Using Linguistic Phenomena to Motivate a Set of Rhetorical Relations

Alistair Knott
Department of Artificial Intelligence, University of Edinburgh
80 South Bridge, Edinburgh EH1 1HN, Scotland
Email: A.Knott@ed.ac.uk

Robert Dale
Human Communication Research Centre, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW, Scotland
Email: R.Dale@ed.ac.uk

May 5, 1993

Running head: Motivating Rhetorical Relations

Abstract

The notion that a text is coherent in virtue of the 'relations' which hold between the elements of that text has become fairly common currency, both in the study of discourse coherence and in the field of text generation. The set of relations proposed in Rhetorical Structure Theory (Mann and Thompson [14]) has had particular influence in both of these fields. But the widespread adoption of 'relational' terminology belies a certain amount of confusion about the relational constructs themselves: no two theorists use exactly the same set of relations; and often there seems no motivation for introducing a new relation beyond considerations of descriptive adequacy or engineering expedience.

To alleviate this confusion, it is useful to think of relations not just as constructs with descriptive or operational utility, but as constructs with psychological reality, modelling real cognitive processes in readers and writers.

This conception of rhetorical relations suggests a methodology for delineating a set of relations to work with. Evidence that a relation is actually *used* by speakers of a language can be obtained by looking at the language itself—in particular by looking at the range of cue phrases the language provides for signalling relations. It is to be expected that simple methods will have evolved for signalling the relations we find most useful.

This paper presents a bottom-up methodology for determining a set of relations on the basis of the cue phrases which can be used to mark them in text. This methodology has the advantage of starting from concrete linguistic data, rather than from controversial assumptions about notions like 'intention' and 'semantics'.

1 Introduction

The notion that a coherent text is more than simply a sequence of sentences is widely accepted; however, a number of different approaches have been taken to explaining what it is that provides textual coherence.

One fairly common idea is that there are implicit **relations** between the sentences of a text, so that (for example) the content of one sentence might provide justification, elaboration or explanation for the content of another. These relations are held to bind a text together, and thus contribute to its overall coherence. They are sometimes made explicit by the speaker's use of particular **cue phrases**, linguistic elements like *more precisely* or *for example*; but in other cases the hearer is left to infer the particular relation that holds between two sentences on some other basis. We will refer to theories that make use of ideas of this kind as theories of **discourse relations**: Grimes' [4] **rhetorical predicates**, Martin's [15] **conjunctive relations** and Hobbs' [7] **coherence relations** would all fall into this category.

Another fairly common idea in the literature is the notion that at least some kinds of texts have an underlying hierarchical structure. Again, this idea is not in itself controversial; but different approaches have been taken to providing an explanation of the nature of these hierarchies, and the relationships they embody. We might call such theories theories of **discourse structure**. Recent examples of this kind of theory include that of Grosz and Sidner [5], who define textual hierarchy in terms of the **dominance** and **satisfaction-precedence** relationships between the underlying intentions of the speaker; and that of Polanyi [17], who proposes related notions of **subordination** and **co-ordination**.

Some theories make use of both of these ideas: this is to some extent true of both Hobbs' notion of discourse structure, which permits trees of relations, and Grimes' predicates, which can take as arguments anything from lexical items to 'large semantic subtrees'. In the field of computational linguistics, and particularly in the subfield of natural language generation, Mann and Thompson's Rhetorical Structure Theory (RST) [14] has been the most influential theory of this kind. A particularly interesting characteristic of RST is the claim that the same rhetorical relations that hold between individual sentences also account for the relationships between larger segments of text; indeed, the theory claims that the same relationships can hold irrespective of the size of the **text span**s involved. This step makes it possible to characterise the structure of a text in terms of a hierarchy of rhetorical relations that connect together the parts.

RST is not itself a 'computational' theory; like many other theories of discourse structure, it calls extensively on the intuitions of a discourse analyst. However, it has recently been adapted for procedural use in a number of natural language generation systems: following early systems by Hovy [8] and Moore and Paris [16], many subsequent text planners have incorporated aspects of the theory (for example, Vander Linden *et al* [28] [27], Rösner and Stede [19] [20]).

In the present paper, we take RST as our starting point. We begin in Section 2 with a brief description of RST's conception of relations,[1] and an attempt to explain why it has such appeal in the domain of text planning. In Section 3 we outline some problems which result from the widespread adoption of the relational concept; these arise from the lack of a standardised and well-motivated set of relations. Section 4 surveys some current attempts to motivate a standard set of relations, and notes some of their shortcomings; finally, in Sections 5 and 6, we present a new methodology for motivating a set of relations, and based on this methodology, we provide a preliminary taxonomy.

2 Rhetorical Relations

According to RST, a written text is coherent (partly) in virtue of the relations which hold between its sub-parts or **span**s. The same relations can link spans as small as clauses or as large as whole paragraphs or chapters: RST gives a catalogue of twenty-three relations, which allegedly suffice to analyse 'the vast majority' of English prose texts. A sample RST analysis is given in Figure 1.[2]

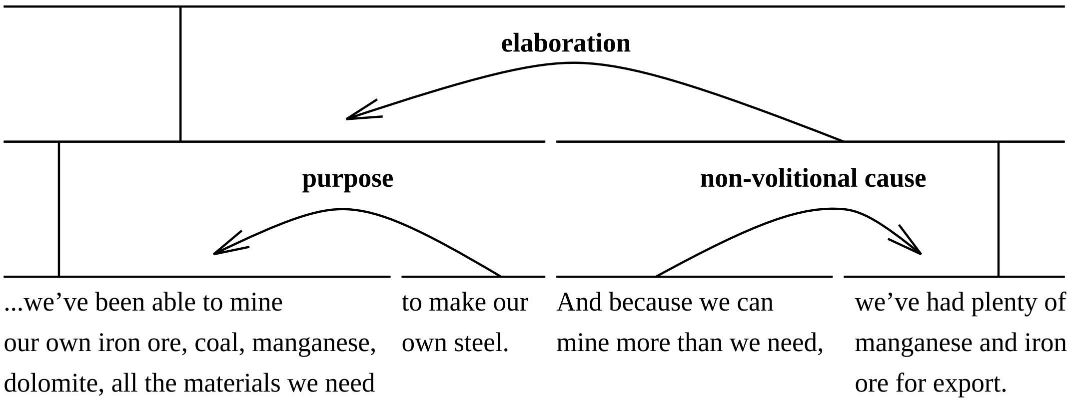

Figure 1: Example of a Rhetorical Structure Tree (from Mann and Thompson [14])

In RST, relations are defined principally in terms of the likely effects of the two related spans on the reader; for instance, whether understanding one span increases the reader's readiness to accept the other, or whether understanding both spans allows the reader to recognise a particular semantic relation as holding between them. This emphasis on effects is clearly central to the writer's intentions in creating a text; the purpose of a text is to achieve a particular effect on a reader. However, not all relations considered in RST are

[1] We will use the terms 'relation' and 'rhetorical relation' interchangeably from now on.

[2] The relations illustrated here link spans asymetrically: in each case a **nucleus** span (which has an arrow pointing to it) is related to a **satellite** span (from which the arrow originates). The claim in RST is that satellite material is 'less important', and that it serves a subsidiary function with respect to the nucleus.

purely 'rhetorical' in this sense: some relation definitions make reference to subject-matter as well—whether a span presents an action or a state, for instance, or whether there is a causal relationship between the situations presented in the spans.

RST's relations are attractive constructs for several reasons. For one thing, referring to intentions permits a more fundamental analysis of text: speech act theory (Austin [1], Searle [25]) makes clear the descriptive inadequacies of purely truth-functional analyses. For another thing, the relations allow a uniform compositional account of text. They can apply between spans as small as clauses, or between larger spans such as paragraphs or groups of paragraphs, based on the assumption that intentions to produce small pieces of text are qualitatively similar to intentions to produce larger pieces. On the basis of this assumption, it is possible to describe text recursively, using a hierarchical structure of relations.

Given these features, it is clear why rhetorical relations have proved popular in text generation. Hovy [10] calls them 'the basic building elements of coherent text'. The particular role that relations play in text planning varies from system to system: in some (eg Hovy [8]), they act as STRIPS-style planning operators which are used to structure a text whose content has already been determined by some prior process; in some (eg Moore and Paris [16]) they are also used to guide content selection; and in some (eg Vander Linden *et al* [28], Rösner and Stede [19]) they are lower level guides to the syntactic realisation of sentences.

3 The Proliferation of Relations

Despite the attractions of the *concept* of rhetorical relations, there is little agreement between researchers about which *individual* relations to make use of. A text planner needs to specify a set of relations to work with; however, there are significant differences between the sets used by the various systems. Although the work inspired by RST takes the initial set of relations as proposed by RST as a starting point, each individual researcher uses relations not found in the original theory. These are sometimes specialisations of existing RST relations (for example, Rösner and Stede's CAUSAL-CONDITION, which is a more specific version of the original RST CONDITION relation); sometimes generalisations over more than one relation (for example, Vander Linden *et al*'s RESULT, which is an amalgam of the RST relations VOLITIONAL-RESULT and NON-VOLITIONAL-RESULT); and sometimes completely new relations (for example, Rösner and Stede's UNTIL, which is customised for instructional texts, where a given action must be performed until a certain condition is met). In other cases, relations with the same name are defined differently—for instance, Hovy's SEQUENCE has a nucleus and a satellite while the RST SEQUENCE is multinuclear; and Rösner and Stede's SEQUENCE does not cover sequences of instructions while RST's SEQUENCE does.

Recently, Hovy [9] collected a large set of relations from other researchers, and organised them into a taxonomy: the idea was to produce a 'standardised and covering set of relations'. The taxonomy has subsequently been reworked by Maier and Hovy [12] and Hovy *et al* [11].

The surprising feature of these classifications is the number of relations they involve: the latter study uses a systemic network with over 80 relations at the leaves and over 170 relations in the network as a whole. This is clearly a substantial departure from RST's original set of relations.

3.1 Reasons for the Proliferation of Relations

This proliferation of relations is partly to be expected: its origin can be traced to a number of features of RST. Firstly, Mann and Thompson are themselves quite flexible about the set of relations in their theory; they are more concerned with establishing the rhetorical relation in general as a useful tool for text analysis.

> Relation definitions have the status of applications of the theory rather than elements of the theory. One might want to change or replace the definitions. . . such changes are to be expected and do not cross the definitional boundaries of RST.
>
> Mann, Matthiessen and Thompson [13], p48.

Secondly, the set of relations in RST is 'in principle open'. Mann and Thompson claim that a small number of relations is sufficient for analysing the great majority of texts; and that the exceptions involve rare relations seldom seem in normal texts. Yet little is said about how many rare relations there might be, and how rare they have to be before being discounted. The size of the working set of relations is specified as 'small', but there is no absolute indication of what 'small' means in this context.

Thirdly, RST relations are claimed to be 'purely descriptive' constructs. The motivation for the set of relations in RST is that it allows an analysis of virtually every text: you can represent nearly every text as a tree of relations, using the 23 relations RST provides. Relations are thus treated as a device for representing the structure of texts. But this criterion serves to justify a great many different sets of relations, of which the RST set is only one. For instance: the space of texts which RST can describe can equally well be described by a smaller set of relations obtained by amalgamating relations in the standard set (for example, by combining the four relations in the 'cause cluster'[3] into a single CAUSAL relation), or by subdividing the existing relations to obtain a larger set (for example, by splitting CONDITION into TEMPORAL CONDITION and CAUSAL CONDITION). More radically, one could use just two relations (for example, CAUSAL and NON-CAUSAL) to describe all the texts which RST can currently account for. (This sounds extreme, of course, and we would want to say that there is a loss of explanatory value if we take a step of this kind; but the point is this: if relations are taken simply as useful *descriptive* constructs, there is no way of justifying the granularity of the description they provide.) Finally, it is possible to define relations which cut right

[3] These are the relations VOLITIONAL CAUSE, NON-VOLITIONAL CAUSE, VOLITIONAL RESULT and NON-VOLITIONAL RESULT.

across the RST set. For instance, Scott and Souza's LIST relation [24] is hard to describe in terms of the relations given by Mann and Thompson [14]: Scott and Souza use this relation to expand the satellite of a superordinate relation into a number of elements, where Mann and Thompson might be more likely to use multiple satellites for the same nucleus. Either system is 'descriptively adequate' for the same corpus of texts: any text which can be analysed by one can be analysed by the other. But from this criterion alone we have no reason to choose between the systems.

3.2 Problems with the Proliferation of Relations

It is clear that RST is very flexible about the set of relations to be used by analysts. Such flexibility is even held by some to be an advantage of the theory:

> Generally a new domain has dictated modifications to the inventory of relations, but this very adaptability is one of its most useful features.
>
> Vander Linden *et al* [28]

At the same time, giving analysts such a free hand to choose relations creates some serious problems.

3.2.1 The Unfalsifiability of the 'Relational Hypothesis'

Flexibility about individual relation definitions tends to undermine the empirical content of the relation construct in general. The hypothesis that 'virtually any text can be analysed by representing its rhetorical relations' becomes much less strong if relations can be created whenever they are needed: it is hard to think what evidence could be found which could disprove it. Rhetorical relations at this point hardly seem to be saying any more than speech act theory; that we must take intentions into account when representing text. The extra claim in RST—that text is coherent by virtue of the relations between its intentions—is virtually unfalsifiable without a method for specifying what is to count as a relation in the first place.

Even incoherent texts can be analysed according to the relations between the intentions in their spans. For instance, the text in (1) seems incoherent at first sight:

(1) John broke his leg. I like plums.

Yet we could still define a relation which holds between the intentions underlying the spans in this text: perhaps we could call the relation **inform-accident-and-mention-fruit**. The relation might be multinuclear, with one nucleus intended to convey information about an accident, and the other intended to convey information about fruit. Clearly, we do not want to include these sorts of relations in any principled set of 'rhetorical relations'.

To place the relational claim on a sounder empirical footing, we need to tighten the constraints on relationhood—for instance, by giving a fixed set of relations, or by giving rules for picking out a set of relations from some larger set of 'potential relations'.

3.2.2 What are Relations Modelling?

The proliferation of different and incompatible sets of relations seems odd from a theoretical point of view: should we really have such a choice about which relations to use? Theoretical constructs are typically treated as corresponding to *real phenomena* underlying the data they describe, rather than as being purely synthetic: otherwise there seems little point in using the constructs at all. But in the present case, the emphasis on descriptive adequacy alone leaves no room for such realism: consequently, little attention is paid to the question of what relations actually model. This means it is possible that different researchers envisage different roles for their relations (see the next section for examples), in which case it may make for confusion to use the label 'rhetorical relation' as a general term.

More seriously, there is a danger that researchers will look no further than descriptive adequacy in deciding on a set of relations to work with. This might happen if RST's existing set were adopted, since it is just one of many descriptively adequate sets. It will also happen if RST's set is altered simply to achieve descriptive adequacy in a different domain. In such cases, it is hard to speak of relations modelling anything at all; and hard to see why the construct is introduced in the first place.

Of course, the real situation is not quite as arbitrary as the preceding discussion would imply. There is still a reasonable consensus between researchers about which relations to use. But this very fact suggests that intuitions are at work which are not being acknowledged: while researchers rarely attempt justification for their choices beyond an adherence to the needs of descriptive adequacy, their choices are not as diverse as they might be given *only* this criterion.

4 Motivating a Set of Relations: Some Current Initiatives

Unless a more explicit approach is taken towards justifying a set of relations, the research programme based on RST may lose much of its momentum. What is needed is a way of giving the concept of a relation more substance; two current suggestions to this end are discussed below.

4.1 Relations as Operational Constructs in Text Planning

As mentioned at the outset, relations, while originally purely descriptive constructs, have been adapted to the procedural task of text planning. In Hovy's [8] text planner, for instance, relations are used as planning operators, which control how textual elements are incorporated together into a single coherent paragraph. A hierarchy of relations is built up, with clause-sized semantic elements at its leaves. New elements are worked into this structure by 'expanding' one of the leaves by a process of adjunction; that is, by replacing the leaf with a relation, whose nucleus is the semantic element originally at the leaf and whose satellite is a new semantic element. The final tree of relations is traversed left-to-right, with each leaf in turn being passed to the sentence generator PENMAN, and the sentences are linked with appropriate connectives.

Clearly, the set of relations used in such a text planner is a major factor in determining the nature of the texts which it produces. So one way to justify a set of relations is to consider its effectiveness when used as a component of a text planning system like the one just outlined.

Hovy's current classification of relations (Hovy *et al* [11]) seems at least partially justified in this way. His relations are arranged into a hierarchy—a top-level distinction between IDEATIONAL, INTERPERSONAL and TEXTUAL relations is motivated on theoretical (Hallidayan) grounds; but lower down, the practical uses of relations within the planner are given more emphasis. Indeed, it would be difficult to argue *a priori* for each of the hundreds of relations lower in the hierarchy.

One important use of relations in the planner is as guides to the surface constructions to be used in the text which it produces. Relations are associated with the textual devices which can be used to mark them (such as sentence connectives): this means that once a structure of relations has been built, the step of converting it to a surface string of text is simplified. Textual realisation methods can also be used to arrange relations in the hierarchy: in an earlier paper (based on a slightly different taxonomy), Hovy suggests that cue words which can mark a given relation can also be used to mark any of its subordinates (Hovy [9]). For instance, *then* (which marks a SEQUENCE relation) can also be used to mark the subordinate TEMPORAL-SEQUENCE relation).

It is clearly expedient for the text planner to use relations which can be associated with methods of textual marking. But there remain some questions about this approach. Firstly, does the hierarchy really conform to the principles of substitutability of cue words? Hovy's examples are very token, and only involve a handful of words. More importantly, are there any good theoretical reasons why a taxonomy of relations *should* be built on the basis of the cue words which can be used to mark them? Hovy admits that his evidence from cue words is 'nonconclusive'. Further work is needed here, both in the detailed justification of the taxonomy, and in the justification of the policy of relying on cue phrases as evidence for a classification. In the present conception, relations are still to be seen mainly as heuristic constructs; it is not clear what theoretical or explanatory significance they have.

4.2 Relations as Psychological Constructs

An alternative approach to motivating a set of relations begins with a direct suggestion about what relations are: it takes them to be *psychological* constructs, which humans make use of when creating and interpreting text. This approach can be seen as sharing Hovy's conception of relations as constructs used in the task of text planning, but making an additional stipulation—that relations are to be those constructs which *people* use to create text, not just some arbitrary set which a text planning machine can use.

Associating relations with the *human* text planning mechanism endows them directly with explanatory significance. It is plausible that texts have the structure they do *as a result of* the strategies we use for constructing them. Texts are human artefacts, constructed by humans for humans: an explanation of the way they are structured can be expected to make reference to the psychological machinery involved in processing them. The difficult question is deciding which set of relations we use as part of this machinery—if indeed we use relational constructs at all.

Hobbs [7] was perhaps the first to suggest an explicitly psychological account of relations. He refers to his relations as 'text-building strategies', employed by the writer to ease the reader's comprehension process. The relations are motivated by a study of the type of inferences a reader needs to make in order to understand a discourse: thus they are seen as being tailored to the reader's needs. It is suggested that relations might eventually be explained in terms of very general principles of 'cognitive economy'—but what these principles are is not considered in any detail.

Sanders, Spooren and Noordman [21] [22] also see relations as modelling psychological phenomena; but they go further than Hobbs, giving concrete suggestions about what these principles might be, and looking for experimental evidence for them. They seek to describe the space of relations using a small number of 'cognitively basic' concepts, which act as orthogonal parameters in defining individual relations. The primitive concepts are as follows:

- 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**.

- 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.

- 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.

- 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 expecation derives from a causal basic relation; or a contrast, where the basic relation is additive.

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.

Support for this set of relations is provided by two 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 where there was disagreement over which relation to use, it tended to be over the value of a single parameter only.

These experiments provide good support for Sanders *et al*'s use of parameters to divide up the relation space. They also embody a sound methodology for justifying a set of relations: statistical tests on analysts and naive subjects. But the actual parameters they choose present some problems.

Firstly, the empirical evidence for the 'source of coherence' parameter is less clearcut than for the other parameters: both analysts and naive subjects are apparently less attuned to it.

Secondly, the status of the 'order of spans' parameter as a 'cognitively important' primitive might be questioned. Sanders *et al* cite several studies which show the *effects* of presenting information to subjects in different orders; but such results do not permit the assumption in subjects of a specific mechanism attuned to span order. The *a priori* grounds for supposing such a mechanism seem weaker than in the case of the 'basic relation' parameter: causal relations are likely to play an important part in our knowledge representation systems, and in the way we reason with them. On the other hand, 'order of spans' seems tailored particularly to the explanation of linguistic phenomena. Clearly analysts can use it, and writers are sensitive to it—but that doesn't *imply* that it is itself a 'cognitive primitive'.

Finally, it might be thought that Sanders *et al*'s methodology of psychological experiments is too blunt a tool to derive all the subtleties in a classification of relations. All in all, twelve types of relations are picked out in a principled way—yet Hovy's text planner seems to require hundreds of relations. The addition of new orthogonal parameters would soon boost the number of relations, but as we have seen, there are already potential problems with the evidence for some of the current parameters. Furthermore, there is no reason *a priori* to expect the space of relations we make use of to be neatly parameterised.

4.3 Summary

Hovy *et al* and Sanders *et al* adopt very different approaches to the motivation of a set of relations. In Hovy's work, a taxonomy is built to reflect surface linguistic phenomena, but its explanatory significance is open to question; in Sanders' work, relations are given an explanatory role as psychological constructs, but the experimental methods for investigating these constructs are fairly blunt.

Our research, described in the following sections, attempts to combine elements from both these approaches: it suggests that evidence for 'psychologically real' relations can be sought in a study of linguistic phenomena.

5 A New Approach to the Motivation of Relations: Looking for Linguistic Evidence for Cognitive Strategies

If people actually *use* a particular 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.

A variety of sources of evidence suggest that languages are 'optimised' in some ways so as best to reflect the communicative needs of their users. For instance, Zipf [29] has shown that commonly used words tend to be shorter than uncommon ones: this clearly makes for communicative efficiency, because it minimises the average length of the words we use. He suggests how a language might evolve so as to maintain this position, through the introduction of new words or abbreviations, and the abandoning of obsolete ones. If relations are useful constructs in the communication process, we might expect language to have evolved in a similar way, to give us simple and convenient ways for marking them. If this is indeed the case, then we could use the relevant linguistic strategies as *evidence* for relations.

Note that this argument only goes through if we assume relations are constructs people actually use. If they are thought of simply as devices for describing text (as they are for instance in RST), there is no way of arguing *why* they should be linked to a certain category of linguistic phenomena. It may be convenient, but other sets of relations are equally justified in principle. The assumption of psychological reality, coupled with the assumption that language is somehow optimised for our communicative purposes, together constitute an *argument* why we can take linguistic marking techniques as evidence for a set of relations.

If the assumptions are accepted, a new methodology for determining a set of relations can be suggested. Studying the means available for marking relations in a given language should be able to tell us about the relations which people actually make use of. The methodology might be described in Hallidayan terms as using the **cohesive** devices a language affords as evidence for a psychological theory of text **coherence**.

The appeal of this approach is not just to do with its emphasis on explanatory adequacy. It also has the advantage of starting from concrete linguistic data—the methods of signalling relations actually used in text. Typically, classification of relations has been done by defining them using concepts like 'intention' and 'semantics', whose representation is itself the subject of much debate. Using the present methodology, it is only after developing a taxonomy of linguistic data that is it necessary to characterise relations using these more abstract terms: a given linguistic strategy can be taken as evidence for the *existence* of some structuring mechanism which corresponds to it, even before we have a way of defining it in theoretical terms at all.

Having argued that a set of relations can be motivated by appealing to certain linguistic strategies, it remains to identify a set of linguistic strategies to study. We turn to this below.

5.1 Cue Phrases: Linguistic Markers of Relations

We concentrate here on a single category of linguistic data, that of **cue phrases**. These are perhaps the most obvious means of signalling relations in text, and many other researchers make reference to them under one name or another (Cohen [2]: 'clue words'; Grosz and Sidner [5]: 'cue phrases'; Schiffrin [23]: 'discourse markers'). It is not certain that all of these researchers have the same conception of the role of cue phrases, or even that they agree about what counts as a cue phrase—for the moment we can think of them pre-theoretically as 'phrases whose function is to link spans of discourse together'. Examples of cue phrases are *because*, *meanwhile*, *but*, and *on the other hand*.

We have taken this group of linguistic phenomena as the focus of our study for a number of reasons.

- Cue phrases form a reasonably homogeneous syntactic group, which it is possible to isolate as the 'object of investigation'.

- Cue phrases are very versatile in the range of relations they can signal. For instance, *because*, *meanwhile*, and *on the other hand* all signal very different relations. There is also quite a large number of them, so a detailed taxonomy is a realistic target.

- Cue phrases provide *simple* methods for marking relations. They are mostly single words or stock phrases that we might call **idiom chunks**; it is often possible to suggest (following Zipf) that they have become simplified because they correspond to constructs which are in continual use when we construct text. The evolution of a cue phrase is sometimes implicit in its etymology: for instance, it is possible to imagine how words like *anyway* or *because* have evolved from more complex formulas and phrases.

- There is substantial empirical evidence that cue phrases are used by readers to facilitate the construction of a mental representation of text. (See for example Segal *et al* [26], for an account of the role of the phrases *and*, *then*, *so*, *because* and *but*.) Again, this suggests that a link can be drawn between cue phrases and cognitive representations of text.

6 A Data-driven Methodology for Determining a Set of Relations

In the previous section, it was argued that linguistic devices (in particular, cue phrases) can be taken as evidence for relations, provided thse are conceived as constructs which people actually use when creating and interpreting text. This section reports on our initial investigations based on this argument. First, a suitable definition is sought for 'cue phrases'; then, a corpus of cue phrases is obtained; finally, the corpus is analysed and taxonomised. The resulting taxonomy seems promising as the foundation for a taxonomy of relations.

6.1 Firming Up the Notion of 'Cue Phrase': A Test for Relational Phrases

The first task is to formulate a precise definition for the class of phrases which is under investigation. Some definitions of the category of cue phrases exist already, but these are often expressed in theory-laden terms, the only independent characterisation being lists of examples: this is the case, for instance in Grosz and Sidner [5], and in Cohen [2]. Cohen's list of examples is more extensive than Grosz and Sidner's, but neither study claims to have a complete list.

In an attempt to come up with a precise yet theory-neutral definition, we have developed a test which picks out a certain group of phrases as they occur in natural discourse. In order to avoid any terminological confusion, we refer to the class of phrases which pass this test as the class of **relational phrases**. Given any candidate phrase situated in a passage of text, an analyst can decide whether it is a relational phrase by following the procedure outlined below.

1. Isolate the phrase and its **host clause**. The host clause is the clause with which the phrase is immediately associated syntactically; for instance, if the passage of text to be examined is

   (2) ...John and Bill were squabbling. John fell *because* Bill pushed him: that was how it all started ...

   then the isolated phrase and clause would be

   (3) *because* Bill pushed him.

2. Substitute any anaphoric or cataphoric terms in the resulting text with their antecedents, and include any elided items. For the above clause, this would give:

   (4) *because* Bill pushed John.

3. If the candidate phrase is indeed a relational phrase, the resulting text should appear **incomplete**. An incomplete text is one where one or more extra clauses are needed in order for a coherent message to be framed.

(5) Because Bill pushed John

is incomplete in this sense: it requires at least one other clause in order to make a self-contained discourse. Even the fact that it could appear by itself on a scrap of paper (say as an answer to a question) does not make it complete; the question is essential context if it is to be understood.

Note that it is only additional clausal material which is to be removed in the test. Any additional contextual information necessary for the comprehension of the clause (for instance, knowledge of the referents of definite referring expressions like *John* and *Bill*) can be assumed to be present.

4. If the candidate phrase is removed, the remaining clause should *no longer* be incomplete. In other words, it should form a coherent mini-discourse on its own. Thus the text

   (6) Bill pushed John.

   can stand by itself, without the need for extra clausal material.[4]

This test is designed to give a reasonably objective way to pick out a set of phrases to act as the object for further study, without relying too much on any particular theory of syntax. There are of course many cases where it is hard to apply the test. For instance, some phrases involve anaphora which is not easily substitutable: the phrase *for <u>this</u> reason* requires a substitution like *for the reason that* ..., which is very ungainly. In such cases, it is preferable to leave the anaphoric expression in the phrase. Indeed, as Halliday and Hasan [6] point out, many *bona fide* relational phrases clearly derive from phrases involving anaphora (consider words like *therefore* and *thereby*); so this analysis is justified. Another problem in the application of the test is to decide whether the required context for a mini-discourse is linguistic or non-linguistic. For instance, consider this discourse:

(7) *But* you can't just leave us here!

It is possible to imagine this discourse with no previous utterances at all. All the same, it needs to be interpreted as a reaction to a previously existing propositional attitude (in this case, perhaps an intention to leave), and this is arguably a kind of discourse context.

6.2 Gathering a Corpus of Relational Phrases

Using the test, we gathered a corpus of relational phrases. The source texts were all passages taken from academic articles and books. Note that different registers of text might contain

[4] A final point: phrases which refer to places in the original text (such as *in the next section, as already mentioned*) will pass the test as it stands—but this is only because their referents have been expressly removed through the operation of the test itself. We will not consider such phrases here, though they are of interest in their own right.

different sets of relational phrases: for instance, phrases like *just then*, *whereupon* and *sure enough* occur in narrative discourse but are relatively rare in academic articles. We have concentrated for the moment on academic texts in order to limit the scope of the investigation.

226 pages of text were analysed, from twelve different authors. This yielded a corpus of around 200 phrases: see Appendix A for a list of these. There was found to be extensive use of a core of phrases across all the authors: for instance, *and*, *since*, *if*, and *but* were used by all twelve; *on the other hand*, *however*, and *also* were used by eleven; and *then*, *for example*, *because*, *when*, and *although* were used by ten.[5]

The following sections report the results of analyses carried out on the corpus.

6.2.1 The Syntactic Diversity of Relational Phrases

Relational phrases fall into four syntactic classes (as defined by Quirk *et al* [18]):

- **Coordinators:** these always appear in between the clauses they link; the clauses can be in separate sentences or in the same sentence. If in the same sentence, no punctuation is required in addition to the coordinator; and if combined in a sequence with other cue phrases, coordinators always appear leftmost in the sequence. For example:

  (8) a. An object may move, *but* it remains the same object...

  b. A general rule is needed to prevent comparative constructions. *Or* some rule is needed that will say: 'if a word cannot ...

- **Subordinators:** these introduce subordinate clauses in complex sentences. The subordinate clause can be on the left or the right of the main clause, but the subordinator is always on the left of the subordinate clause. For example:

  (9) a. *Although* it is common sense that labels are related, this is a difficult idea to explicate.

  b. One further illocution should be considered *before* we discuss some variants.

- **Conjunct adverbs:** these modify whole clauses, and can appear at different points within them, although there is often a default position for particular phrases. There are also syntactic constraints on exactly which positions conjunct adverbs can occupy: at the beginning of a clause, between subject and verb, between any auxiliary verbs, between auxiliary verb and main verb, after a copula if there is one, before a sentential complement if there is one. For example:

[5] Of course, it might be objected that *if ... [then]* is really the relational phrase, rather than simply *if*; our test for relational phrases does not capture this as it stands, but we intend to return to this matter in subsequent work.

(10) a. The parallel between permissibility and possibility has been exploited by many linugists. There are, *however*, two important distinctions between them ...

b. We will select only those hypotheses we deem relevant. *As a consequence*, our discussion differs from the usual views ...

- **Phrases which take sentential complements:** these often introduce a particular intentional stance with respect to the content of the clause they introduce. For example:

(11) a. An act that is physically impossible cannot occur. (...) *It follows that* the language used ...is often straightforward.

b. *It may seem that* we are making too much of orientation; but characteristic orientation is not an idiosyncrasy.

6.2.2 The Space of Relational Phrases

One finding that emerged from a study of the corpus was that the relational phrases fell into two groups. It was possible to envisage a compositional semantics for some phrases: for instance, the semantics of the phrases *many years later*, *a few years later*, and *twenty-five years later* can be seen to depend on the semantics of the constituent phrases *many years*, *a few years*, and *twenty-five years*; and these in turn depend on the semantics of the determiner phrases *many*, *a few*, and *twenty-five*.

Other phrases in the corpus, which we might refer to as simple phrases, are impossible to break down in this way. This might be because they are single words, or alternatively because they are idiom chunks, which are defined precisely as multi-word phrases whose semantics is not compositional. Examples of idiom chunks in the corpus include *on the other hand* (in contrast with the ungrammatical *off the other hand*), *after all* (in contrast with *before all*), and *given that* (in contrast with *taken that*).

A great many phrases seem to be *partly* compositional—for instance, the meanings of *on the one hand* and *on the other hand* can be thought to hinge on the meanings of *one* and *other*, but not on the meaning of *hand*: the phrase *on the other foot* is nonsense as a marker of a relation. There are many other phrases of a similar 'semi-compositional' status; but there seem to be no hard-and-fast rules for working out how such phrases are formed, and it is easiest at the outset simply to treat them as unanalysed atomic elements.

The existence of compositional relational phrases has an important consequence: it makes the class of relational phrases infinite in size. Phrases like *a very very ...very many years later* are technically members of the class, even though in practice they will never occur. This means that in order to describe the class, it is necessary to lay down rules for how compositional relational phrases can be constructed. These rules will be syntactic in nature.

For instance, the following two rules are helpful in expanding the charted space of relational phrases:

- There is a class of words which modify all subordinators and only subordinators; these words are *even*, *just*, *except*, *only* and *especially.* Thus we can construct relational phrases like *only where*, *except before*, and *just on the grounds that.* There are exceptions to this rule (for instance, **except in case*), but it still provides a useful generalisation.

- Temporal phrases can also be modified in a systematic way. The conjunct adverbs *earlier*, *afterwards* and *later*, as well as the phrases *before* and *after* (which can be conjunct adverbs or subordinators), can all be modified by any expression denoting a length of time; for instance *three days after*, *a minute earlier*, and *some time before.* The modifiers always precede the head phrases.

The general syntactic concepts of **head** and **modifier** can be used to analyse any phrase, regardless of its syntactic category. Compositional relational phrases can typically stand alone without modifiers—for instance, *later* and *after* by themselves are still relational phrases. In what follows, modifiers have been stripped wherever possible.

6.2.3 Taxomonising Relational Phrases According to Function

Thus far, we have used a simple test for detecting relational phrases in text, and on the basis of this we have gathered a corpus of relational phrases. The phrases have been classified according to their syntactic properties, so that an idea can be obtained of the complete space of phrases. But since we are principally interested in relational phrases as evidence for 'psychologically real' discourse structuring strategies, a classification of phrases according to their function in discourse is our central objective.

In keeping with the data-driven methodology adopted thus far, the classification will be made by means of a simple linguistic test, rather than by making theoretical claims about the semantics or pragmatics of the phrases in the corpus. The test is to do with **substitutability**. Very broadly, if two phrases are inter-substitutable in a passage of discourse then they should be classified in the same category. If one phrase can always be substituted for another, but not *vice versa*, then the latter phrase should be classified in a category subordinate to that of the former phrase. In this way a taxonomy of synonyms and hyponyms can be constructed.

The test for substitutability must first be slightly elaborated, however. We are not interested in whether two phrases can take the same *grammatical* position in a clause; rather, we are interested in whether they have the same function in signalling discourse relations between the clause and other units. For instance, a conjunctive adverb like *nevertheless* might have the same function as a coordinator like *but*, but the latter can only appear at the beginning of a clause, so straightforward substitution will not always be possible. In view of this, we should allow candidate phrases to be substituted in the clause in a different position, if necessary.

The substitutability test can be expressed thus:

1. Take any relational phrase from the corpus and place it in context—that is, in a host clause, along with any other clauses needed to form a complete mini-discourse.
2. Remove the relational phrase, and insert a different phrase from the corpus (the **candidate phrase**) into the same clause, at any appropriate position.
3. If the resulting discourse is consistent with the sense of the original discourse—regardless of which context chosen in step 1— then the candidate phrase is **substitutable** for the original phrase.

Note that substitutability relations can be mutual. For instance, the synonyms *later* and *afterwards* can be substituted for each other:

(12) John went to bed. *Later/Afterwards* he remembered he had not put out the cat.

But sometimes the relationship is not bidirectional, if one phrase is more general than the other. For instance, *then* is substitutable for *later*, as in (13); but *then* can also mark presentational sequences, whereas *later* cannot, as can be seen in (14).

(13) John went to bed. *Then* he remembered he had not put out the cat.

(14) John is not right for the job. For a start he's not punctual—*then/*afterwards* there are his manners to consider.

7 A Taxonomy of Relational Phrases

The substitutability test can be used to construct a hierarchical taxonomy of relational phrases. Ideally, such a taxonomy should have certain properties:

- All the relational phrases in a language should figure in it.
- All the substitutability relationships between phrases should be explicit in it.
- No phrase should appear in more than one category unless it has two separate meanings.

We have carried out a classification of the corpus of relational phrases according to these requirements. A portion of the taxonomy, dealing with phrases which signal position in some sequence, is given in Figure 2: the rest appears in Appendix B.

Some examples of the substitutability relationships in Figure 2 are given below.

- *Moreover* can be substituted for *above all*; note that the reverse is not always true.
- In turn, *next* can be substituted for *moreover*; and *next* is substitutable by *n-thly*, where *n* is some small number.

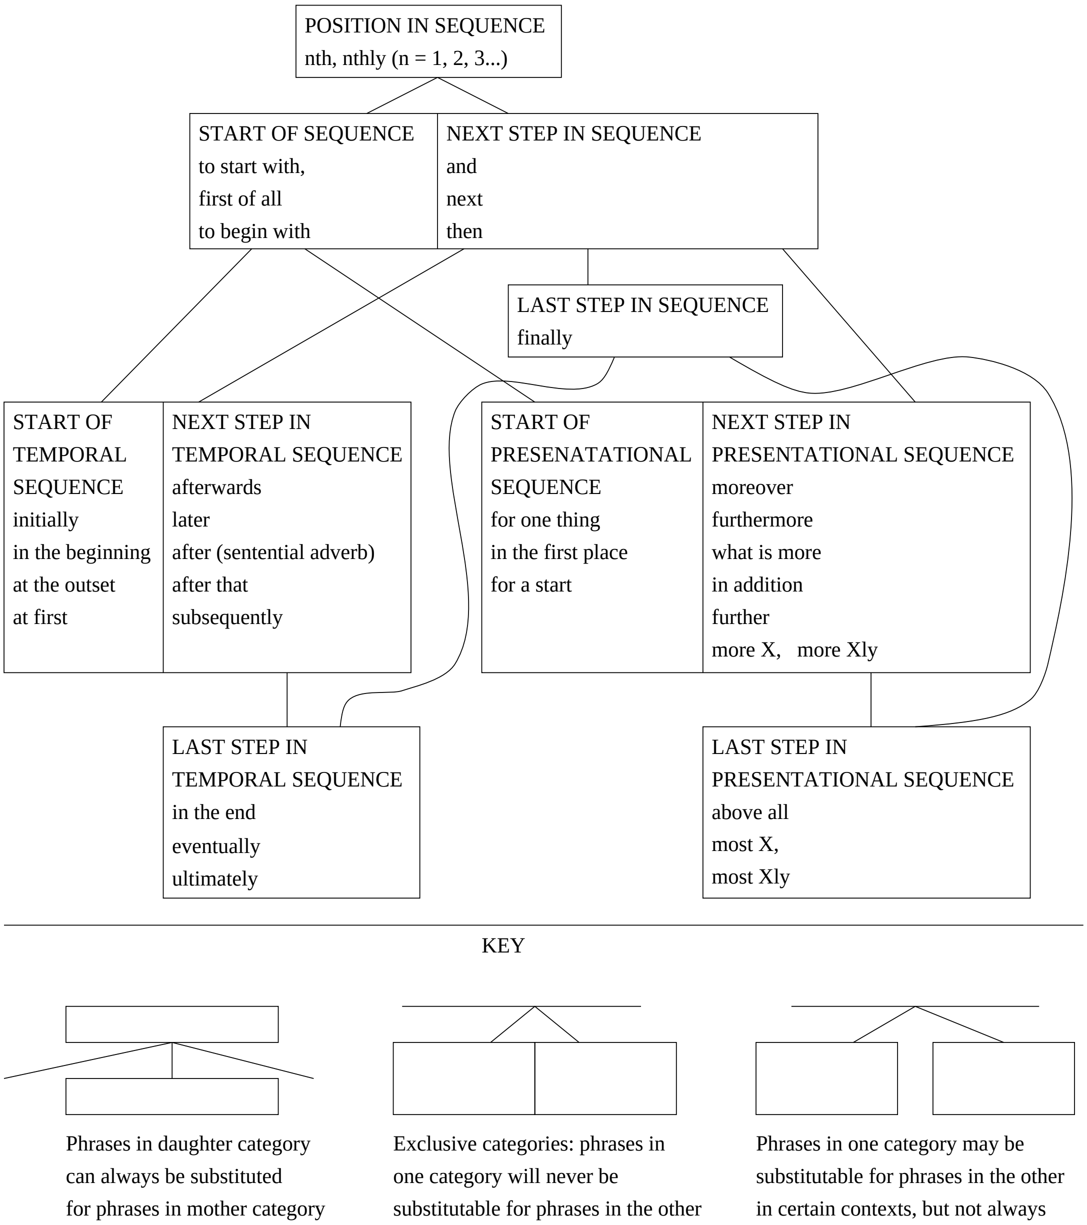

Figure 2: Sequences

- *Above all* is also substitutable by *finally*.

- Note also which substitutions are never possible: for instance, *first of all* and *next* (though these phrases are both substitutable by *n-thly*); or *in the end* and *above all* (though both phrases are substitutable by *next*).

The taxonomy is not yet perfect, and does not yet completely satisfy the criteria set out at the beginning of this section. But at least a reasonably clear method exists for querying and improving it: questions about the inclusion or placing of a given phrase will be decided on the evidence of concrete linguistic examples.

Because the taxonomy is a 'pre-theoretical' one, and makes no reference to theories of discourse structure, it can serve as a useful test-bed for such theories. Do they have a way of accounting for the pattern of phrases it exhibits? An answer to this question requires more research—but a few preliminary observations can be made here.

Firstly, it should be noted that the categories of cue phrases can often be mapped roughly back onto relations in RST's original set. Relations like SEQUENCE, CONTRAST, CIRCUMSTANCE, CAUSE and RESULT all find natural correlates in the taxonomy.

At the same time, there are significant differences. These are sometimes to do with granularity: for instance, cue phrases marking a sequential relation do not form a single group, but are themselves taxonomised. Again, the distinction between VOLITIONAL-RESULT and NON-VOLITIONAL-RESULT finds no correlate in the taxonomy. On the other hand, some RST relations do not appear to have any relational phrases associated with them. EVALUATION, for instance, seems to fall into this class. And ELABORATION, while a prominent relation in many RST analyses, seems to have no single phrase to mark it—a particularly surprising finding.

The taxonomy also yields some interesting observations about Sanders *et al*'s parameterisation of the space of relations. For instance, the concept of 'negative polarity' is a useful one: relations involving a lack of similarity (marked by phrases like *whereas*) and those involving a lack of causality (marked by phrases such as *notwithstanding that*) can all be marked by the phrases *though* and *although*—what is more, these phrases are not sensitive to the order of the spans being related. The class of relations marked by these latter phrases (see figure 6) seems to correspond to Sanders *et al*'s 'negative polarity' relations. (However, note that *but* also marks relations signalling a lack of causality or a lack of similarity; yet as Elhadad and McKeown [3] point out, and the taxonomy also shows, *but* and *although* are not intersubstitutable.)

There is also a certain amount of evidence for Sanders *et al*'s **source of coherence** parameter (which, it will be recalled, takes a value of 'semantic' or 'pragmatic'). *As a result* always indicates a semantic relation

(15) John felt sick. *As a result,* he stayed in bed.

(16) * John stayed in bed. *As a result,* he felt sick.

while *it follows that* always indicates a pragmatic relation:

(17) John's been in bed all day. *It follows that* he feels sick.

(18) * John feels sick. *It follows that* he's been in bed all day.

However, for *negative polarity* relations, there are no phrases of this kind which make the semantic/pragmatic distinction explicit. This may suggest that readers and writers are not normally so concerned with the distinction in this case, and explain why evidence for this parameter was harder to find.

As these observations show, the taxonomy can be a useful reference point when deciding on a set of relations.

8 Conclusion

This paper has presented a new methodology for motivating a set of discourse relations. Relations are treated as psychological constructs, operative in the processes of reading and writing; and evidence for these constructs is sought in a study of the linguistic means available or marking them explicitly in text. The methodology is based around two simple tests: one is for picking out what we have termed **relational phrases** in discourse, the other is for deciding on the substitutability relationships between the phrases. On the basis of these tests, a taxonomy of relational phrases can be built up for a given language, which can be used as the basis for a classification of discourse relations.

Treating relations as psychological constructs gives them an explanatory role in a theory of discourse structure: they form part of an account of why discourse is structured the way it is. This avoids the problems of arbitrariness which beset many current classifications of relations. At the same time, linking relations to relational phrases allows a fine-grained taxonomy to be constructed—one which can be justified systematically at every level by appealing to concrete linguistic evidence.

The taxonomy as it stands is of linguistic data rather than of relations proper: no relation definitions have yet been framed. This is the next task: to define the categories which the phrases serve to pick out. Whether this can be done in terms of any of the existing theories of discourse structure is a question which we intend to address in future work.

References

[1] J L Austin. *How to Do Things with Words.* Oxford: Clarendon Press, 1962.

[2] R Cohen. A computational theory of the function of clue words in argument understanding. In *Proceedings of the 10th International Conference on Computational Linguistics*, pages 251–258, 1984.

[3] M Elhadad and K R McKeown. Generating connectives. In *COLING-90*, pages 97–101, 1990.

[4] J Grimes. *The Thread of Discourse.* The Hague: Mouton, 1975.

[5] B J Grosz and C L Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, pages 175–203, 1986.

[6] M Halliday and R Hasan. *Cohesion in English.* Longman, 1976.

[7] J A Hobbs. On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Stanford University, 1985.

[8] E H Hovy. Planning coherent multisentential text. In *Proceedings of the 26th ACL Conference*, pages 163–169, Buffalo, 1988.

[9] E H Hovy. Parsimonious and profligate approaches to the question of discourse structure relations. In *Proceedings of the 5th International Workshop on Natural Language Generation*, Pittsburgh, 1990.

[10] E H Hovy. Unresolved issues in paragraph planning. In R Dale, C Mellish, and M Zock, editors, *Current Research in Natural Language Generation*, pages 17–45. Academic Press, 1990.

[11] E H Hovy, J Lavid, E Maier, V Mittal, and C Paris. Employing knowledge resources in a new text planner architecture. In *Proceedings of the Sixth International Workshop on Natural Language Generation*, pages 57–72, Trento, Italy, 1992.

[12] E Maier and E H Hovy. A metafunctionally motivated taxonomy for discourse structure relations. In *Proceedings of the 3rd European Workshop on Natural Language Generation*, pages 38–45, Judenstein, Austria, 1989.

[13] W C Mann, C M Matthiessen, and S A Thompson. Rhetorical structure theory and text analysis. Technical Report RR/89/242, Information Sciences Institute, Marina del Rey, CA, 1989.

[14] W C Mann and S A Thompson. Rhetorical structure theory: A theory of text organisation. Technical Report RR/87/190, Information Sciences Institute, Marina del Rey, CA, 1987.

[15] J R Martin. Conjunction: The logic of English text. In J S Petöfi and E Sözer, editors, *Micro and Macroconnexity of Texts*, pages 1–72. Helmut Buske Verlag, Hamburg, 1983.

[16] J D Moore and C L Paris. Planning text for advisory dialogues. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, pages 203–211, 1989.

[17] L Polanyi. The linguistic discourse model: Towards a formal theory of discourse structure. Technical Report 6409, BBN Laboratories Inc, Cambridge, MA, 1986.

[18] R Quirk, S Greenbaum, G Leech, and J Svartik. *A grammar of contemporary English.* Longmans, London, 1972.

[19] D Rösner and M Stede. Customising RST for the automatic production of technical manuals. In *Aspects of Automated Natural Language Generation*, pages 199–214. Springer-Verlag, 1992.

[20] D Rösner and M Stede. TECHDOC: A system for the automatic production of multilingual technical documents. In *Proceedings of KONVENS '92.* Springer-Verlag, 1992.

[21] T J M Sanders, W P M Spooren, and L G M Noordman. Towards a taxonomy of coherence relations. *Discourse Processes*, 15:1–35, 1992.

[22] T J M Sanders, W P M Spooren, and L G M Noordman. Coherence relations in a cognitive theory of discourse representation. *Cognitive Linguistics (to appear)*, 1993.

[23] D Schiffrin. *Discourse Markers.* Cambridge University Press, 1987.

[24] D R Scott and C S de Souza. Getting the message across in RST-based text generation. In R Dale, C Mellish, and M Zock, editors, *Current Research in Natural Language Generation.* Academic Press, 1990.

[25] J R Searle. *Speech Acts.* Cambridge: Cambridge University Press, 1969.

[26] E M Segal, J F Duchan, and P J Scott. The role of interclausal connectives in narrative structuring: Evidence from adults' interpretation of simple stories. *Discourse Processes*, 14:27–54, 1991.

[27] K Vander Linden. The expression of local rhetorical relations in instructional text. In *Proceedings of the 30th annual meeting of the Association for Computational Linguistics*, pages 318–320, Newark, Delaware, 1992.

[28] K Vander Linden, S Cumming, and J Martin. Using system networks to build rhetorical structures. In *Aspects of Automated Natural Language Generation*, pages 183–198. Springer-Verlag, 1992.

[29] G K Zipf. *The Psychobiology of Language: An Introduction to Dynamic Philology.* G. Routledge and Sons, 1936.

A The corpus of relational phrases

Phrase	Category
above all	Conj-adverb
actually	Conj-adverb
after	Subordinator
after all	Conj-adverb
afterwards	Conj-adverb
all in all	Conj-adverb
also	Conj-adverb
although	Subordinator
and	Coordinator
anyway	Conj-adverb
as a consequence	Conj-adverb
as a result	Conj-adverb
as soon as	Subordinator
at any rate	Conj-adverb
at first blush	Conj-adverb
at first view	Conj-adverb
at the outset	Conj-adverb
because	Subordinator
before	Conj-adverb
by comparison	Conj-adverb
by the same token	Conj-adverb
certainly	Conj-adverb
consequently	Conj-adverb
correspondingly	Conj-adverb
despite the fact that	Subordinator
either	Conj-adverb
equally	Conj-adverb
even	Conj-adverb
even then	Conj-adverb
every time	Subordinator
except insofar as	Subordinator
first	Conj-adverb
firstly	Conj-adverb
for a start	Conj-adverb
for instance	Conj-adverb
for the simple reason	Subordinator
further	Conj-adverb

Phrase	Category
accordingly	Conj-adverb
admittedly	Conj-adverb
after	Conj-adverb
after that	Conj-adverb
again	Conj-adverb
all the same	Conj-adverb
alternatively	Conj-adverb
always assuming that	Subordinator
and/or	Coordinator
as	Subordinator
as a corollary	Conj-adverb
as long as	Subordinator
as well	Conj-adverb
at first	Conj-adverb
at first sight	Conj-adverb
at the moment when	Subordinator
at the same time	Conj-adverb
before	Subordinator
but	Coordinator
by contrast	Conj-adverb
by the way	Conj-adverb
clearly	Conj-adverb
conversely	Conj-adverb
despite that	Conj-adverb
earlier	Conj-adverb
else	Coordinator
essentially, then	Conj-adverb
even so	Conj-adverb
eventually	Conj-adverb
except	Conj-adverb
finally	Conj-adverb
first of all	Conj-adverb
for	Subordinator
for example	Conj-adverb
for one thing	Conj-adverb
for this reason	Conj-adverb
furthermore	Conj-adverb

Phrase	Category
given that	Subordinator
however	Conj-adverb
if ever	Subordinator
if only	Subordinator
in a different vein	Conj-adverb
in addition	Conj-adverb
in case	Subordinator
in contrast	Conj-adverb
initially	Conj-adverb
in particular	Conj-adverb
in spite of that	Conj-adverb
in that case	Conj-adverb
in the case of X	Conj-adverb
in the first place	Conj-adverb
in this way	Conj-adverb
inasmuch as	Subordinator
indeed	Conj-adverb
it follows that	Phr.w/scomp
it might seem that	Phr.w/scomp
last	Conj-adverb
later	Conj-adverb
likewise	Conj-adverb
merely	Conj-adverb
more Xly	Conj-adverb
most Xly	Conj-adverb
much sooner	Conj-adverb
neither is it the case	Phr.w/scomp
next	Conj-adverb
nonetheless	Conj-adverb
not because	Conj-adverb
not that	Conj-adverb
notwithstanding that	Subordinator
now	Subordinator
obviously	Conj-adverb
on condition that	Subordinator
on one side	Conj-adverb
on the contrary	Conj-adverb
on the one hand	Conj-adverb

Phrase	Category
hence	Conj-adverb
if	Subordinator
if not	Conj-adverb
if so	Conj-adverb
in actual fact	Conj-adverb
in any case	Conj-adverb
in conclusion	Conj-adverb
in fact	Conj-adverb
in other words	Conj-adverb
in short	Conj-adverb
in sum	Conj-adverb
in the beginning	Conj-adverb
in the end	Conj-adverb
in the meantime	Conj-adverb
in turn	Conj-adverb
incidentally	Conj-adverb
instead	Conj-adverb
it might appear that	Phr.w/scomp
just as	Subordinator
lastly	Conj-adverb
let us assume	Phr.w/scomp
meanwhile	Conj-adverb
merely because	Subordinator
moreover	Conj-adverb
much later	Conj-adverb
naturally	Conj-adverb
nevertheless	Conj-adverb
no doubt	Conj-adverb
not	Conj-adverb
not only	Conj-adverb
notably	Conj-adverb
notwithstanding that,	Conj-adverb
now that	Subordinator
of course	Conj-adverb
on one hand	Conj-adverb
on the assumption that	Subordinator
on the grounds that	Subordinator
on the one side	Conj-adverb

Phrase	Category
on the other hand	Conj-adverb
once	Subordinator
once more	Conj-adverb
or else	Coordinator
overall	Conj-adverb
presumably because	Subordinator
provided that	Subordinator
put another way	Conj-adverb
reciprocally	Conj-adverb
second	Conj-adverb
similarly	Conj-adverb
simultaneously	Conj-adverb
so	Subordinator
specifically	Conj-adverb
subsequently	Conj-adverb
summarising	Conj-adverb
suppose	Phr.w/scomp
supposing that	Subordinator
surely	Conj-adverb
that is to say	Conj-adverb
the more often	Subordinator
then again	Conj-adverb
thereby	Conj-adverb
third	Conj-adverb
this time	Conj-adverb
thus	Conj-adverb
to begin with	Conj-adverb
to start with	Conj-adverb
to summarise	Conj-adverb
to the degree that	Subordinator
too	Conj-adverb
ultimately	Conj-adverb
unless	Subordinator
we might say	Phr.w/scomp
when	Subordinator
where	Subordinator
wherein	Subordinator
while	Subordinator

Phrase	Category
on the other side	Conj-adverb
once again	Conj-adverb
or	Coordinator
otherwise	Conj-adverb
plainly	Conj-adverb
previously	Conj-adverb
providing that	Subordinator
rather	Conj-adverb
regardless of that	Conj-adverb
secondly	Conj-adverb
simply because	Subordinator
since	Subordinator
so that	Subordinator
still	Conj-adverb
such that	Subordinator
summing up	Conj-adverb
suppose that	Phr.w/scomp
sure enough	Conj-adverb
that is	Conj-adverb
the fact is that	Phr.w/scomp
then	Coordinator
thereafter	Conj-adverb
therefore	Conj-adverb
thirdly	Conj-adverb
though	Subordinator
to be sure	Conj-adverb
to conclude	Conj-adverb
to sum up	Conj-adverb
to take an example	Conj-adverb
to the extent that	Subordinator
true	Conj-adverb
undoubtedly	Conj-adverb
until	Subordinator
what is more	Conj-adverb
whenever	Subordinator
whereas	Conj-adverb
wherever	Subordinator
yet	Coordinator

B A preliminary taxonomy of the corpus of relational phrases according to function

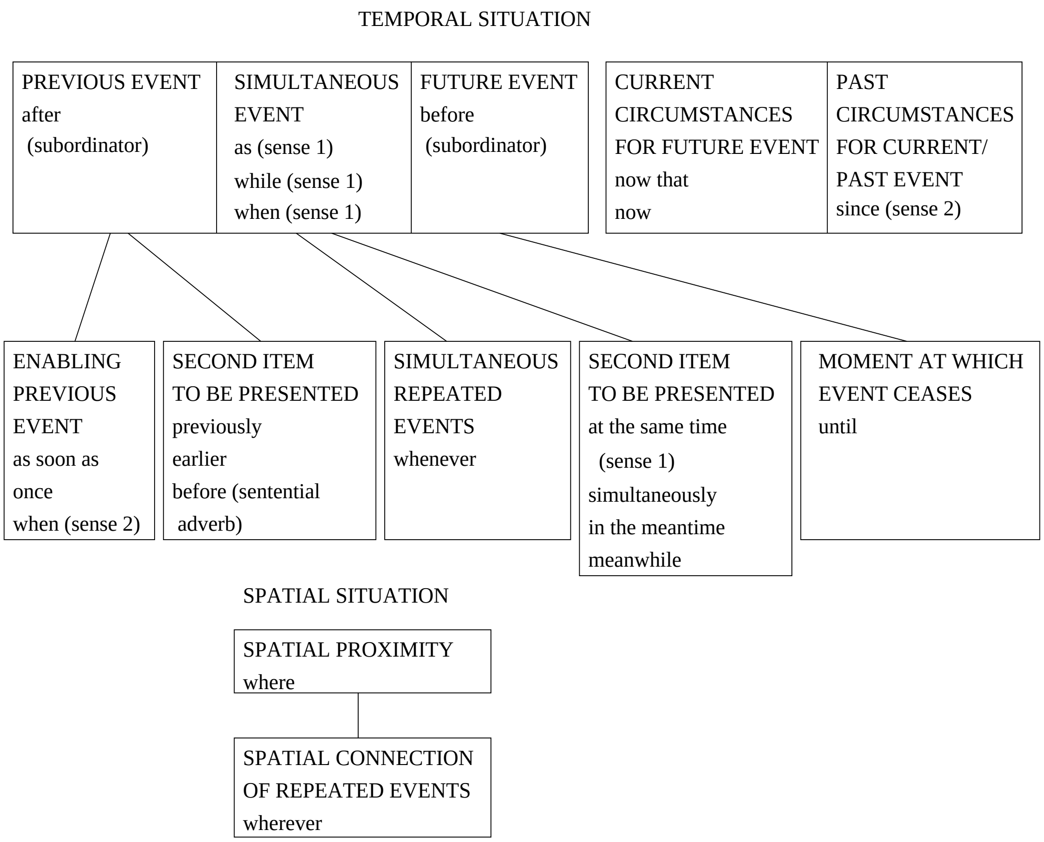

Figure 3: Situating an Event

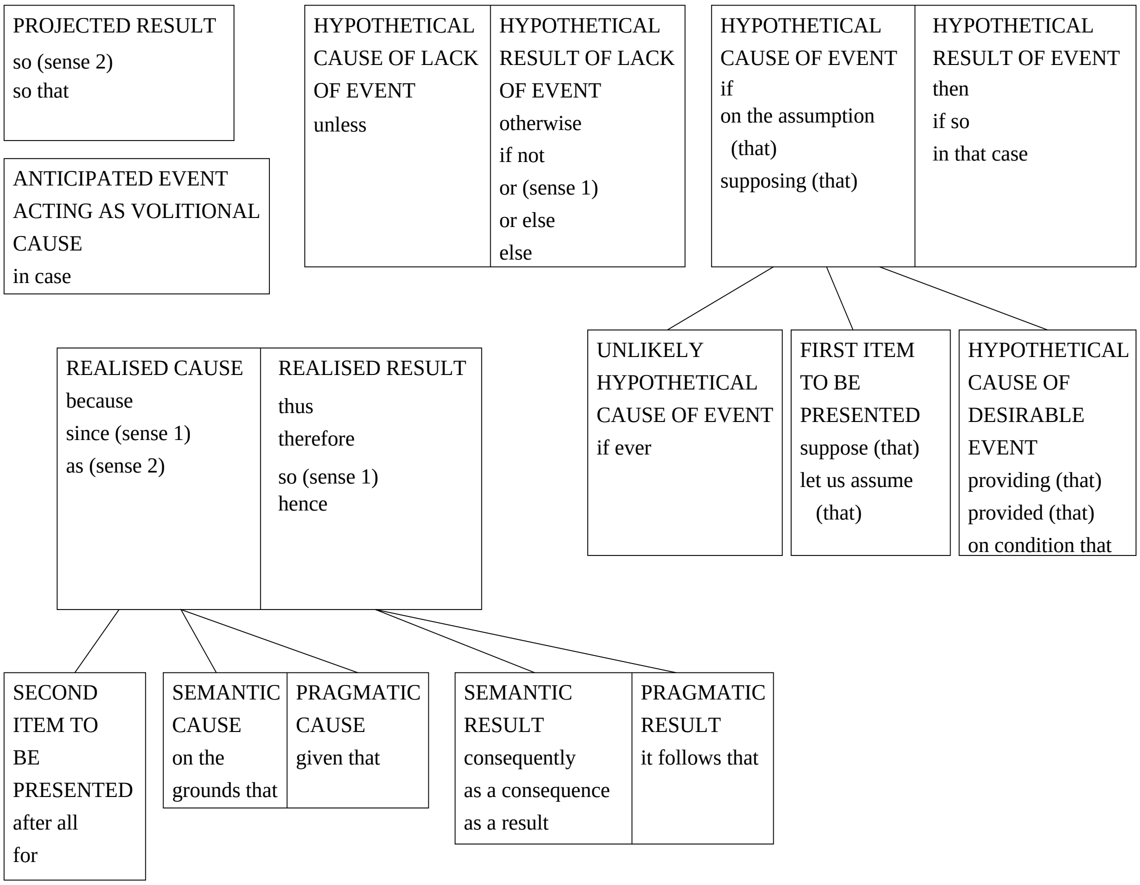

Figure 4: Causal/Purpose Relations

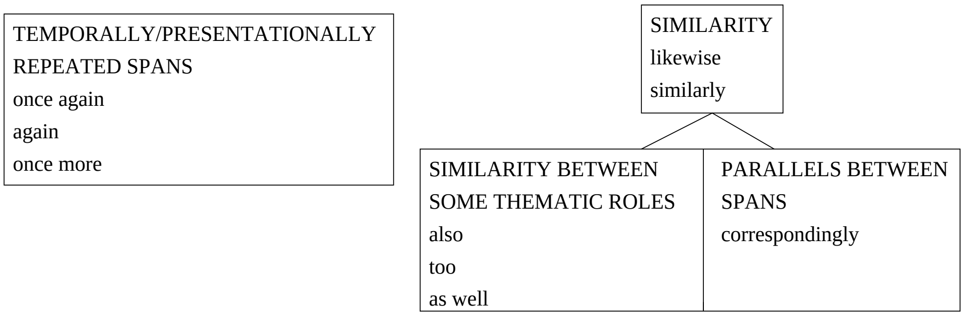

Figure 5: Similarity Relations

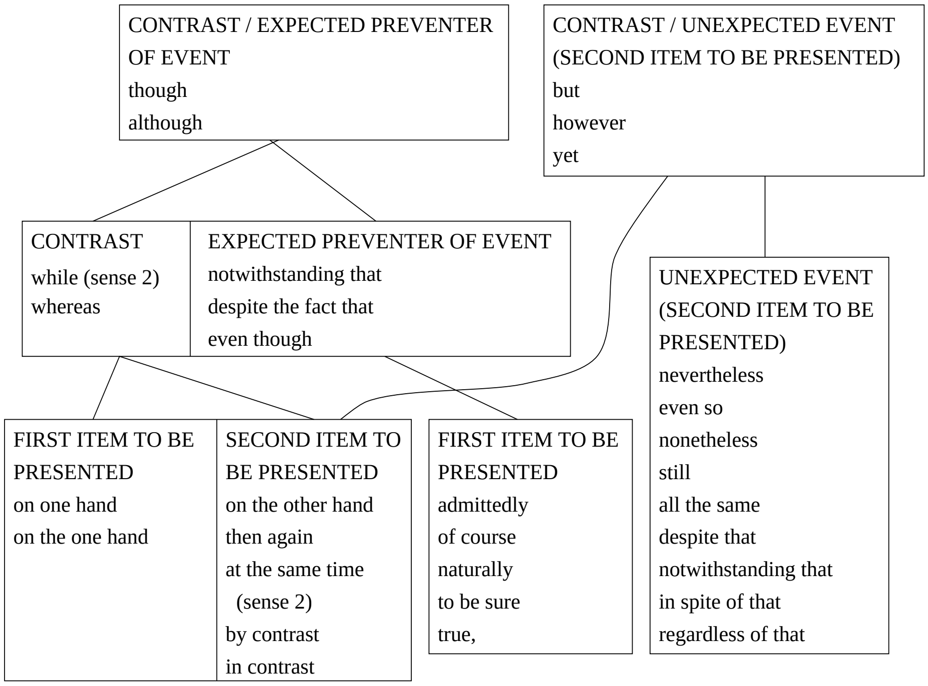

Figure 6: 'Negative Polarity' Relations

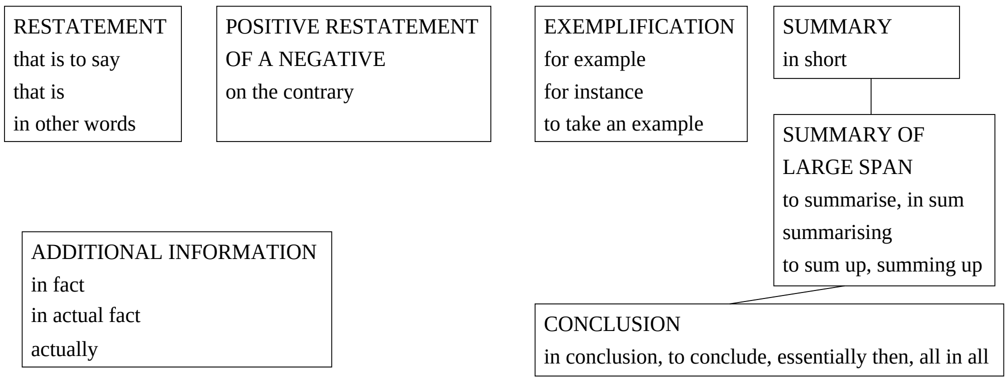

Figure 7: Clarifications