Thoughts About Discourse Comprehension

Nils Erik Enkvist

What is Involved in Understanding Discourse?

There are two very basic problems buried under my question. The first has to do with understanding and comprehension. The second problem is the meaning of 'discourse'. It has often been convenient to make a distinction between *text* and *discourse*. One of the common ways of doing so is to define *discourse* as *text plus situational context*, and conversely *text* as *discourse minus situational context*. To cite once again a classic example: if there is a sign saying 'No Smoking' on the wall of a lecture room, the sign qualifies as discourse. But if a linguist takes down the sign and puts it on his desk in order to analyse its syntax, it turns into text. Those fond of rigorous definitions might object that even on the linguist's desk, 'No Smoking' has the context of a syntactic example. We can satisfy the practical linguist, though not necessarily the philosopher, if we define *situational context* as a *context of basic function*, that is, the context in which the utterance was originally used in real-life communication. 'Was originally used', here, may be modified into 'was or could be used' to allow for invented, imaginary examples. But we might add that discourse linguists generally try to avoid inventing examples. They prefer authentic materials to inventions of armchair data where the linguist fancies weird sentences and passes subjective judgements on their linguistic status.

My 'basic function' also turns out to be an oversimplification if we start thinking, for instance, about people who quote texts such as the Bible or classic poems to support an argument. Problems of allusion and quotation have been studied at great length, but I must here withstand the temptation of discussing such multiple contexts and contextual transfers.

There are other basic terms that may need explication. One of them is *communication*. Communication can be broadly defined as the transfer of information from one individual or group to another. A *natural language*—a spoken, and perhaps also written, language which can be learned by children through exposure in their immediate environment—is not the only code of communication. Non-linguistic semiotic systems such as traffic signs also communicate. The same is true of artificial languages patterned on natural ones, such as Esperanto. They supposed to be easy to learn for international communication. Symbolic languages such as formulae in chemistry or mathematics or formal logic try to maintain a one-to-one relationship between symbol and meaning, and to avoid polysemy and figurative use. Substitute languages such as sign languages for the deaf, comprise of concepts which are to some extent patterned on those of natural languages but their different expressive machinery has given them a life of their own. Ideographic writing systems such as Chinese, again, functions in such a way that the same set of written symbols can represent different forms of speech. None of these semiotic systems or languages need remain absolutely constant. True, there is frozen, ritual language, in prayers for instance. And Dr. Zamenhof, who developed Esperanto, decreed (in vain) that nobody was to change the basic rules he had stipulated for his language. Still, the relentless development of society and language tend to compel changes even in frozen texts: the Lord's Prayer for instance has been given different forms, and even in its shape of 1611, its meaning will not be quite the same for us as it was for, say, Shakespeare.

But natural languages in the narrow sense—here, human codes for communication exposed through speech and perhaps also by writing or electronic media, or through a combination of all three—are actually only one type of the natural codes through which meanings are transferred. To begin with, we might distinguish between language and paralanguage. Language is structured in ways amenable to description in terms of today's linguistics, whereas paralanguage so far defies simple structural description. Among today's paralinguistic features we often include matters such as tempo, speed, pausing patterns, voice colour (husky, strident, strained, etc.), loudness, and the like. The history of linguistics teaches us that features once regarded as paralinguistic, such as intonation, may be amenable to descriptions in linguistics proper, once their systematic and structured use for the conveyance of meanings has been discovered and described.

Beyond paralanguage are inarticulate noises, such as cries expressing pain or chuckles expressing happiness. Still, their meanings are often situationally describable. If we make a noise when our dentist drills into a nerve he will know what we are trying to signal. Meanings of inarticulate noises can be shared by groups speaking different languages. Still, we should beware of calling all of them universal.

So, there is a vast number of codes of behaviour which carry meaning. In some cultures you are expected to remain standing when you enter somebody's dwelling, in others you must sit down at once to avoid offence. In some cultures, bad news, even of the death of one's dearest, should be told with a smile. Doing the wrong thing signals, at best, ignorance. In many cultures, ignorant foreigners, identifiable for instance by accent or dress, are more readily excused than natives. If no excuses exist, such aberrant behaviour turns into bad manners or even wilful offence. When cultural codes change with time, there may arise conflicts between generations, the old conservatives and the young innovators. In some cultures you take your shoes off when visiting, in others you keep them on; I know of one particular society (namely my own) where older people accuse the younger generation of bad manners when the young move around in stockinged feet. Those who wish to rationalise this particular generation gap will say that the young are used to children with muddy boots. Generally, traditions and fashions of clothing are a constant source of discussion and protest. In many countries of Western Europe you can guess a person's age, or at least solidarity with a given generation and perhaps political views, from whether that person takes part in a TV discussion in T-shirt and jeans or in a chalk-stripe business suit, shirt and tie. Such features, and countless more, all carry specific meanings for those who are familiar with the traditions of the society in question. And those who move between cultures and between traditions have excellent opportunities for causing offence, even unwittingly, if they do the wrong thing.

This presents us with another troublesome basic concept, namely *meaning*. Here I must cut an endless discussion very short by simply suggesting, with the structuralists, that meanings presuppose potential contrasts between alternatives from among which a speaker or writer may choose. If there is no choice and no alternative, meanings cannot arise. If there is a society where everybody always drapes herself in the same way in similar animal furs, contrasts in clothing cannot be used to express meanings such as social status. What is completely predictable cannot carry meaning. In Modern English spelling, as information theorists like to note, the letter *q* is always followed by *u*. Hence the *u* has no meaning and could be left out. If we stop worrying about traditions we might just as well spell *qeen* and *qote* instead of *queen* and *quote* without losing anything essential. Note that one highly meaningful contrast is that between speech and silence. Some cultures tolerate and even value silence, in situations where other cultures regard it as an insult.

There is, however, another consideration involved. If we reduce linguistic signals to the minimum needed to express contrasts and thus meanings, we at the same time increase the risks of mistakes. When working with accounts for instance, we cannot afford errors: every number must be absolutely right. In natural-language discourse, however, there is a certain amount of redundancy, of 'unnecessary' information, as well as a set of devices to express hesitation and to correct what went wrong. Therefore we should beware of too readily defining what may seem like unnecessary, redundant repetition as needless and communicatively irrelevant. Some linguists have cited structures such as the Latin *omnium bonorum virorum* with a genitive ending attached to every word as an example of needless bother. But in fact this is not necessarily so. Repeating a case ending shows that the words belong together, and makes it possible to separate them from each other for instance in poetry, as the Romans indeed used to do. And the Finnish *suuressa vanhassa valkoisessa talossa* ('in the large old white house') the repetition of the inessive *-ssa* not only cements the words into a phrase, but also signals that the whole phrase has the function of an adverbial in its sentence.

Concerning communication, again, most messages should presumably be interpreted in their literal sense. But many are figurative; yet others ironical and carry underlying meanings opposite to those on the surface. Ironical utterances should be provided with irony signals, either linguistic or situational-cum-pragmatic: if I say 'what a lovely day' when it is raining cats and dogs, people who know me will realise that my statement can hardly be literally true. I shall not here concern myself with the fine line distinguishing irony from a lie, however interesting this problem may be, not least in contrasting communicative traditions in different cultures.

All this implies what one might call an optimistic teleological stance. Societies, their conventions and their languages have been a long time developing. What is feasible and what works is at a premium, what is awkward and risks communication failures has often been lost by erosion in the self-therapy of language. Note for instance that the principles of economy are different in natural languages from what they are in symbolic systems of, say, logic or mathematics or bank accounts. Messages could be shortened if they were expressed in a language with less built-in redundancy and if speakers and writers avoided all needless constructions and all repetition. But as we have already noted, the cost of such apparent economies would be a vastly increased risk of misunderstanding, comparable to making errors in a bank account. Therefore, simplistic attempts at applying criteria of economy to natural languages and to natural-language communication are bound to be unrealistic. Relating the efficacy of a grammar to the number of rules (the fewer the rules we need, the better the description) is also hazardous. The outcome depends on the form of rules. And then if we simplify our description of a language at one point, it is apt to expand at another. No description has been achieved of any language with the completeness needed for a total overview of descriptive economy.

It is yet another matter that the economies of models of language may well be construed on principles that differ from those of the actual neurophysiological processes of the human brain. There was a classic debate between behaviourists, who claimed that we learn language by imitation and reinforcement of success, and linguists of the generative-transformational persuasion, who claimed that utterances arise through grammatical rules made possible by an innate human linguistic capacity. A sceptic might still ask to what extent people actually learn certain common utterances (say, greetings, sentences occurring frequently in weather reports, and the like) by direct imitation. A person who uses certain utterances very often may well insert them into his discourse as ready-made prefabs, instead of going through the bother of generating them anew every time from scratch. Phrase-books for tourists for example try to collect utterances potentially useful as prefabs.

Underlying all such considerations is the obvious fact that man is an animal in constant quest for meanings. Whatever result of a potential choice we hear or see or feel we tend to interpret it as meaningful. If we cannot assign factual, referential meanings to a piece of discourse, we try to interpret it figuratively. This is what makes metaphor possible. If all attempts at finding referential meanings fail, we go on to establish the pragmatic meaning that piece of discourse has to us in its own specific situational context. In fact, we analyse all utterances as symptoms of a person's character, state, and behaviour. After such a multiple search for meanings it is hard to imagine that any linguistic, or indeed social, behaviour could be completely meaningless. How we respond to interpretative success and failure is an interesting question in itself. Such responses vary with individual and situation. For instance, if I fail to understand an article on physics, I blame my own incompetence. But if I fail to understand a text expressly written for readers like me, such as instructions for my new outboard motor or a paper on the linguistics of discourse, I blame the author. My guess is that those who hate modern poetry are unable or unwilling to adjust to the kind of mental exercise required of the reader, and transfer their feelings of inadequacy on the poet. Many rhetoricians have made a distinction between writer-friendly and reader-friendly discourse, and indeed, composition teaching might well concentrate on showing what is required for a text to be reader-friendly in its particular communicative situation. I shall return to this question in connection with information structure and shared knowledge.

Success Concepts

We have already seen that it is expedient to distinguish various levels of linguistic and communicative success. Partly, these levels presuppose different kinds of competence, partly the greater or less mastery of them will define and colour the linguistic and communicative prowess of a speaker or a writer.

The first prerequisite of understanding discourse is an ability to assign sounds to phonemes within the relevant phoneme inventory. Every student knows that foreign languages make use of strange sounds—noises we are tempted to call them—which may need long periods of intense ear-training to be properly appreciated. Again, an example. In native Finnish there is only one sibilant phoneme, /s/, which is, very roughly, a sound halfway between an English /s/ and /sh/. I say 'native Finnish' because educated Finns know for instance English can distinguish between /sh/ and /s/ and use their non-native /sh/ in recent loan-words and foreign placenames. Native Finnish lacks voiced sibilant phonemes. When a Finn is beginning to learn English one of his first problems is to learn to hear the difference between four sibilants where his native Finnish only has one. (Some Finns in fact voice their /s/ between vowels, but as such voicing fails to carry meaning it passes unnoticed.) The Finns' sibilant problem may stay with them for life, even through highly successful careers in an English-speaking environment.

So, one requirement for learning a foreign language is to make the necessary distinctions between sounds. A Finn cannot make a difference between *lies* and *lice* or *seat* and *sheet* without laborious ear-training and pronouncing drill. We might call this level of competence *phonological competence*. It is predominantly a matter of perception, at least at the initial stages of language learning where the learner cannot as yet predict sounds from his knowledge of the vocabulary (for instance by knowing that *mice* is an everyday English word whereas *mize* is not). And one can improve one's perception by practice, just as a few days in Kruger Park will teach a visitor to spot animals hidden to the untrained eye. The difficulties increase when the foreign language and the mother tongue make use of very different phonological devices. For a European, the tones of Chinese and Vietnamese are hard to hear and make, as are the clicks of several South-African languages. Note that the presence of word tone in, say, Chinese is part of the total economy pattern of the language: using tone makes it possible to use fewer phoneme sequences and a more restricted phonotaxis, and instead increase the repertoire of available words and structures by adding tone.

The next level is, broadly speaking, *syntactic competence*. I say 'broadly speaking' because I am here including lexis as part of syntax: what goes into a grammar and what into a dictionary is a choice made by the linguist rather than by the language. Phonological competence also merges into syntactic competence: many languages distinguish questions from statements through intonation, that is, by phonological means. Syntactic competence involves the ability of distinguishing between words and word forms (the latter a crucial point in synthetic languages with a rich morphology) and between different syntactic patterns.

The highest and most comprehensive level might be called *pragmatic competence*, the ability of using syntactic structures in a communicatively successful manner. By this time the reader may have asked whether there is a specific 'semantic competence', an ability to operate with meanings. The answer is that all levels of competence, phonological, syntactic and pragmatic, are involved with meanings: phonological competence with meaningful distinctions between sounds and sound patterns, syntactic competence with meaningful distinctions between words and syntactic structures, and pragmatic competence with meaningful distinctions arising from the use of different syntactic structures. I already mentioned the need for understanding irony as different from straight non-ironical language. Another area where we need pragmatic competence is the distinction of levels of politeness, most basically in simple overt matters such as modes of address, and more subtly in anticipating what will be understood as good manners and what as shocking, discourteous or indecent by a specific set of communication partners.

In applied linguistics, *communicative competence* has been used as a package term encapsulating all the different kinds of competence that go into successful communication. And many linguists have studied *communication strategies* in the sense of approaches adopted by speakers whose command of the relevant language is inadequate for normal expression. One might for instance try gesturing or even speaking in one's own language and thus transferring the onus of understanding on one's hearer (adopting the anecdotal imperialist attitude embodied in 'everybody understands English if I speak it loud enough'). And then, success in communication is not a matter of linguistic skills alone. Despite many years of language study at school, some speakers fail to manage even simple speaking situations. But there are others who do very well by supplementing a minimal vocabulary and next to no grammar with a splendid talent for uninhibited gesture and mime.

A practical question that arises in such situations is the level of tolerance of shortcomings in the speech and writing of others. I already touched upon this very subtle business in connection with behaviour, and the same principles apply to communication through language. People of the same social class and group will expect their communication partners to follow their own traditional patterns, including taboos: a young man may tell bawdy stories to his peers but not to a prim elderly aunt. Those who indulge their linguistic habits outside their own class run the risk of being branded as rude or uneducated or offensive, or as over-refined and supercilious. It is risky to break the pragmalinguistic conventions (which here include syntactic and, in speech, phonological ones) of one's hearers and readers. However, a foreign accent may often serve as an excuse. If somebody is obviously a foreigner, he or she cannot be expected to grasp all the niceties of civilised behaviour. To language teachers and learners this poses a problem. What should be the optimal level of ambition when we learn a foreign language? In each specific situation, where is the point of diminishing returns, beyond which the hearer's expectations, and requirements, rise without any benefit to communication? Often, we learn a language for some specific purpose which enables us to define the desired competence level. But often, as in the study of foreign languages at school, we aim at a loosely-defined general competence where reading comprehension and tourist phraseology is readily weighted at the expense of more creative spoken communication. Many of us have been embarrassed because our fluent command of a few foreign language phrases and a reasonable pronunciation have caused an overestimation of our skills. In tests at school and university, skills are measured on a pre-set scale reflecting the testers' opinions of the relevant language. In authentic communication, skills are related to the receptor's expectations in a specific situation.

Problems of this kind also raise the question of hierarchisation of levels such as comprehensibility, syntactic competence, and pragmatic competence. In interpreting actual discourse, do we first segment it into phonemes, next analyse its syntax and syntactic meaning, and finally place it in a situational context to get at pragmatic meaning supported by paralinguistic features? Even common sense tells us that this cannot be so. Actual discourse comprehension must be based on a highly complex interplay of the different levels and their concomitant types of competence. With a computer analogue we might speak of parallel processing of discourse within a number of linguistic subsystems. Understanding discourse is largely a matter of anticipation. We hear the initial phonemes of an utterance, and at once set about to predict what is most likely to follow, not only in phonology but also in syntax and in pragmatics. Indeed, simple experiments show that pragmatics and syntax can override phonology: try looking at your watch and asking, 'whappimeisit?', and your communication partner is more likely to say something like 'five past ten' than 'what is pime?'. Note that such pragmatic competence, anticipating who is likely to say what in a given situation, is based on what we might call cultural competence. We must have some experience of how people behave, verbally and otherwise, in relevant situations within the relevant culture.

Grammaticality, Acceptability, Appropriateness

Three areas of competence used for judgements of communicative success are Grammaticality, Acceptability, and Appropriateness. In another paper (Enkvist 1990) I have called them 'success concepts'.

First, grammaticality. Traditionally, linguists have made a distinction between what is 'correct' and what is 'wrong'—most of us will remember applications of such clear dichotomies from our schooldays. Indeed, in culturally homogeneous and stable societies there readily arises a consensus of 'right' and 'wrong' in language. An interesting problem is on what authority, on what grounds, such distinctions are made. Though the schoolteacher's judgements are usually conservative, nobody can claim that correctness equals age: Chaucer for instance used multiple negatives of a kind modern schoolteachers equate with sin. Appeals to authority are also common: the Fowlers recommend x and stigmatise y, and the Fowlers cannot be wrong. But where did the Fowlers and their colleagues get their authority, their feeling for correctness in English? Though people like to appeal to so-called authorities, the ultimate authority is social approval—what a social group believes about its language is 'correct' for that group and expressed by trustees such as the Fowlers. What is 'wrong' is what foreigners say before they have learned the language, or what other social groups say 'because nobody has succeeded in teaching them to speak correctly'. Note that in situations dominated by an educated elite, educated usage becomes the norm. But there are situations where social groups get mixed, in armies for instance, and then using an elite norm may be to invite trouble.

In another perspective, linguists too must worry about correctness. When a grammarian writes a grammar he or she must decide what to approve of and include, or what to exclude, leave out, and perhaps warn against if he or she is of the normative persuasion. Most grammarians in the history of linguistics have relied on their own socially conditioned intuitions, on the views of other grammarians, and on examples out of a body of approved, often literary, texts. Decisions concerning grammaticality became even more pressing with the advent of generative-transformational grammar. Now the grammarian's job was to set up formalised rules generating sentences. And the only way of deciding whether a rule was good or not was to see what it did. If it generated good sentences only, the rule was all right; if it generated garbage, it had to be revised. But the problem is how to distinguish good sentences from garbage before the rules are there. This is another chicken-or-egg-type of problem.

That such a simplistic view of right and wrong could persist in a science as sophisticated as linguistics owed largely to Saussure's (1955[1915]:36-39) distinction between *langue* and *parole* and Chomsky's (1965:3-15) related distinction between *competence* and *performance*. *Langue* and competence were norms, predominantly social and predominantly mental respectively, which steered people's linguistic behaviour. *Parole* and performance were what people actually did when communicating. By defining linguistics as being mainly concerned with *langue* and competence, and by usurping the privilege of deciding what was part of this *langue* and competence, a linguist could side-step the issue and himself decide on what was proper and what was not. Of course, this was a crude way of cutting a Gordian knot, ignoring the well-known fact that linguistic competence consists not of a single norm but of many, the choice depending for instance on the situation. To return to my old example: a sergeant uses one norm when addressing a squad of recruits, and (one would hope) a very different norm when talking to his family over breakfast. Nor would we long tolerate a person who speaks precisely the way he or she writes, and vice versa.

There have actually been two contrary opinions about such matters. One breed of linguists and grammarians pretend that the basic system of the language they are describing is homogeneous enough to be squeezed into one single description. In practice, they restrict themselves to one of the many existing norms, often the general norm of educated writing which is tacitly assumed to be closest to the *langue*. Another breed of linguists are prepared to accept the fact that what we call a language is in fact a vast collection of different sublanguages, situational styles, historical variants, regional dialects, socially conditioned variants, and individual idiosyncrasies. If so, a grammar of one single norm obviously cannot cover all these sublanguages. Therefore the best grammarians are prepared to recognise that they are dealing with a restricted spectrum of language, which they define and limit so as best to serve the future users of their grammars. Or, worse, to satisfy their own prejudices. One sometimes hears linguists trying to end arguments with a categoric 'In my language such a construction is impossible, so forget it'. But in reality, one of the most basic characteristics of a natural language is that it allows of variation in many dimensions, while still working well as a medium of communication. Ignoring variation means ignoring the essence of language.

So, grammaticality, in terms of yes or no, may be a practical aid to grammarians restricting themselves to one variant of language. In linguistics, however, it needs much refinement. One classic supplementary concept is *acceptability*. To study acceptability we must ask a suitably selected group of informants how they respond to certain test samples of language. Would they approve of them and use them, or not? And are they natural enough to be modified in certain specific ways, so that the informant can, for instance, turn an affirmative sentence into a negative or a question? Acceptability experiments were performed in the 1960s at University College London by Lord Quirk and his collaborators (see Quirk 1966; Greenbaum & Quirk 1970). What they showed was that sentences can be placed on a scale of acceptability: there are sentences that are approved or rejected by practically all informants, and there are sentences where the opinions remain divided.

How do people in fact judge the acceptability of a certain sentence? Presumably, one of the tacit criteria is contextualisation. The informants try to think of a situation in which it would seem natural for them to use that particular sentence. If they succeed, the sentence is all right; if not, it is no good; and between these two extremes there is a twilight zone of hesitation.

If so, the relationship between acceptability and situation is crucial enough to require further scrutiny. Some people have wide ranges of linguistic experience helping them to contextualise sentences. Others may lack such experience. Uncontrolled factors also readily seep into work with informants. One might imagine that informants (in actual fact students) who spent the morning reading modern poetry would respond differently from students who spent the morning studying Gibbon or Macaulay, or, why not, physical chemistry. Such arguments suggest the need for yet a third success concept, a 'contextualised acceptability' that we might call *appropriateness*. In setting up such a concept we are assuming that utterances have different acceptability ratings in different situations: what the sergeant says when drilling recruits is appropriate for the barracks but not for his family, and the other way round. What has been called 'style' is in fact a special type of appropriateness: a sonnet style involves language appropriate to sonnets, a scientific style to papers in science, a legal style to laws and statutes, and so forth. A full communicative competence will involve the mastery of appropriateness conventions across the entire spectrum of language a speaker is likely to need. In addition to the general norms of appropriateness within a family and everyday activities, professional people will thus have to master the language appropriate to their professions. And to appreciate modern poetry, a reader must cope with its appropriate linguistic conventions which may differ greatly from those in other kinds of communication.

Comprehension in toto

To understand our fellows, we need not necessarily rely on utterances that satisfy strict requirements of grammaticality, acceptability and appropriateness. On the contrary: as any close look at unwashed transcripts of actual speech shows, we use lots of ill-formed sentences and ill-formed chunks of discourse. But still we can be perfectly understood. Even sceptics are easy to convince by simple experiments: play a bit of videotape of a fluent speaker, and then give your audience a literal transcript of the text. It is an everyday experience that speech which sounds perfectly well-formed and even elegant on video may contain large numbers of false starts, repetitions, corrections, structure shifts, anacolutha, sentence fragments, stutterings and the like. They look bad on paper but often pass unnoticed in an actual communication situation. Some of them even contribute to the impression of fluency.

Such observations show that linguistic well-formedness, at whatever level we define it, is not a good criterion, or necessary condition, of comprehensibility in discourse. One reason is context. The situational context of discourse often contributes decisively to discourse understanding. Important contextual features are often signalled through deixis, that is, by anchoring the text to time, place and scenario through the use of specific pointers. What such deictics point at may be irrecoverable when the situational context is removed. For instance, if I show my spectacles to a live audience and say 'without these I could not manage' those hearing my text from tape have no chance of knowing what I am talking about. And if my son and I are mounting an outboard motor onto our boat, he will know precisely what 'two more inches' will mean. The context of situation is enough, and 'will you please move the motor two more inches to the left', in the midst of the action, would seem odd at least. Deixis is the most obvious, but by no means the only, way through which context enters into discourse. All kinds of allusion, intertextual references to other texts, agreement with or protest against conditions of politeness in the particular communication situation, and the like enter into the total meaning of discourse.

Texts also have both a referential meaning and a pragmatic meaning. If, in the midst of a lecture given in English, I insert three minutes of Finnish, only those who know Finnish will capture the referential meaning of my Finnish passage. That is, of the way in which my Finnish contributes to my argument and to its description of something outside the communication situation. But my Finnish still carries pragmatic meaning. My non-Finnish hearers might suspect that I have lost my mind, that I have forgotten where I am and who I am talking to, or, more charitably, that I wish to illustrate a point in my argument. In such instances my non-Finnish listeners must resort to pragmatic meaning. If we include paralanguage—gesture, voice colour, loudness, tempo and the like—among factors affecting pragmatic meaning, there may arise conflicts between reference and pragmatics. We may be nasty to our conversation partner but use the paralanguage of the sweetest friendship, or we may praise him with the paralanguage of hatred. Paralanguage then becomes an irony signal, a warning that the utterance should not be taken literally.

But if a linguistic description of a piece of discourse cannot tell us whether that piece of discourse is comprehensible or not, what criteria can we use instead? The best we can do is to look at the main purpose of communication, which is to transfer information from one person to others.

Meanings, we have noted, can only arise through choice between alternatives. If there is no choice, there is no information. We can postulate that the purpose of information transfer is to lead the receptor of a message to a scenario, or world of discourse, which resembles the world the speaker or writer wishes to convey. I say 'resembles', or is isomorphic with, rather than 'is identical with' because we cannot take for granted that the hearer or reader will build up a world quite the same as that of the speaker/writer. For communicative success it suffices that the worlds are similar enough in their relevant features to make people respond in the same way. In fact, as we have noted, people with different backgrounds may respond very differently to one and the same text: what is polite to some may be rude to others. And think of a non-physicist reading physics, or a non-linguist reading linguistics. They will understand something of the text but not reconstruct fully the world the writer was trying to describe. They do not have the background needed for complete scenario-building around the message.

Such considerations also explain our need for terminologies and of formal languages in science and scholarship. The purpose of terminology and of formulae, for instance in mathematics, physics, chemistry and logic is to maximise isomorphies between the scenarios of a speaker/writer and his information partners. Scientific and scholarly terms can be defined with a minimum of fuzz and potential discrepancy between the sender's and the receiver's scenarios as long as they are trained to use the same terminology. Fuzz is apt to enter into everyday, non-terminological language.

If we were to sum up these principles we might say: A text is comprehensible to those who can build around it a scenario in which it makes sense. A text can make sense in two ways. Those believing in a truth-functional semantics might say that a text makes sense if its scenario might be true ('might be' rather than 'is' to allow for fiction and fantasy). And those wishing to base their semantics on pragmatics might say that a text makes sense if the scenario it has created conforms to views of human behaviour which the receptor can accept and understand.

In these terms, when we hear or read a text, it restricts its scenario. When a text begins, we do not know what it wants to say and what its scenario will be. But as the text goes on, it successively narrows down its scenario and makes it increasingly specific. And it does so by eliminating alternatives. Hearing or reading a simple sentence such as *Susie flew to Paris*, we begin by noting that the scenario is concerned with Susie, thus eliminating all the world's billions of non-Susies, and perhaps also all the Susies known to the conversation partners who seem contextually irrelevant. *Flew* eliminates all other verbs or phrases, including those related to travel such as *walked* and *drove* and *cycled* and *went by train*. *To* eliminates prepositions such as *from* and *through* and *over*, and *Paris* eliminates all other possible expressions of place. At the same time, in English, the syntactic pattern expresses functions: we know for instance that *Susie* is the syntactic subject, and also the semantic agent, of the sentence. From all possible worlds we have been led to one specific world or scenario: the successive eliminations of alternatives have built up a scenario, a world, in which Susie flew to Paris.

I mentioned parallel processing, and to get a full picture of the communication process we should reconstruct the complex elimination process at all levels: phonology, syntax, pragmatics. We should also note that discourse comprehension often involves a reconstruction of the elements eliminated, that is, of the elements missing from the text itself but potentially present as features of the background. To understand the full implication of Susie's *flying* to Paris, we should know that, in other scenarios, she might have taken a car or a bicycle or even walked.

Information Structure and Text Strategy

Such considerations also bring up the question of information structure in discourse and in sentences. Indeed, under labels such as theme-and-rheme, topic-and-comment or presupposition-and-focus, information structure has become one of the central concerns of linguistics over the past quarter-century. One of the fundamental distinctions students of these matters like to make is between information a speaker/writer thinks the receptor already knows and has present in his mind, and information supposedly new to the receptor. A speaker or writer assumes that the receptor knows what everybody knows because of the human condition and life under the same stars; what people of the same cultural group know because of their group membership; what has been mentioned in the discourse; and what is present in the situational context. In English and in many other languages, unmarked sentence-initial elements usually refer to shared, old information, and new information is placed later in the sentence, the most important new information often going last. If new information is placed early in the sentence, it must be marked. A structure such as *It was by plane that Nils came to Durban*, is used in situations where a speaker assumes his hearer knows that Nils came to Durban; the new information is that he came by plane. And to warn the hearer that new information is placed early in the sentence there must be a special marker, namely cleft.

As the sentences of a coherent piece of discourse are not independent, but are linked to each other in a way steered by a uniform principle of discourse organisation, we need a concept such as *text strategy*. By 'text strategy' I mean the overall pattern of text organisation which governs the treatment of old and of new information. Before linguists learned to acknowledge that sentences are not autonomous but owe their form to their environment in text and situation, the difference between, say, *Charlie ate the apples.* and *The apples were eaten by Charlie.* was dismissed as a matter of 'emphasis', 'style' and the like. Though such sentence pairs are equivalent in the sense that if one of them is true, so is the other, they cannot freely replace each other in actual discourse. If you ask *What did Charlie do?. The apples were eaten by Charlie.* would be a weird answer; *the apples were eaten by Charlie*, on the contrary, or *the apples Charlie ate* (and the peaches Mary put into the refrigerator) would be acceptable as answers to *What happened to the apples and the peaches?.* As I am citing a passive in my example I should anticipate a usual question and add at once that another reason for the use of passives is agent repression: in their reports, scientists are used to saying *2 cc sulphuric acid were added to the solution* rather than *on Monday morning my assistant Susie Brown then added 2 cc sulphuric acid to the solution.* Because who did the adding, and when, is presumably irrelevant to the result. Agent repression occurs in many types of texts, laws and statutes for example.

The important lesson taught us by studies of information structure and text strategy is that the sentence is not autonomous. It is subject to forces in text and situation (discoursal parameters you might call them). Text strategies govern the linearisation of texts and sentences because they regulate the order in which alternatives are eliminated to lead the receptor to the desired scenario. Text strategies are exposed not only through element order and the syntax governing element order, but also through lexical choice. Compare *Nils is Elizabeth's father.* when talking about Nils, and *Elizabeth is Nils's daughter.* when talking about Elizabeth; or other converses such as *Betty is older than Susie/Susie is younger than Betty*, and countless other such converses. In fact, one might suggest that such lexical converses exist precisely to make possible different linearisations manifesting different text strategies. Instances such as *the post office is closed on Sundays/the post office is open on weekdays* both place the same old information, *the post office*, first. The former sentence answers the question *when is the post office closed?*, the latter *when is the post office open?*. But in such instances there is more text-strategic tolerance: both answers might be acceptable for either question, though one might be regarded as more appropriate to a specific context than the other.

For the composition teacher, all these matters should be of vital interest. Unfortunately they have been neglected for simpler matters of correctness. The point is that to produce receptor-friendly text, a speaker/writer does wisely in being aware of the distinction between old (shared) and new information. Generally, she does well in choosing her words and structures so as to place old information first. There are strategies permitting or even favouring placing new information early, or giving nothing but new information in sentence fragments. But they should be used with skill and, when need be, due marking. For a writer, the difficulty is to guess what of her message is already known to the reader, who lacks chances of protesting the way one could in face-to-face communication. This is why it is easier to write a letter to somebody we know well than, say, a letter to the editor of a paper with a very varied circle of readers. It is instructive to compare, for instance, editorials of different types of newspapers and try to figure out what knowledge they expect of their readers. And one may suspect that one reason why schoolchildren, and even adults, who are good at storytelling may produce very poor written versions of their stories. They may lack the ability and experience of correctly estimating amounts of shared knowledge. In face-to-face situations estimates are easier: people's faces show when we go wrong, and the floor may be open for direct questions. Perhaps explicit training in the use of text-strategic, lexical and syntactic devices handling distinctions between knowledge old and new should be included in the school syllabus.

Style

One of the concepts we should briefly discuss is *Style*. Ever since the ancients, much of rhetoric has been preoccupied with style, and countless literary scholars and critics have spoken and written about style, usually, however, basing their dicta on intuition rather than on precise, verifiable (or falsifiable to speak with Karl Popper) analyses. Indeed, literary scholars have often been hostile to attempts at linguistic analyses of style. In brief, style is an area where hermeneutics and positivism meet and readily clash. One of the problems in stylistics is to what extent such clashes could be turned into a merger.

To approach style, we might ask how our impressions of what we call the 'style' of a text actually arise. One useful answer is, through comparison with our memories of texts we regard as comparable (cf. Enkvist 1964; 1973a; 1973b). When we hear a new text, for instance a sermon, we compare the emerging text with our memories of other sermons. If we decide the language of the emerging text is similar to that of other sermons, the preacher is producing his discourse in the traditional style of sermons. If not, he is original; he might even be shocking if he uses the barracks language of a sergeant. What decides our choice of background texts for comparison is, obviously, our personal experience, our culture (are sermons a traditional type of discourse in our culture?), and, in the last instance, the situation (a clergyman delivering a sermon in a church). So, when we hear or read, we instinctively, or sometimes deliberately and consciously, compare the language with that of previous experiences of comparable texts. Our impression of the style arises out of this process of comparison. I already used the term 'frozen discourse' for ritual texts which must conform, even literally, to tradition. There are situations where conformity is at a premium and where good manners require a reproduction of patterns traditional within the culture and social group. And there are other situations were we praise originality and departures from tradition, as in certain types of poetry, in advertising, and so on. Those literary scholars who like to speak about intertextuality might define style as one of the results of intertextual comparison. Those interested in historical styles may define stylistic appropriateness in relation to patterns existing during a certain historical period.

In a wider semiotic sense, styles can exist in any area of human activity in which there is a choice affected by situation. The line between function and style is also often writ in water. If a basic function can be satisfied by various designs, their differences may count as styles. An airport for instance has a set of basic functions, such as moving and inspecting aeroplanes, passengers and luggage. But these requirements can be satisfied by various kinds of architectural solutions, which might then count as styles. Similarly, we might say that various types of text and discourse—weather reports, laundry lists, laws, statutes, operating instructions, and so forth—can be written with various degrees of stylistic latitude, as long as they satisfy their primary purpose. And the potential variation of styles will grow when the function of the discourse gets increasingly less stringently definable. A novel for instance usually contains different text types—dialogue, description, collage and so forth—which may all have distinctive styles, as may their manner of fitting into and contributing to the total design. Also the border between context and style may vary with the beholder's perspective. If one poet praises his monarch with a sonnet and another poet with an ode, the same function, praise of monarch, can be said to appear in two styles, sonnet and ode.

The linguist who wishes to base his work on falsifiable data must first define the body of text with which he wants to compare the text he is examining. (We might call this body of text a 'norm' as long as we strip 'norm' of all value judgements: here a 'norm' is not necessarily good or bad, beautiful or ugly.) He or she must then opt for a definite model for describing the texts. This model must be adequate to bring out the potential similarities and differences that define the style. And the linguist may then have to devise statistical methods for showing the significance of the similarities and differences between text and norm. Such an approach cannot do justice to the processual aspect of style, because it does not imitate the receptor's ongoing comparison between the emerging text and the possible interplay of competing norms or norms that supplement each other. So, it remains impossible, at least at our present stage of knowledge, to imitate hermeneutics by positivist analyses. But stylistic and stylo-statistical analysis are valid as methods of linguistic description, as well as valuable supplements to subjective impressions of styles.

There is yet another point that arises in connection with style. If we define style as a choice of expression affected by situational context, we should also regard the different codes that occur in code-switching situations as styles. If a person speaks standard High German in certain situations but switches into another kind of German—such as dialect or Swiss German—in other situations, these variants of German satisfy our definition of styles. The same applies to languages. If a person speaks English in one range of situations, Afrikaans in another, and Zulu in a third range of situations, these three different languages then take over the function of styles. In this sense all studies of code-switching turn out to be studies of style.

So, style is part of meaning, not merely meaningless surface decoration. For an appreciation of styles we must know the relevant set of norms. This in turn argues for the maintenance of canons, that is, exemplified sets of norms, as fundamental to all studies of language and literature. What should go into the canons is, however, a matter of culture and society, rather than of linguistic science as such.

Academy of Finland

References and Further Reading

Chomsky, Noam 1965. *Aspects of the Theory of Syntax.* Cambridge:MIT.

de Saussure, Ferdinand 1955[1915]. *Cours de linguistique générale.* Publié par Bally, Charles et al. Paris: Payot.

Enkvist, Nils Erik 1964. On Defining Style. In Enkvist, Nils Erik, John Spencer & Michael J Gregory (eds): *Linguistics and Style.* London: Oxford University Press.

Enkvist, Nils Erik 1973a. *Linguistic Stylistics.* The Hague & Paris: Mouton.

Enkvist, Nils Erik 1973b. *Stilforskning och stilteori.* Lund: Gleerup.

Enkvist, Nils Erik 1982. Categories of Situational Context from the Perspective of Stylistics. In Kinsella, Valerie (ed): *Surveys 1. Eight State-of-the-Art Articles on Key Areas in Language Teaching.* Cambridge: Cambridge University Press. (Cambridge Language Teaching Surveys.)

Enkvist, Nils Erik 1987. A Note Towards the Definition of Text Strategy. *Zeitschrift für Phonetik, Sprachwissenschaft und Kommunikationsforschung* 40,1:19-27.

Enkvist, Nils Erik 1988. On the Interpretability of Texts in General and Literary Texts in Particular. In Sell, Roger D (ed): *Literary Pragmatics*. London & New York: Routledge.

Enkvist, Nils Erik 1989. Connexity, Interpretability, Universes of Discourse, and Text Worlds. In Allén, Sture (ed): *Possible Worlds in Humanities, Arts and Sciences*. Berlin & New York: Walter de Gruyter. (Research in Text Theory 14.)

Enkvist, Nils Erik 1990. Success Concepts. In Lindeberg, Ann-Charlotte, Nils Erik Enkvist & Kay Wikberg (eds): *Nordic Research on Text and Discourse. Nordtext Symposium 1990*. Cbo: Cbo Academy Press.

Enkvist, Nils Erik 1994. The Epistemic Gap in Linguistic Stylistics. In Winter, Werner (ed): *On Languages and Language. Societas Linguistica Europaea Presidential Addresses*. Berlin & New York: Walter de Gruyter.

Quirk, Randolph & Jan Svartvik 1966. *Investigating Linguistic Acceptability*. London, The Hague & Paris: Mouton.

Greenbaum, Sidney & Randolph Quirk 1970. *Elicitation Experiments in English: Linguistic Studies in Use and Attitude*. London: Longmans.