1

Language in use

A sample of language

Given a sample of language, there are all kinds of things we can say about it. Take, for example, a familiar public notice:

KEEP OFF THE GRASS

This, to begin with, is something in English, as distinct from French or Arabic or Chinese or any other language. It consists of four words, all in capital letters, and all, we might more expertly add, monosyllabic. If we have been fortunate enough to have had some instruction in linguistics, we might then go on point out that the words combine to form a grammatical unit called a **sentence**, and a sentence furthermore of an imperative as distinct from a declarative or interrogative kind consisting of two main **constituents.** The first is a verb phrase consisting of the two words *keep off*, the second a noun phrase which itself consists of two constituents, a definite article *the* and a noun *grass*. Noting these grammatical features, we might think up a number of other sentences that seem to be structured in the same way: *Put out the light*, for example, or *Turn off the tap*, only to realize perhaps that appearances are deceptive and that these are actually not quite the same, but interestingly different. For these two structures can also take the form of the alternative sequences *Put the light out*, *Turn the tap off*, but *Keep the grass off* will not do. So examining the properties of this sample of ours as a sentence might lead us into a fascinating excursion into the mysteries of grammatical analysis.

But although linguists might delight in examining our sample in this way, this is not the kind of thing people would customarily do. Languages are traditionally recorded for us in analytic terms:

grammars display the range of possible structural combinations in sentences, and dictionaries provide us with the meanings of words separated out and listed in alphabetic order. These can be said to represent the encoded resources of form and meaning that speakers of a particular language know and draw upon intuitively when they use it. But they do not correspond with how speakers actually experience it as use. When they come across a public notice, they do not see it as a sample of language and analyse it into its formal constituents. They take note of it only to the extent that they recognize its purpose, as something not to analyse but to act upon. In other words they treat it as a **text**.

What is a text?

A text can be defined as an actual use of language, as distinct from a sentence which is an abstract unit of linguistic analysis. We identify a piece of language as a text as soon as we recognize that it has been produced for a communicative purpose. But we can identify a text as a purposeful use of language without necessarily being able to interpret just what is meant by it. It is a fairly common experience to come across texts in an unknown language which we nevertheless recognize as public notices, food labels, menus, or operating instructions, and to be frustrated by the inability to understand them. Clearly we would generally need to know the language a text is in to be able to interpret it. But this is not the only condition on interpretation. We may know what the language means but still not understand what is meant by its use in a particular text.

Consider again the public notice 'KEEP OFF THE GRASS'. We may know well enough what the word *grass* **denotes** (and should we be in any doubt we can consult a dictionary to find out). But what the word denotes is not the same as knowing what it is meant to **refer** to when it occurs here in the phrase *the grass*. The definite article *the* signals that what is being referred to is a matter of shared knowledge. The grass. But which grass? Obviously, one might say, the grass in the vicinity of the notice. So what we do is to establish reference by relating the text to the **context** in which it is located. But then the question arises as to how far this vicinity is meant to extend. Does *the grass* refer just to the particular patch where the notice is placed, or to other patches nearby as well, or to the whole park? The range of reference is not specified in the language itself. We make assumptions about what it is on the basis of what we know about public notices of this kind and how they are conventionally meant to be understood. In other words we relate the text not only to the actual situational context in which we find it, but to the abstract cultural context of what we know to be conventional.

And by relating text to context we infer not only what the notice refers to, but also what its purpose is. We recognize that it is intended as a prohibition, although whether we choose to pay any attention to it is another matter—and one we shall be taking up later.

The same point can be made about other notices we come across in daily life. Thus we recognize that the texts 'HANDLE WITH CARE' or 'THIS SIDE UP' refer to a container on which they are written and function as requests, that 'WET PAINT' refers to some surface in the immediate vicinity that has been newly painted, and functions as a warning. Similarly, when we see the label 'KEEP AWAY FROM CHILDREN' on a medicine bottle, we take this as a specific warning in reference to the particular contents of the bottle, rather than, say, as a piece of general advice to keep clear of young people at all times. When we come across notices and labels, then, we make sense of them by relating the language to the immediate perceptual context where they are located, and to the conceptual context of our knowledge of how such texts are designed to function. We cannot make sense out of them simply by focusing on the language itself. In the case of simple texts like notices and labels, establishing the language-context connections is usually a fairly straightforward matter. With other texts, even apparently simple ones, making such connections is not so easy, as anybody who has had the experience of assembling furniture from a set of instructions is likely to testify.

Text and discourse

The simple texts we have been considering so far all serve an obvious utilitarian purpose: notices, labels, instructions are designed to be directly acted upon and to get things done. But of course not all texts are so simple in form or so straightforward in function. Although, as we have seen, not all texts extend beyond the sentence, a great many of them do: travel guides, information leaflets, newspaper articles, interviews, speeches, reports, poems, and so on. Some of these have an obvious utility function but others are meant to serve a range of different social purposes: to give information, express a point of view, shape opinion, provide entertainment, and so on. These functions, furthermore, are frequently combined in complex ways: a travel guide, for example, may provide information, but is also designed to promote the attractions it describes; and what is presented as a factual account in a newspaper article will usually reflect, and promote, a particular point of view.

Whether simple or complex, all texts are uses of language which are produced with the intention to refer to something for some purpose. We identify a stretch of language as a text when we recognize this intention, and there are times when the intention is made explicit as when a text is labelled as a *notice*, or *instructions*, or *report* or *proclamation*. But recognizing a text is not the same as realizing its meaning. You may not know what is being referred to in a particular text, or in part of a text; or you may know full well what is being referred to, but fail to see what communicative purpose lies behind the reference. In the case of simple texts, like public notices, it will be a straightforward matter to match up intention with interpretation, but in the case of more complex ones, like newspaper articles, such matching can, as we shall see later, prove to be highly problematic.

People produce texts to get a message across, to express ideas and beliefs, to explain something, to get other people to do certain things or to think in a certain way, and so on. We can refer to this complex of communicative purposes as the **discourse** that underlies the text and motivates its production in the first place. But at the receiving end readers or listeners then have to make meaning out of the text to make it a communicative reality. In other words, they have to interpret the text as a discourse that makes sense to them. Texts, in this view, do not contain meaning, but are used to mediate it across discourses. Sometimes, of course, as with the notices we have been considering, the mediation is relatively straightforward: what the text means to the reader will generally match up with what the producer of the text meant by it. Obviously we must generally assume that texts will serve to mediate some convergence between discourses, or otherwise no communication would take place at all, but, as we shall see, the degree of convergence varies a good deal. As we all know from our own experience, no matter how explicitly we think we have **textualized** what we want to say, there is always the possibility that it will be interpreted otherwise.

So the term **discourse** is taken here to refer both to what a text producer meant by a text and what a text means to the receiver. Of course what somebody means by producing a particular text may well relate to broader issues of what social and ideological values they subscribe to, and another way of thinking of discourse is indeed to focus on such broader issues and look at how texts can be used to express, and impose, certain ways of thinking about the world. This is something we shall return to later in the book (in Chapter 7).

Spoken and written text

For the moment, the point to be made is that texts are the perceptible traces of the process, not itself open to direct perception, of mediating a message. In conversation, these traces are typically fragmented and ephemeral, and disappear as soon as they are produced to serve their immediate discourse purpose. They can, of course, be recorded, but do not need to be, and usually are not. Thus, participants in spoken interaction produce and process text as they go along and there is no need for it to be retained as a record for it to mediate their discourse, and this mediation is regulated on-line to negotiate whatever convergence between intention and interpretation is required for the purpose. Written text, on the other hand, is not jointly constructed and construed on-line in this way. It is typically designed and recorded unilaterally in the act of production by one of the participants, the writer, as a completed expression of the intended message. The text is then taken up and interpreted as a separate process. The mediation, therefore, is displaced and delayed and this obviously will often make a convergence between intention and interpretation more difficult to achieve.

And there is a further difficulty. When people communicate they do not only produce linguistic texts. In speech, they make use not only of language but of **paralanguage**—tones of voice, varying stress, pauses, and so on, and what they say is accompanied by facial expression, or gesture, as part of the message they intend to get across. In written communication, too, how a text is given a particular shape by choice of typeface, or its arrangement on a page, may suggest significance over and above what it signifies linguistically. And it may be **multimodal** in that the text is accompanied by, and related to, other modes of communication —pictures, diagrams, charts, and so on.

It is the lack of direct correspondence between text and discourse that makes communication so indeterminate, and so intriguing. Life would in some ways be much easier if we could pin things down more precisely, if all we needed to do to communicate was to assemble a combination of linguistic forms of fixed meaning and transmit them for dismantling at the receiving end. A text would then signal its own meaning, whatever the context or purpose of its production. But when we use language we do not just present the meanings that are encoded in it, we exploit them as a potential resource for making meaning of our own. The encoded meanings are **semantic meanings** and are what are described in dictionaries and grammar books. To know a language is to know what they are. But in using a language we not only put this knowledge on display but also act upon it as appropriate to our communicative intentions: in other words we always make this semantic meaning serve a **pragmatic** purpose.

Illustration from a crowded train

To illustrate: let us suppose that you overhear the following **utterance** in a conversation between two people in a crowded train.

> He has put it in a safe place and it will not be found.

As linguistic data, we might note that as far as grammar is concerned this is a complete and well formed sentence of English. The present perfect in the first half (*has put*) and the passive in the second (*will ... be found*) are produced in conformity with grammatical rule, there is agreement, as required, between the singular pronoun *he* and the following verb, the past tense forms are morphologically well formed (*put* not *putted*, *found* not *finded*), word order is as it should be, and so on. As far as **lexis** is concerned, we can attest that the words that occur in the sentence are quite normal English ones (*put, safe, place, find*). But we do not only recognize that this is a regular and well formed example of English. Since we know what semantic meaning is signified by the grammatical and lexical forms we have recognized, we are able to decode what has been encoded and assign it a meaning, as a sentence. However, we are still in the dark about what this person is actually talking about. Who is 'he' and what is 'it'? These pronouns have an established denotation: *he* encodes the **semantic features** of singular and masculine, and *it* the semantic features of singular and inanimate. But although a knowledge of these denotations narrows down the possibilities, it does not tell you who or what is being referred to. What is meant by the language will continue to be pragmatically elusive.

This utterance is, of course, only a fragment of conversation, one piece of the text, that the two people in the train are producing in the process of enacting their discourse. They know what they are talking about because they have established the context of shared knowledge and assumption that the actual language they produce keys into. If we are not a party to this context and only have the linguistic trace of their discourse to go by, we cannot interpret what they mean by what they say. When they leave the train, they take their meaning with them, and we will not be able to recover it, no matter how closely we analyse the actual language.

Conclusion

To summarize. When people communicate with each other, they draw on the semantic resources encoded in their language to key into a context they assume to be shared so as to enact a discourse, that is, to get their intended message across to some second person party. The linguistic trace of this process is the text. In the case of conversation, the text is jointly produced as the discourse proceeds by overt interaction, and it typically disappears once it has served its purpose. In the case of writing the text is unilaterally produced and remains as a permanent record. But it is still only a discourse trace, and what is meant by it has to be inferred by interpretation, and this inevitably raises the question of how far this interpretation corresponds with the intentions that informed the discourse which gave rise to the text in the first place.

In the normal circumstances of use, of course, we only pay attention to text in order to realize its discourse function and so we tend to think of the two as the same thing, as indeed do some linguists, and talk about the meaning of a text as shorthand for what it means to us or what it might mean to the writer or speaker. But although we normally experience text as part of the discourse process, it is perfectly possible to focus on the text alone. This is after all what proofreaders generally do when they scrutinize a piece of writing to identify typographical errors, or wordings that do not conform to established code conventions. More interestingly, texts can also be subjected to close analytic study to find out patterns of actual usage which those producing them may be quite unaware of, an area of linguistic description we shall be returning to in a later chapter. But before we get on to this, it will be necessary to say a little more about how the concepts introduced in this chapter—semantic and pragmatic meaning, sentence, utterance, text, context, and discourse—figure in a general model of communication.

2

Communication

Grammar and communication

As was pointed out at the beginning of the first chapter, linguists have traditionally focused their attention primarily on the internal properties of languages, on how meaning is **formally encoded** in lexis and grammar. The description of such properties can be said to be an account of what people know of their language, an account of their **linguistic competence**. It was this competence that enabled us to describe the grammatical features of the utterance in the train in the preceding chapter. Of course that description made use of terminology which might well be unfamiliar with people competent in English: knowing the grammar of a language is not the same as knowing how to describe it—that is the business of the grammarian. But the point is that anybody competent in English would recognize that the utterance conforms to the **encoding conventions** of the standard language, that it exemplifies a well formed sentence in English, which it would not do if it had taken the form, for example:

> They has it in a safe place put and it will not to find.

The original utterance, we can say, exemplifies a possible sentence in English and this second one does not.

So on hearing this remark in the train, one judgement we can make on the basis of our linguistic competence is whether it is grammatically and lexically possible or not, that is to say, in accordance with the encoding conventions of the language. But we can also recognize degrees of possibility. So if the utterance had been: