considerable effort and frustration, making decision-making an integral part of our normal, everyday tasks (Janlert, 2001). It is not an altogether new situation. When I veer to the right to allow passage for someone coming the other way, I am well aware that this is *both* a physical, evasive move, *and* that the other person will (normally) take notice of my maneuver and act (or refrain from action) accordingly. My physical action is at the same time a message that might be paraphrased as “I’m going this way now (so don’t you do the same).” In a world where each object, event and action can in principle be traced and observed from any other point, each act is necessarily potentially also a message, and the actor must take that into account. To be successful in such a world, an actor will be hard pressed to design his actions as much from the communicative point of view as from the physical point of view.

Viewed from the opposite end, it is rather obvious that communication is being increasingly reified, so that you are simultaneously creating “things” in the process. Even though the newest communication products often tend to be not very tangible, they share other properties with material things, such as permanence, objectivity, combinability (after the fact), searchability, processability, etc. Moreover, a considerable proportion of what certainly looks like messages (issued by human beings) are, because not mediated by any person, in fact “real” actions (and not just speech acts). They really turn on the TV, shut down the chemical plant, set the missile on a new course, etc. The demands for the physicality of “real” actions are also lowered (again without people starting to think in terms of speech acts) as people perform more and more of their work in the digital worlds. It seems we are now thinking so loud, we are meaning so intensely, that the difference between thinking and acting is dissolving.

And so the border between the meaningful and the meaningless, between communication and action is slowly beginning to evaporate. There are obviously still non-symbolic objects in the world, and there are still, and will probably always remain, very non-thing-like communications. Still, the world will become more and more meaningful, events will more and more be interpreted also as information, as messages.

### **Degree of authoring**

Different texts, different media differ in the degree to which they are designed and produced with a purpose of being interpreted and understood in a particular, intended way. Novels and movies are strongly authored. An improvised TV interview, somewhat less. If we just put up a fixed TV camera (no cameraman, no cutting, panning, zooming, etc.) and people in view are not aware they are being filmed, we have very low degree of authoring (still, there is the decision of when and where to put the camera). As soon as people become aware of the camera they become actors and co-authors. When it comes to physical actions and objects, complete absence of authoring is becoming less common. Very strong authorship may be declining too, if new media become popular. Then there is the current vision of *mobility*: the more movable entities there are around, the more problematic it will be to carry through a fixed design.

### **Contextual strategies**

There are two possible general approaches for acquiring context intelligence. One is *walking in*, approaching the target from without (in terms of some proximity measure). The other is being *thrown in* right at the target, then exploring outwards with the target as starting point. Walking in works best for linear discourse and hierarchical discourse. In other types of non-linear discourse there are several paths to the target, maybe even an infinite number, so a straightforward walk in may be less successful. *Sampling* the context will anyway have to suffice in cases where context is rich and many-dimensional, and time and cognitive resources are limited.

Then there is the closely related but conceptually different issue of whether context pickup is continuous (data driven) or on demand. The walk-in approach can hardly rely on context on demand; we don't know what to ask for before we know the target. A general observation here is that we may not be able to decide that this or that aspect of the environment will be or not be an important piece of the context for a so far unknown target. Certainly there are media, such as certain genres of movies, with conventions about how to signal that some event or object will prove important for something that will happen later (what that will be, the viewer doesn't know yet). Less authored, less orchestrated situations do not supply such cues.

The thrown-in approach, by contrast, strongly suggests demand-driven search for context. Still, this is compatible with a more continuous mode of working: we are thrown into a particular situation, but instead of directly attending to what is immediately confronting us we begin by taking in the



environment, savoring the ambience, etc — *then* we focus on what is before us.

If we assume contextual continuity, the walk-in approach and continuous context pickup will be a rather efficient way of solving the context acquisition task. “Context awareness” could then perhaps be identified with a continuous background process of building and maintaining context knowledge without prescience as to what it might prove useful for.

We should, however, probably prepare for operating largely in an event-driven manner because we, and our objects and applications are going to be *thrown* into one situation after another. Modern transportation (e.g. elevator, airplane, metro) throws you from one place to another; turning on the TV (or zapping between channels) throws you into the middle of a film, a news report, a documentary, a docu soap, or what is it?; interrupts from phones, e-mail, advertisements, etc. throw you into new situations; searches on the web throws you into new, unknown material; etc. Backtracking is becoming an increasing problem from the human point of view. Since we have limited capacity to keep a stack of pending, interrupted activities, we tend to arrange for events and signals that will bring unfinished tasks to the fore again. That makes us even more event-driven, to the point that when we have finished a task and there are no signaling events, we do not know what to do.

At some points, we might arrange to have context services or context servers available, providing quick updates of where we (or our computer applications) are, who these people are, what is going on, etc. In the information context, *meta information* may be seen as a way of meeting this demand for compact context reports. Having been at a particular target before, we (or our applications) may be able to *recall* the context from contextual memory. Instead of extensive context exploration assuming no previous knowledge of the context, it may be possible to pick up context cues that trigger recall of stored context knowledge, at least usable as a first approximation of the actual context. Large parts of that might be stereotypical. It seems that Context Aware Computing at this time is primarily operating under some such assumption: pick up some fairly simple cues, on the basis of which some stereotypical context is inferred and selected (from a fairly limited number of anticipated contextual possibilities.)

### Literature

Aristotle (c. -350) *Poetics*. <http://classics.mit.edu/Aristotle/poetics.html>

Bakhtin, M. (1981) *Dialogic Imagination*. University of Texas Press, Austin [in Swedish: Bachtin, M. (1991) *Det dialogiska ordet*. Anthropos, Gråbo.]

Bloomfield, L. (1967) *Language*. Unwin University books (first edition 1933).

Bruner, J. (1986) “Two modes of thought,” in *Actual Minds, Possible Worlds*. Harvard University Press, Cambridge, Massachusetts.

- Calvino, I. (1989) "Levels of reality in literature," in *The Literature Machine*. Pan Books, London.
- Clark, H. H. (1996) *Using language*, Cambridge.
- Context-Aware Computing (2001) Special Issue of *Human-Computer Interaction*, vol. 16, nos. 2-4, 2001.
- Csikszentmihalyi, M. & E. Rochberg-Halton (1981) *The meaning of things*. Cambridge University Press, Cambridge.
- Goodman, N. (1968) *Languages of Art*. Bobbs-Merrill, Indianapolis.
- Grice, H. P. (1975) "Logic and Conversation," in P. Cole and J. Morgan, eds., *Syntax and Semantics*, vol. 3, Academic Press.
- Janlert, L.-E. (2001) "Postmodern decisions," in Å. Grönlund, ed., *Electronic Government – Design, Application, and Management*. Idea Group Publishing, London.
- Reddy, M. (1979) "The conduit metaphor—a case of frame conflict in our language about language," in A. Ortony, *Metaphor and thought*, Cambridge University Press, Cambridge.