Chapter 3

Strategies for Motivating a Set of Relations

> Lacking a clear way to continue our description of the underlying conceptual system, we turned to analyses of the parts of it frozen by social convention into the English lexicon.
>
> *Miller & Johnson-Laird,* Language and Perception*, p697*

3.1 Introduction

In the previous chapter, a great many sets of coherence relations of different kinds were described, and the need was emphasised for a means of justifying one such set over the others. The justification should ideally perform several functions. It should provide an independent definition of relations, which makes clear *why it is* that relations underlie discourse coherence. It should also impose some kind of limit on the *size* of the set of relations, so that a theory explaining coherence in terms of relations is not unfalsifiable. It should also make clear how a given set of relations is *more appropriate than any other set* for the theoretical task it is to perform.

This chapter is given over to examining various strategies that have been proposed in the literature for justifying a set of relations. The first of these is to pick a set of relations that allows an adequate description of all the discourses that the theory purports to explain. The second strategy suggests that connective **cue phrases** such as *because* and *however* can be used to motivate a set of relations. A final strategy is based on the contention that relations model psychological constructs used in human text processing. Each of these strategies has advantages and shortcomings, which will be discussed. In the second half of the chapter, a new strategy for justifying a set of relations is proposed, which (it will be argued) combines the advantages of the above methods, and minimises their shortcomings. According to this strategy, cue phrases can be used as evidence for relations *precisely if* relations are considered as psychologically real entities. This is the central claim of the thesis: the rest of the chapter will be devoted to clarifying it, arguing for it, and defending it against various objections.

3.2 Choosing a 'Descriptively Adequate' Set of Relations

The most fundamental rationale for a set of coherence relations is that they enable an analysis of all the texts which the theory using them purports to account for. Some sets of relations are tailored specifically to particular registers of discourse (for instance, those in Rösner and Stede's (1992) planner, and in that of Vander Linden *et al* (1992) are just for processing instructional texts) while other sets (such as that of RST[1]) are claimed to be almost register-independent. In all cases, a claim is made that a given set of relations is sufficient to enable an analysis of a previously-specified set of texts. Clearly, any theory invoking relations must address the question of how well they can be made to fit texts, and so must make this claim in some way.

However, it remains to be determined what it means to say that a set of relations is 'sufficient' to analyse a text. As noted in Section 2.2.3, such a claim must be interpreted in the light of a theory of span structure, which specifies the places in a coherent text where relations are expected to be found. RST makes a rigorous prediction, holding that *each clause* in a text is linked to some other portion of it by a rhetorical relation; the same goes for each 'composite' span created by a schema application (except the top span, of course). Other theories impose lesser constraints: for instance, Grosz and Sidner's (1986) relations link 'discourse segments', typically composed of several clauses—which means that there are many clauses between which relations are not construed to apply. At the other end of the scale, accounts like Vander Linden *et al*'s (1992) suggest that relations can even exist *within single clauses*—although they do not yet specify exactly when such relations will be found.

How should we decide how densely a text must be filled with relations in order for it to be 'adequately described'? It seems that this question can only be addressed when we know more about what it signifies for two text segments to be connected by a relation. It makes no sense to set a criterion for 'descriptive adequacy' which cannot *itself* be further justified; except perhaps as a working hypothesis.[2] On the other hand, if we had a theory about text processing in which relations played a role, this might well be able to tell us where to *expect* to find relations in a text. Say the theory states that a reader links each paragraph to the previous one by means of relations. The notion of descriptive adequacy would be different for this theory than for one which claimed that the processing of individual clauses is mediated by relations.

There is a second problem with descriptive adequacy as the sole criterion for judging a set of relations. Many different sets of relations can be used to describe any given text: the subtlety of the distinctions between relations is not constrained. For instance, while Mann and Thompson choose to split 'causal' relations into five sepa-

[1] RST analyses are claimed for 'virtually every text' (Mann and Thompson (1988) p20), though exceptions are mentioned for some registers, such as legal documents and some kinds of poetry.

[2] This suggestion is at odds with Mann and Thompson's (1988) position. They claim that the decision about the size of text units to be related is 'arbitrary'—their only caveat being that the units should be defined in some theory-neutral way (p6).

rate groups (VOLITIONAL-CAUSE, NON-VOLITIONAL-CAUSE, VOLITIONAL-RESULT, NON-VOLITIONAL-RESULT, and PURPOSE), others have chosen a less fine-grained distinction: Scott and de Souza (1990) group all these relations together for the purposes of textual realisation. Vander Linden *et al* combine the VOLITIONAL-RESULT and NON-VOLITIONAL-RESULT categories to create a simple RESULT relation, but keep PURPOSE separate. Each of these models seems feasible as a 'descriptive' framework for text: it is equally feasible to split up Mann and Thompson's relations still further, creating relations like IMMEDIATE-VOLITIONAL-CAUSE, DISTANT-VOLITIONAL-CAUSE and so on. If we are looking to justify a set of relations, we need to have a way of deciding on an appropriate level of detail.

It could be argued that although a *tight* specification of the requisite level of detail is not feasible, some sort of approximation can still be arrived at. For example, it is possible to describe any text just using the relations 'CAUSAL' and 'NON-CAUSAL'; we might argue that such a description is a clear case of 'descriptive inadequacy'. But this is to miss the point: Grosz and Sidner use just two relations, but their relations are adequate (it is claimed) to achieve the task for which they were designed, namely accounting for the pattern of pronominalisation in texts. The essential point is, again, that it is only when relations are given some sort of *theoretical role* that we can even begin to talk about descriptive adequacy. The standards of adequacy are set by the demands of the theory in which the relations figure. The theory will determine what information about a text relations are supposed to capture; we can *then* ask whether the description they provide is in fact sufficient to capture that information.

In short, the criterion of descriptive adequacy, while in some sense essential for any set of relations, is far from being able to stand on its own as a method of justification.

3.3 Associating Relations with Cue Phrases

Cue phrases[3]—clausal/sentence connectives such as *but* and *because*—have provided another source of evidence for justifying a set of relations. Such phrases are often conceived as *signalling* relations in a text: thus, for instance, *because* can be used to signal the presence of a CAUSE relation:

(3.1) Jane fed Lars *because* he was getting so hungry.

It is important to recall that we are interested in looking at the *set* of cue phrases, in order to motivate a *set* of relations. As emphasised in Section 1.2, this task is distinct from the 'first-order' task of identifying relations in actual texts. A study of cue phrases will certainly not suffice for this latter task; it is widely accepted that relations can be

[3] What I am calling 'cue phrases' have been given many different names in the past: 'conjunctive elements' (Halliday and Hasan (1976)), 'clue words' (Cohen (1984)), 'cue phrases' ((1986)), 'discourse markers' (Schiffrin (1987)), 'meta-technical utterances' (Zuckerman and Pearl (1986)). Various different ways of defining them have been suggested, and the different definitions pick out slightly different sets of phrases. For the moment, the class of cue phrases can be identified by typical examples, such as *however*, *then*, *previously*, *or*, *next*, *while*. I will propose a more rigorous definition of my own in the next chapter.

unmarked in text, so cue phrases will only be a partial source of information in this regard.

3.3.1 An Atttractive Source of Evidence for Relations

Cue phrases would certainly be convenient as a source of evidence for a set of relations. There are a large number of them, and the differences between them can be quite subtle; associating cue phrases with relations would yield a sophisticated classification of relations. For instance, Elhadad and MeKeown (1990) have noted some subtle differences between *but* and *although*, phrases which are often inter-substitutable:

(3.2) a. He failed the exam *but* he's smart. Let's hire him.
b. *He failed the exam *although* he's smart. Let's hire him.

We would expect these differences to be reflected in the set of relations if we constructed it to mirror the set of cue phrases.

At the same time, using cue phrases as a source of evidence would give us a way to decide at what level of detail to *stop* making distinctions. If alternative cue phrases exist to pick out two similar relations (as in the above cases), then they can be distinguished. But there would be no need to create two separate relations if no cue phrases exist for distinguishing them in text. Thus, for example, there would be no need to distinguish between 'FEMALE-VOLITIONAL-CAUSE' and 'MALE-VOLITIONAL-CAUSE'.

Finally, cue phrases simply provide an *extra source of information* when it comes to working out relation definitions. Many systems for justifying a set of relations start 'from first principles', without any preconceptions about which relations are eventually going to be decided on. Thus Maier and Hovy's (1991) taxonomy of relations is based on Halliday and Hasan's (1976) analysis of the functions of language as 'ideational, interpersonal and textual'; Hobbs' (1985) classification is based on similar abstract considerations about 'the situation in which discourse between a speaker and a listener takes place'. It is certainly essential to have such high-level concerns in mind when working out relation definitions—but the task would be considerably eased if we could also make use of information about the applicability of cue phrases. This is because we can find out about all the situations in which a relation can be used, *in advance of working out its definition*—we just need to examine the range of ways that the appropriate cue phrase can be used in discourse. For instance, having posited that a relation exists which is signalled by the cue phrase *although*, we can consider the different contexts in which *although* can be used, and try to abstract from these to create a definition of the relation itself. The ability to draw on concrete linguistic examples is likely to be of considerable help.

3.3.2 Previous Work with Cue Phrases

As was seen in Chapter 2, a number of researchers have made use of cue phrases in determining a set of relations. While Halliday and Hasan (1976) are only interested in classifying the linguistic resources available for signalling relations, other theorists have used the cohesive resources in a language as 'evidence' for a set of underlying coherence relations, which can themselves be marked or unmarked. The work of Ballard *et al* (1971) and Longacre (1983) follows this strategy, appealing to surface syntactic phenomena in the motivation of 'deep' interclausal relations:

> It is our contention...that a surface taxonomy of form within a language determines a similar taxonomy of deep relations, and that the two taxonomies stand and fall together.
>
> *Ballard* et al *(1971), p75*

According to this approach, the existence of a cue phrase in a language is testimony to the existence of a particular type of interclausal relation, and this relation can apply between clauses even in the absence of the cue phrase. Among the attractions of this method, Ballard *et al* note one I already mentioned—that it enables us to fix the level of detail of the analysis:

> The deep grammar...stops short of dissolution into general semantic or logical categories. It stops in fact where the structure of a given language indicates a cut-off point in that it sets up no more deep structure categories than are required to account for surface encodings.
>
> *Ibid.*

Martin (1992) follows a similar strategy, suggesting that the relations in a text are in principle markable by surface conjunctions. His classification of cue phrases is extremely detailed, and is a clear testimony to the benefits of cue phrases for creating a subtle taxonomy.

Several other researchers make reference to cue phrases when putting forward a set of relations. For instance, Hobbs (1985) uses cue phrases as an informal method of deciding on which of his relations applies in a text—if you can insert *then* between two segments, then the OCCASION relation 'is an excellent candidate'. But he is emphatic that such tests do not figure in the *definitions* of the relations he presents. Scott and de Souza (1990), in a study of how RST relations can be textually marked, found that a large number of them map closely onto cue phrases. Sanders *et al* (1992, 1993) also identify 'prototypical markers' for each relation in their taxonomy, although the main justification for the taxonomy (to be examined in Section 3.4) is not concerned with linguistic issues. Even Mann and Thompson's (1988) relations, which are expressly defined without reference to surface linguistic phenomena, can often be associated with classes of cue phrase.

3.3.3 Problems with Reliance on Cue Phrases

To sum up: there are many advantages to be gained in using cue phrases to decide on a set of coherence relations; and cue phrases have been quite widely used for this purpose. However, an important problem remains for all of these attempts to date: the decision to link relations to cue phrases *itself* needs to be justified. Without such a justification, we still lack a reason to prefer the chosen set of relations over other possible sets.

To illustrate the problem, consider the text in 3.3.3:

(3.3) Bob cooked supper that night. His wife had been working hard at the office all day.

The current suggestion is to account for the coherence of this text by saying that the two clauses are linked by 'some relation which can be signalled by a cue phrase'. (In this case, *because* seems the most likely cue phrase for the job.) But why should this *because*-based relation be any more appropriate in explaining coherence than relations like VOLITIONAL-CAUSE or CAUSE-OF-COOKING-EVENT, which don't happen to correspond to cue phrases? How can we use the association with a cue phrase to *argue* against someone who proposes an alternative relation?

Somehow an argument must be given that relations linked to cue phrases tell us something about the text that relations otherwise classified do not. But neither Longacre nor Martin (who are most explicit in their appeal to cue phrases) provide such an argument: in fact, justification seems to remain principally in terms of descriptive adequacy; and it has already been argued (in Section 3.2) that such an appeal is of little use by itself. To reiterate the point made in Section 3.2: the notion of 'descriptive adequacy' cannot really be used on its own to evaluate a descriptive formalism. It needs to go hand in hand with a *theory* about the thing described: then we can ask whether the formalism is adequate to provide a description in the appropriate theoretical terms. So for a sounder justification of the reliance on cue phrases, what is needed is some sort of theory about text, which would show why linking relations to cue phrases makes them *particularly revealing* as descriptive constructs. At that point, it would be reasonable to prefer them over any set of relations which could not be given such a theoretical role. But such a theory still needs to be provided.

3.4 Looking for 'Psychologically Real' Relations

A third method for justifying a set of relations begins immediately by giving them a theoretical role. The central idea is that relations model *psychological constructs*—that is, they tell us something about the psychological processes which occur in people when they create and interpret text. A given relation can then be justified by producing evidence that it is one that people actually use when processing text. This is why it makes sense to use the relation to link two text spans: because it models part of the process which *actually led* to these spans being juxtaposed as they are. By using a relation conceived of as psychologically real, we are not just describing the text in an arbitrary manner, but contributing to an *explanation* of why the text is the way it is. Clearly, this conception of relations gives us exactly the sort of 'theory about text' that we need in order to justify the use of relations as an appropriate descriptive device.

Many theories of relations make an appeal to psychological notions. As mentioned in the previous chapter, Hobbs (1985) thinks of his relations as 'text building strategies', used by the speaker to facilitate the job of the listener. And he makes some suggestions (p22) about the nature of these strategies. To re-iterate the quote given at the outset:

> It is tempting to speculate that... coherence relations are instantiations in discourse comprehension of more general principles of coherence that we apply in attempting to make sense out of the world we find ourselves in, principles that rest ultimately on some notion of cognitive economy.

Mann and Thompson's (1988) relations also embody psychological insights. Relations are functional constructs, associated with the particular effects a writer intends to achieve; relation definitions make extensive reference to the psychological states of the reader and writer. Yet no *evidence* is given that the RST relations are the ones people use: justification of the relations is again purely in terms of their 'descriptive adequacy' in the hands of discourse analysts.

By far the most thorough investigation of the idea that relations are psychologically real comes from Sanders, Spooren and Noordman (1992, 1993). Any theory proposing that relations 'model psychological constructs' must address two key issues. Firstly, some kind of account of human text processing is needed, which makes it clear what role is played in the mechanism by 'relation-like' constructs. Secondly, there has to be some way of *investigating* these psychological constructs, and working out what they are. Sanders *et al* address both of these requirements.

3.4.1 An Overview of Sanders *et al*'s Work

Parameterising the Space of Relations using 'Cognitively Basic' Primitives

Sanders *et al* approach the issue of psychological reality from the perspective of text understanding. Understanding a discourse involves constructing a coherent mental representation of it, and this in turn involves setting up appropriate *links* between the representations of its various segments.

These coherence links, it is argued, are likely to be established using *general cognitive resources*. During comprehension, the segments of the discourse are integrated into a language-independent representation; part of the reader's general framework for making sense of the world. Coherent texts are likely to be structured in such a way as to facilitate this integration—so Sanders *et al* propose that coherence relations should be investigated by looking for the 'cognitively basic' features which must underlie them.

Four 'cognitively basic' primitives are identified, according to which relations can be classified: these are described below.

- Basic operation. Every relation is deemed to have either a causal or an additive component. causal relations are those where a 'relevant' causal connection exists between the spans; all other relations are additive. a is an example of a causal relation; b is an example of an additive one.

    a. The drive to the arrivals was closed so that nobody could leave the terminal.

b. Centraal Beheer's turnover is about 2.4 billion guilders. In 1988 the profits increased from 75 million to 103 million guilders.

- Source of coherence. Every relation is coherent on semantic or pragmatic grounds. It is semantic if the spans are related in terms of their propositional content and pragmatic if they are related because of their illocutionary force. a is an example of a semantic relation; b is an example of a pragmatic one.

  a. Theo was exhausted because he had to run to the university.
  b. Theo was exhausted because he was gasping for breath.

- Order of segments. This distinction only applies to causal relations; they are deemed to have basic order if the antecedent is on the left, and non-basic order if it is on the right. a is an example of a **basic** relation; b is an example of a non-basic one.

  a. The drive to the arrivals was closed, so nobody could leave the terminal.
  b. Nobody could leave the terminal, because the drive to the arrivals was closed.

- Polarity. A relation is positive if its basic operation links the content of the two spans as they stand, and negative if it links the content of one of the spans to the negation of the content of the other span. Negative polarity relations typically involve either a violation of expectation, where the expectation derives from a causal basic relation; or a contrast, where the basic relation is additive. a is an example of a positive relation; b is an example of a negative one.

  a. Because he had political experience, he was elected president.
  b. Although he had no political experience, he was elected president.

These four parameters can combine to form twelve 'complex' relation types. For each type, Sanders *et al* provide one or more sample RST-like relations: for instance, CAUSE-CONSEQUENCE (basic operation = causal; source of coherence = semantic; order = basic; polarity = positive); or CLAIM-ARGUMENT (basic operation = causal; source of coherence = pragmatic; order = non-basic; polarity = positive). Each of the relations is associated with a 'typical' connective word used for marking it.

Evidence for the Parameterisation

Support for the four parameters is provided by a number of empirical experiments. The first of these used discourse analysts as subjects: they were given definitions of all the relations, and asked to decide which relations were appropriate for a number of sample texts. The second experiment used 'naive' subjects who did not know about the relation definitions: they were shown sample texts without explicit connectives, and had to decide which connective word was most suitable. Both experiments were designed to test how much agreement there is on how to use the relations. In both cases it was found that there was a fair amount of agreement between subjects. Equally importantly, where there was disagreement over which relation to use, it tended to be over the value of a single parameter only—this provides support for the independence of the decisions about the different parameters.

However, in both experiments, there were differences in the *strength* of evidence for different parameters. In the first experiment, there was hardly any confusion over the value of the polarity parameter. But for negative polarity relations, there was a great deal of confusion about the source of coherence parameter: in fact, this parameter is the least agreed upon for all classes of relations. This finding is replicated in second experiment: agreement is lowest for source of coherence.

A third experiment also used discourse analysts as subjects: here the task was to *compare* the coherence relations in a number of different sample texts, and to group those texts which used the same relation. The results showed four distinct clusters of relations: positive causal relations, positive additive, negative relations, and 'conditional' relations (a subtype of causal relation). Again, source of coherence is not well distinguished—there is no evidence for this parameter amongst negative relations, and between positive relations there is not much. Finally, in this experiment, no evidence at all is found for the order of spans parameter.

The last experiment was targeted specifically at the source of coherence parameter. It was hypothesised that confusions regarding this parameter would be lessened if relations were presented in contexts rich enough to disambiguate them. The subjects were again discourse analysts; the task was similar to that in the first experiment. Only those relations with positive polarity were examined. For these relations, under these conditions, it was found that a distinction can indeed be made between semantic and pragmatic relations.

3.4.2 Some Problems with Sanders *et al*'s Parameterisation

While the initial idea that 'relations are psychologically real' provides a very promising method for justifying a set of relations, deciding on exactly *which* set of relations is psychologically real presents problems of its own. The experiments reported above can be criticised on a number of grounds: they do not provide conclusive evidence for Sanders *et al*'s four-way parameterisation.

Problems with Specific Parameters

Firstly, two of the parameters are supported much less strongly than the others. The first three experiments give only weak support for source of coherence. The final experiment only gives support for this parameter for a subset of the relations in the set; and then only under ideal conditions.

Questions can also be raised about order of spans: there was no evidence at all to support this parameter in the third experiment. The first experiment showed that analysts *could* distinguish between basic and non-basic order; but this just shows that they could use the relation definitions they were given, it does not legitimise these definitions.

How do you Choose a Set of Parameters to Test?

The problems with source of coherence and order of spans throw up a more general problem with the experimental methodology—that of *deciding* on a set of parameters to test. Sanders *et al* (1993) take a lot of trouble giving *a priori* justifications for each parameter; but there is still no *systematic* approach towards deciding on the set. For instance, we have no way of knowing that *all* the relevant parameters have been identified. Maybe there are more than four: 'hypotheticality' could be another candidate, in view of the results of the third experiment.

One argument given for the four chosen parameters is that they result in a *productive* taxonomy—in other words, that every combination of parameter values seems to yield a plausible relation. This point is cited as one reason why temporal relations are not included as a parameter. But for a start, Sanders *et al*'s four chosen parameters do not currently result in a *completely* productive taxonomy: the fact that order of spans is only defined for causal relations means that it has four 'empty' slots. Moreover, why should we expect that the cognitive factors which underlie coherence relations will result in a neat parameterisation? It seems perfectly plausible that some factors are only relevant for a *particular* type of coherence relation.

Some Questions about the Experimental Paradigms

Two kinds of experiments are carried out by Sanders *et al.* In one kind, the subjects are discourse analysts and the task makes explicit reference to coherence relations: texts have to be analysed using relations, or sorted into groups on the basis of the relations they use. In the other kind of experiment, the subjects are 'naive' about theories of discourse, and the cue phrases they use are taken as evidence of the relations that they perceive in a text. Both types of experiment are open to question. It is not certain that results obtained from discourse analysts, thinking explicitly about relations, can be taken as evidence for the kind of relations that people normally use when they process text. Neither can the evidence from cue phrases be taken as conclusive: why should we suppose such a tight association between cue phrases and the constructs we make use of when processing text? At the very least, an argument must be given for this policy. Otherwise, we are making unfounded assumptions about the very constructs we are investigating.

The Grain-Size of Relations in the Taxonomy

For Sanders *et al*, the twelve combinations of parameter values do not pick out individual relations; rather *classes* of relations. Thus, for example, two 'prototypical relations' are identified which are causal, semantic, basic and positive—CAUSE-CONSEQUENCE and CONDITION-CONSEQUENCE. How are these individual relations identified? How many of them are required? Sanders *et al* talk about these extra relations as needed to achieve 'descriptive adequacy'; but once again, it is unclear how this criterion is to be interpreted. *What* needs to be described? At the outset, relations were conceived of as modelling the cognitive constructs involved in human text processing. Does this mean that such constructs exist at a finer level of detail than can be specified by the parameters? If so, additional principles are surely needed for picking them out. And if not, then what is the point in refining the taxonomy beyond the level of detail provided by the parameters?

To sum up: a number of problems remain with Sanders *et al*'s justification of a set of relations—although their enterprise of looking for evidence for cognitive text-processing constructs certainly seems the most attractive way of effecting a justification. In the next section, an alternative source of evidence for these cognitive constructs is proposed, which (hopefully) overcomes many of the problems that have so far been raised.

3.5 A New Motivation for Relations: Linguistic Evidence for Psychological Constructs

So far, we have looked at two methods for justifying a set of relations. One suggests using the space of cue phrases in a language to work out a taxonomy. This enables an extensive and detailed taxonomy to be worked out; but it is unclear what explanatory role relations thus justified are to play in a theory of discourse coherence. The other method is based on the idea that relations model the psychological constructs which mediate the production and interpretation of discourse. Here, the theoretical role of relations is clear; but difficulties arise in the attempt to *discover* what these constructs might actually be.

It will be noticed that the advantages and the problems for these two approaches are complementary. The central idea in what follows is that the approaches can be *combined*, so as to capitalise on their advantages and minimise their drawbacks, by taking cue phrases as evidence for cognitive text-processing constructs.

3.5.1 The Central Argument

The claim to be established is the following: that the existence of a cue phrase in a language is good grounds for inferring the existence of a corresponding 'relational' construct in the cognitive apparatus of those who use the language. It will be argued that a language is likely to contain resources for making explicit all the relations which play an important part in human discourse processing. An additional argument will be given as to why cue phrases are a particularly appropriate kind of 'linguistic resource' to study.

The implications of this argument can be illustrated by giving an example. It would mean, for instance, that the existence in English of the word *however* points towards the existence in speakers of English of a *text structuring strategy* which can be signalled in text by using that particular phrase. Again, it is important to note that I am *not* suggesting that cue phrases can be relied upon to identify the relations in *actual* texts—as has already been pointed out, relations will not always be explicitly signalled. My suggestion is rather that the *set* of cue phrases in a given language can be used to determine a *set* of relations: how these relations are identified in particular instances is a separate question.

The claim being made is far from self-evident. Two quite different sets of things are being associated; a set of clause/sentence connectives and a set of cognitive strategies. Why should we assume that there is a connection? Or at least, why should we assume that the connection is strong enough to warrant cue phrases being used as a central source of evidence? Some researchers do assume such a connection without discussion—for instance in Sanders *et al*'s experiments, subjects' use of cue phrases is taken to reflect the relation they are thinking of. Yet it could be that the mapping between cognitive representations and cue phrases is not just a simple one-to-one: the relations people use might be more subtly delineated than is suggested by the range of cue phrases in a language; alternatively, they might be more broadly classified, so that the distinctions between cue phrases give a false impression of the accuracy with which relations are specified. These are real possibilities: if we want to make a direct link between cue phrases and cognitive constructs, we must provide an argument for so doing.

The claim that cue phrases mirror people's text-structuring mechanisms will be supported in three stages. Firstly, a clearer idea will be sought about what such mechanisms might be expected to be like, and what their role could be in the tasks of text creation and text comprehension. Next, a model of the process of 'communication via a text' will be advanced, in which the *communication of relations* between a writer and a reader plays an important role. Lastly, it will be argued that since the communication of relations is an important feature of communication via a text, it is to be expected that language contains ways of making relations explicit.

3.5.2 What Are 'Psychologically Real Relations'?

In this section, the nature of the psychological constructs we are looking for is examined in more detail. We are interested in how people represent the relations between segments of text, for the purposes of text comprehension and text construction, and how the representations are used in these tasks.

In one sense, the idea that 'people use relations to structure text' is almost trivially true: clearly, people form *some* representation of the relations between text segments, because texts are more than just collections of clauses, and people can recognise the difference between a coherent text and a collection of clauses. Somehow, the way clauses are combined in coherent text is being modelled: the real issue to address is what these models are like, and how the models influence text processing. And here there are many open questions. Do people use the same structuring strategies for reading as for writing? Do we really use a 'fixed set' of strategies, or are more general mechanisms in operation? Do we represent inter-segment relations using the same resources we use to represent the content of sentences, or are they treated in a different way? These issues and others will be addressed below.

It is important to note that we are not yet asking about *how to define* the relations which people use—this question is dealt with in the two following chapters. The present concern is rather to ask what sort of psychological constructs relations might be, and what processing purposes they might serve.

Relations In Discourse Interpretation

A basic preliminary observation is that people find it *easier* to process texts if they are more than just a collection of clauses. It is not just that we can recognise the difference between a coherent text and a collection of clauses; we actually *use* some representation of the relations in a text to help us process it. This can be seen, for instance, in a study by Meyer and Freedle (1984). Pairs of texts were prepared which differed only in the relations between groups of sentences: for example, in one text, the elements of content were linked by CAUSAL relations, and in the other, they were just presented as a collection. In this case, subjects' recall was much better for the former texts; this suggests that the presence of CAUSAL relations somehow facilitates text processing or storage.

Further studies show that the signalling of relations in a text facilitates its interpretation. For instance, Haberlandt (1982) shows that reading time is improved by the addition of linguistic markers; Segal *et al* (1991) show that the presence of interclausal connectives in a text helps subjects decide how to classify the connections between its clauses. We should not assume in interpreting such studies that *particular* surface cues mark *particular* coherence relations—this is the very claim that we are trying to justify. But the *general* finding that connectives facilitate discourse processing can be taken as a sign that relations between spans of text are somehow involved in the process: these experiments again show that it is important for readers to work out how text segments are linked together.

Alternative Conceptions of Psychological Reality

Sanders *et al* base their taxonomy of relations on a 'psychologically plausible' account of how the relations in a text are interpreted. They begin by claiming that relations of the kind proposed in RST are implausible as psychological constructs, because they are treated as unanalysed units:

> from a psychological point of view, Mann and Thompson's ideas are not very convincing, because they assume that all relational propositions are cognitively basic. If, for example, a relation like EVIDENCE occurs in a discourse, people interpret the discourse by referring to the cognitively basic notion of the EVIDENCE relation... Such an assumption is rather implausible.
>
> *Sanders* et al *(1992) p4*

The idea that people should have 'in the head' a complex construct such as the EVIDENCE relation is seen as unlikely. For Sanders *et al*, it is more plausible to *decompose* relations according to more 'general' principles of cognition such as causality and polarity: it is then claimed that such principles are used *jointly* to *infer* relations such as EVIDENCE. This is thought to be more plausible than the idea that each relation is a completely separate purpose-built construct.

However, it is not clear that general principles of cognitive organisation are the *only* ones which can be considered 'psychologically plausible'. A much broader conception of psychological reality seems possible, especially when we consider that reading and writing are highly skilled, practised activities.

Psychological studies of skill acquisition frequently point towards a model in which practice at a task leads to the development of *specialised mechanisms*, specifically tailored for performing the task in question. Such mechanisms have been posited in many different domains. For instance, Reason (1979) suggested that everyday tasks like driving and cooking are carried out by a system of **motor programs**, operating with a certain degree of autonomy. Models of linguistic processing also commonly involve sets of specialised constructs evolving during the course of practice: PDP models are a case in point. Consider, for instance, Rumelhart and Norman's (1982) model of the performance of skilled typists: here, every word in the lexicon is associated with an **action schema**, which when activated, sets off a chain of events leading to a sequence of keypresses. Such schemata develop as a typist's skill increases; they are clearly not 'cognitive primitives' in Sanders *et al*'s sense. And yet, they (or something like them) definitely seem to be 'used by people' in performing tasks.

The idea of a set of specialised constructs capturing regularities at the lexical level is extensible upwards to larger levels of structure. It is possible to imagine constructs corresponding to coherence relations evolving in the same way: producing and understanding large pieces of text are highly practised tasks, and it does seem plausible that mechanisms are developed specifically for them. When thinking of psychologically real relations, therefore, we do not have to limit ourselves to thinking about general 'cognitively basic' principles. However, this is not to say that general cognitive principles should not feature at all in a psychological theory of relations. A complete theory might well envisage relations *partly* in terms of general cognitive resources, and *partly* as learned strategies for structuring discourse.

Relations in Discourse Production

Sanders *et al*'s account of psychological reality is mainly based on the task of discourse interpretation: how coherence relations are inferred from passages of text. It is suggested that relations have a role in discourse production as well, but this role is not examined in any detail. This may be because less research has been done on how extended passages of discourse are produced—nevertheless, there are some studies, particularly about written discourse production, and it is instructive to consider how coherence relations can be fitted into existing theories.

Psychological theories of writing are still at a relatively early stage compared to theories of discourse interpretation, the difficulty being how to get an experimental handle on the process. Some studies (eg de Beaugrande (1984)) have used films of writers at work, and analysed pauses and crossings-out; but this source of evidence is unlikely to be rich enough to provide a full account of the processes involved. Most current theories are based on another technique—**protocol analysis**. In this paradigm, subjects are asked to 'verbalise' while they write, about what they are thinking about and how they are performing the task that they have been set.

Again, this source of evidence is questionable: it is unlikely that all the mechanisms involved in writing are amenable to verbalisation. Despite this, some initial theories of written composition have been formulated, which look plausible as far as they go. Perhaps the best known theory is that of Flower and Hayes (eg Flower and Hayes (1980), Hayes and Flower (1980)). This theory draws on spoken protocols, and also on the the notes taken by writers. It is suggested that the task of writing involves three separate sub-processes:

Planning: this process itself has three components.

- In the **generating** component, information relevant to the writing task is retrieved from long-term memory. The topic of the text is the initial search key; thereafter, the search key is the last item to be retrieved. The process stops when irrelevant items begin to be produced, and loops back to previous items to look for other relevant information in the same way.
- In the **organising** component, the most useful generated items are selected and organised into a writing plan. A number of **operators** are used to effect the organisation—for instance, 'identify as a possible first or last topic', 'search for a previously noted topic subordinate to present topic', 'order with respect to a previously noted topic'.
- In the **goal-setting** component, criteria are identified by which to judge the text, and they are stored for later use in **editing**.

Translating: this involves creating sentences out of the organised material.

Reviewing: again, there are two components to this process, **reading** and **editing**. The material so far produced is read, segment by segment, and each segment is edited in turn. The editing process detects and corrects inaccuracies in meaning, violations in writing conventions, and mismatches between the writer's intentions and those apparently expressed in the text.

The order of the processes is roughly as outlined above—although the editing process can cause the system to be re-entered at various different stages. In fact, editing and generating can interrupt any of the other processes.

Much of this model seems little more than a 'common sense' view of writing. But while it is clearly not very elaborate, it seems likely that the writing process is broken down in something like the way proposed. A number of hypotheses correlating stages in the process with the form of subjects' notes and protocols seem to be borne out. In which case it must be asked: how (if at all) do coherence relations fit into the model?

There are several possible answers. For one thing, relations could be involved in the **generation** component, as aids to the retrieval of relevant material from memory. If the content of one portion of text has been decided, then a coherence relation could give us useful information about other items of content which could feature in adjoining portions of text. A set of relations could provide a standard repertoire of methods for accessing relevant content. Evidence for this idea comes from a study of young writers by Bereiter and Scardamalia (1987). Children experience particular problems in finding content in their compositions: this study examines the effect of cueing children with suggestive prompts when they 'dry up'. The prompts given are predominantly cue phrases—for instance *even though*, *also*, *for example*. Given such cues, children are frequently able to produce more material for their compositions. Bereiter and Scardamalia suggest that the cues 'appear to stimulate the children to search for new nodes in memory that meet the logical requirements of the sentence openers' (p62). More mature writers, it might be supposed, will have internalised such cues, so that they can initiate the search for content themselves instead of waiting for an external prompt. We would not want to suggest that these internalised cues can be mapped directly onto surface cues. But this experiment does show that whatever form they take, methods for accessing new and relevant content are of great use in creating text.

Another possible role for relations is in Flower and Hayes' **organisation** component. The operators used in this component are used to structure elements of content already accessed. Many of them work with pairs of topics, for instance by deciding which of two topics to mention first, or whether one topic is subordinate to another. Coherence relations (which might well feature ordering constraints or incorporate hierarchical concepts like subordination) could be involved in making such decisions.[4]

It is interesting to note that the above uses for relations correspond quite closely to the uses found for relations in current text generation programs. Hovy (1988) uses relations as planning operators for working out text structure; Moore and Paris (1989) incorporate relations into planning operators, and in addition use these operators to access new elements of content to be generated. Many other subsequent systems have a similar design. It would be disingenuous to draw conclusions about human text processing mechanisms by looking at the way current text generation programs operate—but when constructs used in generation systems are found to resemble constructs hypothesised in psychological models, it seems worth mentioning the fact.

To sum up: it is plausible to think of coherence relations as modelling a set of strategies used by people to access and organise elements of content when planning text. The conception of relations as planning operators seems useful for psychological modelling as well as in computational systems.

A 'Basic Level' of Relations

One problem for the above idea is that it is hard to find any single relation that is not *sometimes* going to be useful for 'accessing and organising material'. For instance,

[4] A third possible role for relations is in the **translation** component, where the plans thus far built are converted into sentences. If relations were involved in the construction of plans, and at the same time associated with surface linguistic expressions, then the task of translation would be considerably facilitated: the relevant linguistic expressions would be *predetermined* by the plans. However, this idea cannot be used in the present argument, because—again—it assumes the conclusion we are trying to reach; that cue phrases can be used as evidence for relations.

consider a 'possible' relation like 'BROTHER-EVENT', where one span introduces a person, and the other span presents something happening to that person's brother. Such a relation would be helpful, almost by definition, in generating Example 3.4:

(3.4) Lars fixed the boat. His brother had holed it by bashing into the jetty.

Having decided on producing the first clause, BROTHER-EVENT could be used as a search cue to retrieve the material in the second clause. If coherence relations are to be thought of as strategies for accessing and organising material, why shouldn't BROTHER-EVENT count as a coherence relation?

One answer is that this relation is *seldom* a useful one when it comes to producing a coherent text. In most cases, if this relation is used, an incoherent text results:

(3.5) *Lars fixed the boat. His brother was a sergeant in the Danish police.

Furthermore, in the coherent cases, it is likely that an alternative relation can always be found. In 3.4 above, other relations such as BACKGROUND or CAUSE could have been used; and these are more often useful when it comes to structuring text.

But this explanation is not yet completely convincing. It could still be claimed that BROTHER-EVENT is useful *in particular cases*, such as Example 3.4. The relation 'BROTHER-EVENT' must have *some* kind of mental representation: if the reader or writer did not appreciate that it was Lars' *brother* who had holed the boat, then something would be missing from their representation of the text.

Of course, we can appreciate that in this example, Lars' action is not caused by just any event, but by an event involving his brother. Likewise, *each* CAUSE relation will be unique in certain ways: the point is that CAUSE is an *abstraction* from particular instances of relations. When I suggest that CAUSE has 'psychological reality', I am not claiming that we *only* represent relations at this level of abstraction, but that it is used *for some purposes*.

It is useful to take an analogy from another field of psychology at this point: Rosch's theory of categorisation (eg Rosch *et al* (1976), Rosch (1978)). Rosch claims that in order to form strategies for dealing with the infinite variety of stimuli we are faced with in the world, we have to work with abstractions, since we are finite processing devices, and we cannot have a particular strategy for each stimulus. Thus we have to treat some stimuli as equal: Rosch calls the level of abstraction at which we operate **the basic level**. The basic level is thought of as optimising the trade-off between useful categories and general categories. For instance, 'chair' is a basic level category: if we have to write down all the things you can do with a chair, the list will be much longer than the list of all the things you can do with a piece of furniture (a superordinate category), and not much shorter than the list of things you can do with an armchair (a subordinate category). Rosch claims that for some purposes, we work with the concept 'chair' rather than with more specific concepts. Yet at the same time, of course, we can recognise individual chairs and tell them apart. Different tasks call for reasoning at different levels of abstraction.

A similar point can be made for relations. It's just for some components of the text generation process that we abstract to the level of CAUSE or BACKGROUND. This does not mean that we can't recognise and differentiate individual instances of these relations for other purposes—of course we can tell the difference between individual instances of a particular relation. But text generation is a hard task, with lots of simultaneous constraints, demanding lots of processing: the ability to work with abstractions would be a useful one. In order to decide which abstractions to use, it would make sense to look again at the trade-off between utility and generality. Thus it is likely that BROTHER-EVENT is not so good at retrieving relevant information as CAUSE: so it would be less useful to work with this concept. In the same vein, a more specific relation like CAUSE-THROUGH-BROTHER'S-ACTION is not likely to be much more useful than CAUSE as a search cue for relevant material. Thus we would want to include CAUSE in our set of coherence relations, but not BROTHER-EVENT.

The above ideas should help to give substance to the idea that a set of coherence relations can be taken to model mechanisms 'used by people' when they produce text.

3.5.3 The Communication of Relations

Thus far, the production and the interpretation of discourse have been considered separately. In this section, we consider how the two processes come together, in what we might call 'communication via a text'.

For the purposes of the argument, it is not only important that writers use relations when creating a text, and that readers use them when interpreting it, but that they use the same relations. Otherwise it is impossible to argue that it would be helpful for writers to signal the relations they use in surface text. In order to make this argument, the model needed is something like the one in Figure 3.1: here, the *communication of relations* is seen as an intermediate step in the communication of a writer's goals and ideas to the reader. In such a model, the identification of relations is something which

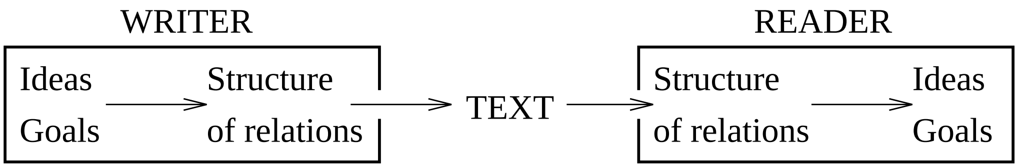

Figure 3.1: A Model of Communication Via a Text

really matters: readers need to be able to do it. It is this which makes it likely that ways exist for identifying relations explicitly.

What arguments can be given for the idea that writers and readers use the same relations? A number of points can be made. Firstly, if we assume that writers learn how to structure their texts by reading other peoples' texts, it seems likely that there will be an overlap in the methods used in reading and writing. More importantly, arguments can be given from the perspective of computational efficiency. The knowledge that a fixed set of relations is used to structure text permits a big reduction in the search space for both reading and writing. When a reader is working through a text, (s)he will know that the *next* segment of text will be related to the current segment in one of a fixed number of ways. If the set of relations includes a CAUSAL relation, then this is one of the alternatives the reader must always consider. Relations which are *not* in the set do not even need to be considered. Imagine a 'possible' relation which is unlikely to feature in the actual set; for instance the 'BROTHER-EVENT' relation mentioned above. The argument is, that at some stage of processing, the question of whether BROTHER-EVENT holds is not one which needs to be asked; whereas a relation like CAUSE is always going to be a possibility. Take a text like 3.6:

(3.6) Lars woke up early one morning. His brother was being noisy in the kitchen.

The obvious interpretation of the relation in this text is CAUSE, with the first span as the nucleus. But if BROTHER-EVENT were a relation in the set, then this would be an alternative: in this case, we wouldn't know which kitchen was involved, or when Lars' brother was being noisy in it. The text would be much more ambiguous, and it would be likely to be harder to process. It is a sign of how heavily we can rely on the conventional set of relations that we find it hard to even imagine a relation like 'BROTHER-EVENT'.

Writers will also profit from the conventional use of one particular set of relations. In example 3.6, the writer knows that enough has been done to disambiguate the relation in question, because the reader will not be expecting a relation like BROTHER-EVENT: so there is no need to make the causal relation explicit. In order to achieve coherence, texts only need to be specified with sufficient detail to allow the reader to work out which of the limited number of relations is being used.

From the point of view of computational efficiency, then, a strong case can be made for the use of a standard, smallish set of relations, by both readers and writers.

3.5.4 The Need to Signal Relations in Text

The final stage in the argument builds on the idea worked out in Section 3.5.3, that the communication of relations 'really matters'. It is shown how this idea can be taken to support the strategy of using cue phrases as evidence for the relations that readers and writers use.

The argument is as follows: if people actually *use* a certain set of relations when constructing and interpreting text, it is likely that the language they speak contains the resources to signal those particular relations explicitly. If people plan texts by building a structure of relations, and understand texts by working out this structure, then being able to mark relations explicitly in text will facilitate the communication process, by making it easier for a writer to indicate to a reader which relation is intended. As a consequence, ways of signalling relations in text would be extremely useful.

Of course, relations do not *always* need to be signalled in text. Often, they will be inferrable without explicit signals, as in example 3.6 above. The inference can be due to contextual information, or to the reader's general knowledge, or to the conventional use of a subset of possible relations. However, it is unlikely that any relation exists that is *always* inferrable in these ways. For one thing, the amount of relevant knowledge that the reader has is not always under the writer's control. Consider a text like 3.7:

(3.7) Bill was laughing. Frank was angry with him...

Here, if we don't know more about the situation, we don't know whether Bill is laughing because Frank is angry, or whether Frank is angry because Bill is laughing. It is useful to have linguistic devices, like the cue phrases *because* or *as a result*, to distinguish these possibilities.

It is not strictly true to say that the reader's knowledge is beyond the control of the writer. In the above example, the writer could have given enough prior context to disambiguate one or the other reading. But using cue phrases is a much simpler way of providing the necessary information; from the point of view of efficiency, it is an attractive option.

So—again from the point of view of efficiency—we can argue that if the communication of relations is important, then simple linguistic means (such as cue phrases) will exist for identifying them in text. If this is indeed the case, then we can look at the range of cue phrases in a language to give us an indication of the relations that people use in constructing and interpreting texts in that language.

3.5.5 Summary

To sum up the argument that has been presented:

- In Section 3.5.2, the idea that 'people use a set of relations when they process text' was fleshed out. A conception of relations as constructs developed during the course of practice was proposed. Rosch's notion of the basic level was invoked, to illustrate how a particular level of abstraction is suitable for particular tasks. It was also used to emphasise that just because we work at this level of abstraction for some purposes, it does not mean that for other purposes we cannot represent the relations in discourse with more finer detail.

- In Section 3.5.3 it was suggested that if readers and writers use relations, then it is likely they use the same set. The argument was on the grounds of efficiency: the conventional use of one particular set would reduce the search space for both reading and writing tasks.

- Finally, in Section 3.5.4, it was argued that if the communication of relations is of real importance in the communication of a writer's ideas via a text, then it is likely that simple linguistic means exist for identifying relations explicitly. Again, the argument was on the grounds of efficiency: language, if it is an efficient communicative tool, should contain such devices.

- If the above arguments go through, then it is permissible to take connective devices like cue phrases as *evidence* for the set of relations that people use.

3.6 Some Objections to the Argument

The argument might be considered quite speculative in some places; particularly those which make reference to considerations of processing efficiency, or which propose the investigation of psychological constructs through an analysis of language. Some of these issues are addressed in this section.

3.6.1 Can you really investigate psychological constructs without doing any psychological experiments?

As a method of studying psychological constructs, the technique being proposed is very unusual. How can we expect to find out about human processing mechanisms just by studying linguistic phenomena? How could we possibly confirm or refute any conclusions we came to? Almost without exception, psychological theories are based on data about how humans perform experimental tasks. In this technique, 'a language' is the source of all relevant data: is this really permissible?

Admittedly, the technique is unconventional as a psychological methodology. However, many objections have already been mentioned concerning the experimental paradigms used to study human text structuring mechanisms. To sum up some of the points in Section 3.4.2: when subjects are asked to talk explicitly about the relations they use while they are processing text, it is not certain that the tasks they perform are exactly those they normally perform. If subjects are 'naive', it is not certain that their understanding of what coherence relations are is sufficiently clear; if they are discourse analysts, it is possible that their intuitions are tainted by biases towards one theory or another. There are methods of studying relations without making subjects think about them explicitly—in particular, recall and reading time can be examined. But these methods are indirect: they can only be used to choose between various hypotheses, not to form hypotheses in the first place. And finding an experimental indicator of the relations used in generation promises to be an even harder task. Cue phrases cannot themselves be used, until an argument is given to suggest why cue phrases and cognitive constructs should stand in a one-to-one relation with each other. But if we have such an argument, then there is little point in conducting an *experiment* using cue phrases as evidence for relations: the argument allows us to look *directly* at the cue phases in the language to find out about the constructs people employ.

To put the positive case for studying language to find out about psychologically real relations: it just seems that language provides an enormously rich source of evidence for studying the ways in which people structure text, and it would be a pity to ignore it. As the next two chapters show, the set of cue phrases in English is very large and diverse, and it is structured in all sorts of interesting ways. Some very subtle distinctions between relations are captured by the different phrases—in comparison, the experimental methodologies outlined above seem relatively impoverished sources of data. Of course I am not saying that psychological experiments cannot *also* be used to investigate the relations people use. I am just arguing that the method of analysing cue phrases is a legitimate (and attractive) alternative.

3.6.2 It might be *useful* if there were a cue phrase for every relation people use: this doesn't mean there *will* be one.

This objection goes to the heart of the argument being advanced. It suggests that we are not permitted to rely on language being perfectly efficient as an instrument for the communication of relations. What grounds are there to suppose that cue phrases *will* exist in a language just because it would be *useful* if they did? For one thing, we need to be certain that language is flexible enough to adapt to the requirements of its users. On top of that, it may not be just a matter of flexibility: maybe efficiency isn't the only constraint on the way languages change. Other conflicting factors may also be involved—for instance, a stylistic trend in favour of terse, simple sentences—such that the end result only *partially* satisfies some constraints. The argument could, in short, be considered Panglossian, affirming that 'language has evolved to be the way it is, and therefore the way it is must be the most efficient'. If this were the case, it might well put the proposed methodology in jeopardy.

We would not want to have to rely on 'the efficiency of language' as an article of faith. Fortunately, there do seem to be at least some indications that cue phrases have evolved through considerations of efficiency. For one thing, many cue phrases are single words, or **idiom chunks**,[5] which are conventionally treated as single words. Moreover, the etymology of many cue phrases suggests that they have evolved from longer, less formulaic phrases. For instance, Halliday and Hasan (1976, p230) list several connective words which were originally more complex phrases, containing anaphora, and making use of the 'compositional' resources of the language. Such words include *therefore*, *thereupon*, *whereupon*, and so on. In addition to this, there are cases of new cue phrases being invented where there is a need for them. For instance, the phrases *iff* and *just in case* have been coined in logical and philosophical genres of text respectively, to replace the longer and more unwieldy phrase *if and only if*. There are many examples of cue phrases evolving in this way: they provide some support for the idea that the set of cue phrases will reflect the set of relational constructs we use.

3.6.3 Cue phrases aren't the only way of signalling relations.

As far as it goes, the argument gives no reason why cue phrases should be studied to the exclusion of other linguistic devices for signalling text structure. In fact, the principal criticism made about the present methodology (Bateman and Rondhuis (1994), Seligman (1994)) has been that it concentrates exclusively on cue phrases.

There are certainly other means of signalling discourse structure. A great many researchers (see e.g. Moens and Steedman (1988), Lascarides and Asher (1993)) have shown that tense and aspect are cues to the temporal structure in a text. Scott and de Souza (1990) have explored a large number of syntactic devices as signallers of RST relations. Several researchers (e.g. Sidner (1983), Grosz and Sidner (1986)) have suggested that the pattern of pronominalisation in a text provides information about its thematic structure. Delin and Oberlander (1992) investigate the discourse structures

[5] Phrases whose meaning is not a function of the semantics of the individual words of which they are composed.

signalled by *it*-clefts. In short, it is uncontroversial that discourse structure can be expressed through a wide range of surface linguistic devices.

In the light of these examples, the present emphasis on cue phrases should really be seen as just a starting point in working out a theory of discourse coherence. Ideally, other cohesive strategies should also be considered. However, the general idea that the investigation should be driven by the range of available resources remains the same. In addition, there are some good reasons for beginning by looking at cue phrases. For one thing, cue phrases are a relatively homogenous set of linguistic devices, whose different effects can be fairly easily compared amongst one another. Furthermore, many alternative cohesive strategies also have cue phrase counterparts. For instance, temporal relations can be signalled by phrases like *previously*, *afterwards* or *while*. Finally, it can be argued that structures which cannot be signalled by cue phrases should not be modelled using coherence relations, but with some other theoretical construct. This argument will be taken up in Section 7.2.

3.6.4 Different languages have different cue phrases.

In cross-linguistic studies, it has been found (eg Ballard *et al* (1971), Longacre (1983)) that different languages often have different sets of cue phrases. Longacre, in justifying his set of relations, uses cue phrases from a number of different languages as evidence. This goes against the argument being suggested here: if a cue phrase doesn't exist in a certain language, we would have to say there was no need for it. This in turn would mean advocating different structuring mechanisms in people speaking different languages.

However, this objection is not insurmountable. For one thing, the differences between the sets of cue phrases in different languages are not *that* great. It is quite surprising how much similarity there is between English cue phrases and those in the Philippine languages which Ballard and Longacre study. And between European languages, which are much more closely related, it is likely that the differences are even smaller.

Furthermore, differences between the sets of relations in two languages might be attributable simply to the different registers of discourse used in these languages. In spoken language, for instance, phrases like *lastly* or *to summarise* are very rare; maybe we would not expect to find them at all in a language with a largely oral tradition, or in one without any writing at all. All this means is that speakers of this language have not needed to internalise the constructs developed by English speakers in the course of learning how to read and write.

Finally—to pre-empt some of the discussion in later chapters—while cue phrases from one language may not always translate directly into cue phrases in another language, it may still be that the dimensions along which cue phrases *vary* are the same in the two languages. Thus, at a more abstract level, interesting similarities may still be found. To take a simple example, there is no single cue phrase in English to translate the German cue phrase *wenn*; two phrases (*if* and *when*) are needed to do the job. However, there are *other* phrases in English—for instance *then*—which seem to manifest just the kind of ambiguity that *wenn* does. See Chapter 6 for further discussion of this idea.

In any case, it is not inconceivable that different language communities use different structuring mechanisms even for producing similar registers of discourse. As mentioned in Section 3.5.3, one psychological role for relations could be as *conventions*, such that the reader knows what set of relations the writer is working with, and the writer knows which relations the reader will be expecting: different languages could make use of different conventions. It is important to remember that 'psychological constructs', in the sense that we are talking about them here, do not *have* to be 'cognitive primitives' posessed by all humans. An equally plausible idea is that relations model a mixed bag of constructs, some of which are cognitive primitives and some of which are acquired through exposure to a particular language community.

3.6.5 What about relations between large segments of text?

One of the attractive features of coherence relations is their insensitivity to span size. Indeed, this feature was taken in Chapter 2 to be central to the notion of coherence relations. However, cue phrases seem primarily designed for linking clauses or sentences together. In linking relations to cue phrases, do we not risk ignoring the issue of higher level relations?

This objection is telling if we adopt a straightforward conception of cue phrases as clause or sentence conjunctions; however, we will actually be working with a more general and informative definition of cue phrases (to be outlined in the next chapter, in Section 4.2). This definition moves away from a syntactic conception of cue phrases, allowing more complex phrases such as *this is because* and *following this*. Stock phrases such as these, which make use of propositional anaphora, can signal relations between large spans of text, and hence (according to the argument in this chapter) can be used as evidence for identifying those relations.

3.7 Summary: A New Proposal for Motivating Relations

In this chapter, a methodology has been suggested for working out a 'standard' set of coherence relations. Coherence relations are thought of as psychological constructs used in planning and interpreting discourse, and it is suggested that the cue phrases in a language can be used as evidence for these constructs.

This argument leaves us with a fairly strong general prediction; namely, that the set of relations which corresponds to the set of cue phrases will suffice to describe the relations in all coherent texts. In the remaining chapters, this prediction will be refined and tested.